Nature's Cognitive Architecture:

From Neuroscience To Artificial General Intelligence (AGI)

D. R. Dison

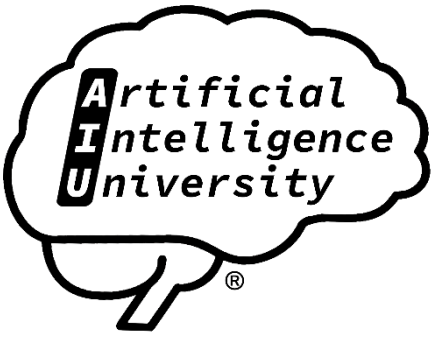

Nature’s Cognitive Architecture:

From Neuroscience
To Artificial General Intelligence (AGI)

D. R. Dison

  https://www.linkedin.com/in/d-r-dison

TABLE OF CONTENTS:

  https://www.linkedin.com/in/d-r-dison

FRONT MATTER

NATURE'S COGNITIVE ARCHITECTURE:

From Neuroscience to Artificial General Intelligence (AGI)

A complete treatment of the brain's regions, interfaces, molecules, dynamics, drives, and plasticity, drawn from the Fifteen Encapsulated Layers of Abstraction of the Human Nervous System and Endocrine System, and carried through to a buildable specification by the SPARCD method: Specification, Pseudocode, Architecture, Refinement, Completion, and Demonstration.

D. R. Dison

Founder, Artificial Intelligence University (AIU)

Open Researcher and Contributor ID: ORCID: 0009-0001-9246-5758

https://www.linkedin.com/in/d-r-dison

Version 20

This volume is both a work of popular science, written to be read cover to cover by any curious person, and a formal engineering specification, written to be handed to a capable implementation engine and built.

The main text (Chapters 1 through 13) is the science.

The six appendices are the specification.

A reader may enjoy the first without the second; an engineer needs both.

A NOTE ON NOTATION

Throughout this document, a construct that is real, functionally necessary, and relied upon in the neuroscience literature, but that has never been given an accepted proper name as a unified concept, is marked on first mention with the tag (PUC), meaning Previously Unnamed Construct.

Where such a construct is identified, a name is proposed.

The proposal of a name is not a claim of discovery.

The anatomy was already known.

What was missing was the word, and a science without a word for a thing cannot reason about that thing efficiently.

Throughout this document, a specification field or property that genuinely does not apply to a given construct is marked (NA), meaning Not Applicable, and is always accompanied by the reason for its non-applicability, so that a reader never mistakes a field that does not apply for one that was merely omitted.

HOW TO READ THIS BOOK

This volume is written for two readers at once: the working neuroscientist who needs an exhaustive connectivity reference, and the beginner who has never taken a neuroscience course.

Every technical term is defined on first use.

Every acronym is spelled out on first use.

The beginner should read Chapter 1, then read the plain-language paragraph at the head of each region, and skip the connectivity lists until curiosity demands them.

The specialist should skip Chapter 1 entirely, and may wish to begin with Chapter 5 on the rhythms and Chapter 6 on the quantitative grounding, Chapter 7 on the molecular messengers, and Chapters 8 through 11 on the dynamics, drives, and plasticity that make the map runnable, which are new to this expanded edition and add the temporal, numerical, and chemical dimensions that a purely anatomical catalogue omits.

THE SPECIFICATION (Appendices 1 through 6: the SPARCD method)

Appendix 1. Specification
What the digital brain must do, and the qualities it must have. Functional and non-functional requirements, user scenarios, constraints, assumptions, and acceptance criteria.

Appendix 2. Pseudocode (in English Readable Code, ERC)
The logic of the system in precise plain English: the tick loop, the per-construct update rule, the drive loop, and the learning rule, written so that a human or a machine can read them without a programming language.

Appendix 3. Architecture
The skeleton of the system: components, data model, the construct registry, the scheduler, the neuromodulatory bus, and how they fit together.

Appendix 4. Refinement
How the design is improved and made trustworthy: test-driven development, the validation targets, the known trade-offs, and the iterative loop that hardens it.

Appendix 5. Completion
How the system is finished and shipped: the full test strategy, documentation, deployment, monitoring, and the definition of done.

Appendix 6. Demonstration
Why the system is programming-language-agnostic and data-structure-agnostic up to this point, so that anyone can build their own at home; the planned implementations (Causalontology 2.0.0, a Python Minimum Viable Product, and a PrologAI implementation); a concrete, staged ladder of demonstrable milestones - smoke tests, sensory front-ends for text, vision, and speech, logic and reasoning trials, emotional trials, a robot body, and a mapping onto the developmental stages of a human child; and a confident, well-founded case for Artificial General Intelligence (AGI), with the route mapped from the ground already taken to the summit.

  https://www.linkedin.com/in/d-r-dison

A NOTE ON THE TWO HALVES

The thirteen chapters and the six appendices describe the same object twice.

The chapters describe it as nature built it, in the language of biology, for a reader who wants to understand.

The appendices describe it as an engineer would rebuild it, in the language of specification, for a machine that must construct it.

Everything the appendices require, the chapters have already supplied.

The specification adds no new neuroscience; it only reorganizes what the science established into the exact form an implementation engine consumes.

CHAPTER 1. THE FIFTEEN LAYERS, IN BRIEF

1.1 The Framework in One Page

The human nervous system and endocrine system span more than fifteen orders of magnitude, from the quark to human culture.

The companion volume argues that this span is organized into exactly fifteen encapsulated layers of abstraction, each a genuine level of nature with its own components, its own governing laws, and its own vocabulary, and each communicating ordinarily only with the layer directly above and the layer directly below it.

The fifteen layers are:

• Layer 1, the Subatomic Layer.

Quarks, electrons, protons, photons.

Governed by quantum mechanics.

• Layer 2, the Atomic Layer.

The chemical elements.

Governed by the periodic law.

• Layer 3, the Ionic and Simple Molecular Layer.

Ions, water, neurotransmitters.

Governed by electrochemistry.

• Layer 4, the Macromolecular Layer.

Proteins, ion channels, receptors, DNA.

Governed by structural biology.

• Layer 5, the Subcellular and Organellar Layer.

Membranes, mitochondria, vesicles, the synaptic density.

  https://www.linkedin.com/in/d-r-dison

Governed by cell biology.

• Layer 6, the Cellular Layer.

Neurons, glia, endocrine cells.

Governed by neuronal physiology.

• Layer 7, the Synaptic and Junctional Layer.

Synapses and gap junctions.

Governed by synaptic physiology.

• Layer 8, the Microcircuit Layer.

Canonical circuit motifs.

Governed by circuit neuroscience.

• Layer 9, the Brain Regions, Nuclei, and Pathways Layer.

The subject of this volume.

• Layer 10, the Neural and Neuroendocrine Systems Layer.

Distributed systems and hormonal axes.

• Layer 11, the Whole Nervous and Endocrine System Layer.

The integrated regulatory organ.

• Layer 12, the Organism and Behavior Layer.

The behaving, experiencing person.

• Layer 13, the Dyadic and Small Group Layer.

The relationship.

• Layer 14, the Community and Society Layer.

Culture and institution.

• Layer 15, the Evolutionary and Species Layer.

Phylogeny and population.

The companion volume further argues that the adjacency rule has seven documented exceptions, called Layer-Skipping Channels, in which a signal travels intact across non-adjacent layers, and that these exceptions obey an economic logic: biology pays the cost of layer-by-layer relay when a message needs an address, and refuses to pay it when the message is a broadcast.

1.2 What Layer 9 Is, and What It Is Not

Layer 9 is the layer of named anatomical structure.

  https://www.linkedin.com/in/d-r-dison

It is the layer of the map.

Its principal unit of analysis is the region: a discrete, anatomically bounded piece of nervous tissue, or a discrete endocrine gland, that performs an identifiable function by virtue of the microcircuits it contains and the connections it makes.

It is essential to be precise about what distinguishes Layer 9 from its neighbours, because these boundaries are where the framework earns its keep.

Layer 8, below, is the layer of the microcircuit motif.

A microcircuit motif is generic: lateral inhibition is lateral inhibition wherever it appears, and the same handful of motifs recur throughout the brain.

Layer 9 is what you get when you take a specific set of Layer 8 motifs, replicate them, arrange them in a particular architecture, and wire them to a particular set of inputs and outputs.

Regional identity is therefore not a matter of possessing a unique circuit.

It is a matter of connectional fingerprint: what a region receives, and to whom it sends.

This is the single most important idea in this volume, and it is worth stating in its strongest form.

A brain region is defined by its interfaces.

Layer 10, above, is the layer of the distributed system: the visual system, the motor system, the hypothalamic-pituitary-adrenal (HPA) axis.

A Layer 10 system is a chain of Layer 9 regions plus the pathways connecting them, considered as a functional whole.

The visual system is not a place.

It is an ordered sequence of places.

So Layer 9 supplies the parts and the wires; Layer 10 assembles them into systems.

1.3 The Organizing Principle of This Volume

The outline that follows is organized neuroanatomically, by classical anatomical division, because that is the organization a reviewer, a clinician, and a student all already carry in their heads.

But functional systems are held in close regard throughout: for every region, the functional system or systems to which it belongs are named explicitly, so that the reader can traverse the material either way.

For every construct at Layer 9, four things are specified.

- The inputs.

  https://www.linkedin.com/in/d-r-dison

From where, and by the name of the interface that carries them.

• The outputs.

To where, and by the name of the interface that carries them.

• The function or faculty.

What the region does, and what capacity of the whole organism would be lost if it were destroyed.

• The system membership.

Which Layer 10 system or systems this region participates in.

Chapter 3 then treats the interfaces themselves as first-class objects, because the central claim of this volume is that in a layered architecture, the interfaces are not the plumbing between the interesting parts.

The interfaces are the interesting part.

1.4 A Word on Interfaces, for the Beginner

An interface, in engineering, is the agreed boundary at which two components meet and exchange something.

A wall socket is an interface.

It has a defined shape, a defined voltage, and a defined promise: plug in a conforming device and it will receive power.

Neither the power station nor the lamp needs to know anything about the other's internals.

They need only agree about the socket.

The brain is built the same way, and this volume takes the position that the named white matter tracts of neuroanatomy are exactly such sockets.

The perforant path is an interface.

The fornix is an interface.

The corticospinal tract is an interface.

Each has a defined origin, a defined destination, a defined cargo, and a defined format.

And, as we shall see in Chapter 3, a great many real interfaces in the human brain have never been named at all, which is why this volume proposes names for them.

CHAPTER 2. THE COMPLETE OUTLINE OF LAYER 9 CONSTRUCTS

2.1 THE CEREBRAL CORTEX

The cerebral cortex is the folded sheet of grey matter covering the cerebral hemispheres, roughly two to four millimetres thick, containing on

  https://www.linkedin.com/in/d-r-dison

the order of sixteen billion neurons, and organized into six horizontal laminae.

A warning about vocabulary is required immediately, because the collision of terms has misled many students.

The six laminae of the cortex are not the fifteen layers of this framework.

They are called layers in the classical literature, and they will be called laminae throughout this volume to avoid the collision.

The laminae are internal to Layers 8 and 9.

The laminae matter here because they are the physical implementation of the cortex's interface discipline.

Lamina 4 is the input port: it is where driving thalamic afferents terminate, and it is thick in sensory cortex and nearly absent in motor cortex, which tells you at a glance which regions listen and which regions speak.

Laminae 2 and 3 are the horizontal port: they send and receive corticocortical connections with other cortical regions.

Lamina 5 is the subcortical output port: its large pyramidal neurons project to the striatum, the brainstem, the spinal cord, and the higher-order thalamus.

Lamina 6 is the thalamic feedback port: it projects back to the thalamic nucleus that supplies the region's input.

This four-port arrangement is repeated across the whole cortical sheet with remarkable constancy, and it is the reason the cortex can be described as running one algorithm on many different data streams.

2.1.1 PRIMARY VISUAL CORTEX (V1, STRIATE CORTEX, BRODMANN AREA 17)

Plain language: the first place in the cortex where seeing happens.

It does not recognize objects.

It detects edges, orientations, and motion, the raw alphabet from which vision is later assembled.

INPUTS (SUBSCRIBE)

• From the lateral geniculate nucleus of the thalamus, by way of the optic radiation (also called the geniculocalcarine tract).

This is the driving input, and it terminates in lamina 4C.

• From the pulvinar, by way of pulvinocortical projections.

A higher-order, modulatory input.

• From V2, V3, V4, and V5/MT, by way of corticocortical feedback projections terminating in laminae 1, 2, and 6.

  https://www.linkedin.com/in/d-r-dison

• From the nucleus basalis of Meynert, by way of the cholinergic basal forebrain projection.

Modulatory; gates plasticity and attention.

• From the locus coeruleus, by way of the ascending noradrenergic projection.

OUTPUTS (PUBLISH)

• To V2, V3, V4, and V5/MT, by way of corticocortical feedforward projections from laminae 2 and 3.

This is the origin of both the dorsal and ventral streams.

• To the lateral geniculate nucleus, by way of the corticogeniculate feedback projection from lamina 6.

Note the numerical asymmetry, which is one of the most striking facts in neuroanatomy: the feedback projection from V1 to the LGN outnumbers the forward projection from LGN to V1 by roughly ten to one.

The cortex tells the thalamus what to expect far more than the thalamus tells the cortex what is there.

• To the superior colliculus, by way of the corticotectal projection from lamina 5.

• To the pulvinar, by way of lamina 5 driving projections.

FUNCTION AND FACULTY

Orientation selectivity, spatial frequency analysis, binocular disparity (the basis of stereoscopic depth), direction-of-motion detection, and colour-opponent processing.

Destruction produces cortical blindness in the corresponding region of the visual field, with preservation of pupillary reflexes, since those are mediated subcortically.

The phenomenon of blindsight, in which a patient with V1 damage can guess the location of a stimulus they report not seeing, is explained by the surviving retino-collicular-pulvinar route, which bypasses V1 entirely.

SYSTEM MEMBERSHIP: The visual system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY with local ATTRACTOR (contour completion). State: activation per feature channel; gain. Rule: tracks thalamic input, scaled by attentional gain, sharpened by lateral inhibition.

DRIVES: Not Applicable, for the reason that this construct serves no drive directly; it supplies the perception on which all drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Critical-period plastic, limited in adult; weights refine with visual experience gated by acetylcholine.

  https://www.linkedin.com/in/d-r-dison

MODULATION: Acetylcholine and norepinephrine set gain.

2.1.2 SECONDARY AND HIGHER VISUAL AREAS (V2, V3, V4, V5/MT)

Plain language: the assembly line that turns edges into things.

Each stage extracts something more abstract than the last.

INPUTS (SUBSCRIBE)

• From V1, by way of feedforward corticocortical projections.

• From the pulvinar, by way of pulvinocortical projections.

The pulvinar is thought to coordinate the timing of communication among these areas.

• From each other, laterally and by feedback.

OUTPUTS (PUBLISH)

• V4 to the inferotemporal cortex, by way of the ventral stream, also called the occipitotemporal pathway or the "what" pathway.

Cargo: object identity, form, colour.

• V5/MT to the posterior parietal cortex, by way of the dorsal stream, also called the occipitoparietal pathway or the "where and how" pathway.

Cargo: motion, spatial location, and the information needed to guide action toward an object.

• To the frontal eye fields, by way of long-range corticocortical projections, for the control of eye movements.

FUNCTION AND FACULTY

V2 processes contours and figure-ground segregation, including illusory contours.

V3 processes dynamic form.

V4 is critical to colour constancy and to the representation of shape; its destruction produces achromatopsia, the loss of colour vision, in which the world appears in shades of grey.

V5/MT is critical to motion perception; its destruction produces akinetopsia, a rare and disturbing condition in which the patient perceives the world as a series of still frames, and cannot, for example, pour tea, because the liquid appears frozen.

SYSTEM MEMBERSHIP: The visual system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY with local ATTRACTOR (contour completion). State: activation per feature channel; gain. Rule: tracks thalamic input, scaled by attentional gain, sharpened by lateral inhibition.

  https://www.linkedin.com/in/d-r-dison

DRIVES: Not Applicable, for the reason that this construct serves no drive directly; it supplies the perception on which all drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Critical-period plastic, limited in adult; weights refine with visual experience gated by acetylcholine.

MODULATION: Acetylcholine and norepinephrine set gain.

2.1.3 PRIMARY AUDITORY CORTEX (A1, HESCHL'S GYRUS, BRODMANN AREAS 41 AND 42)

Plain language: where hearing becomes conscious.

Organized by pitch, like a piano keyboard laid out across the tissue.

INPUTS (SUBSCRIBE)

• From the ventral division of the medial geniculate nucleus of the thalamus, by way of the acoustic radiation (also called the thalamocortical auditory radiation).

Driving input to lamina 4.

• From the contralateral A1, by way of the corpus callosum.

• From the nucleus basalis of Meynert, cholinergic and modulatory.

OUTPUTS (PUBLISH)

• To the secondary auditory cortex and the superior temporal gyrus, by way of feedforward corticocortical projections.

• To Wernicke's area, by way of the auditory-to-language interface in the posterior superior temporal gyrus.

• To the medial geniculate nucleus, by way of corticothalamic feedback from lamina 6.

• To the inferior colliculus, by way of the descending corticocollicular projection, which modulates auditory brainstem processing according to expectation.

FUNCTION AND FACULTY

Tonotopic frequency analysis, meaning that neighbouring cells respond to neighbouring pitches; temporal pattern analysis; and sound localization by integration of interaural time and level differences computed in the brainstem.

Bilateral destruction produces cortical deafness.

SYSTEM MEMBERSHIP: The auditory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY, tonotopic. State: activation per frequency channel; gain. Rule: tracks medial geniculate input scaled by gain.

  https://www.linkedin.com/in/d-r-dison

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Adult plasticity of frequency tuning, gated by attention/acetylcholine.

MODULATION: Acetylcholine, norepinephrine.

2.1.4 PRIMARY SOMATOSENSORY CORTEX (S1, POSTCENTRAL GYRUS, BRODMANN AREAS 3, 1, AND 2)

Plain language: the map of the body's surface, laid out across the brain in the famous distorted homunculus in which the lips and hands are enormous and the trunk is tiny, because cortical territory is allocated by sensitivity, not by size.

INPUTS (SUBSCRIBE)

• From the ventral posterolateral nucleus (VPL) of the thalamus, carrying body sensation, and the ventral posteromedial nucleus (VPM), carrying face sensation, by way of the thalamocortical somatosensory radiation.

Driving input to lamina 4.

The VPL is fed by the medial lemniscus (fine touch, vibration, proprioception) and the spinothalamic tract (pain, temperature, crude touch).

• From M1, by way of dense reciprocal corticocortical connections across the central sulcus.

OUTPUTS (PUBLISH)

• To the secondary somatosensory cortex (S2) and the posterior parietal cortex, by way of feedforward corticocortical projections.

• To M1, by way of reciprocal connections.

This is the anatomical basis of sensorimotor integration: the motor cortex is continuously informed by the sensory cortex about the consequences of its own commands.

• To the striatum, by way of corticostriatal projections from lamina 5.

• To the dorsal column nuclei and the spinal dorsal horn, by way of descending corticosensory projections that gate incoming sensation at its first synapse.

This is why you cannot tickle yourself: the motor system informs the sensory system that the touch is self-generated, and the sensory system attenuates it.

FUNCTION AND FACULTY

Somatotopic representation of touch, pressure, vibration, proprioception (the sense of limb position), and the discriminative aspects of pain and temperature.

  https://www.linkedin.com/in/d-r-dison

Destruction produces contralateral loss of fine touch and proprioception, and, notably, astereognosis, the inability to recognize an object by feel alone.

SYSTEM MEMBERSHIP: The somatosensory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY, somatotopic. State: activation per body-surface channel. Rule: tracks thalamic (VPL/VPM) input; self-touch attenuated by efference copy.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Map plasticity with use and injury.

MODULATION: Acetylcholine, norepinephrine.

2.1.5 PRIMARY MOTOR CORTEX (M1, PRECENTRAL GYRUS, BRODMANN AREA 4)

Plain language: the region that issues the final cortical command to move.

It contains the largest neurons in the brain.

INPUTS (SUBSCRIBE)

• From the premotor cortex and supplementary motor area, by way of corticocortical projections.

These carry the plan.

• From S1 and the posterior parietal cortex, by way of corticocortical projections.

These carry the current state of the body.

• From the ventral lateral nucleus (VL) of the thalamus, by way of the thalamocortical motor radiation.

The VL is the gateway through which both the basal ganglia and the cerebellum deliver their influence to the cortex, and it is therefore one of the most consequential interfaces in the entire nervous system.

• From the cerebellum, indirectly, by way of the dentato-thalamo-cortical pathway terminating in VL.

• From the basal ganglia, indirectly, by way of the pallido-thalamo-cortical pathway terminating in VL and VA.

OUTPUTS (PUBLISH)

• To the spinal cord, by way of the corticospinal tract (also called the pyramidal tract), arising from the giant Betz cells of lamina 5.

  https://www.linkedin.com/in/d-r-dison

Roughly ninety percent of these fibres decussate, meaning they cross the midline, at the pyramidal decussation of the medulla, forming the lateral corticospinal tract; the remainder descend uncrossed as the anterior corticospinal tract.

• To the cranial nerve motor nuclei, by way of the corticobulbar tract, controlling the muscles of the face, jaw, tongue, and larynx.

• To the striatum, by way of corticostriatal projections.

• To the subthalamic nucleus, by way of the hyperdirect pathway, bypassing the striatum entirely.

Cargo: a fast global stop signal.

• To the pontine nuclei, by way of the corticopontine projection, which is the first leg of the great cortico-ponto-cerebellar loop.

• To the red nucleus, by way of the corticorubral projection.

FUNCTION AND FACULTY

Execution of voluntary movement; encoding of movement direction, force, and speed.

Destruction produces contralateral spastic paresis: weakness with increased tone, because the descending inhibition of spinal reflexes is lost.

SYSTEM MEMBERSHIP: The motor system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-INTEGRATOR issuing commands. State: activation per movement parameter. Rule: integrates plan from premotor/basal-ganglia/cerebellum, emits descending command.

DRIVES: Serves whichever drive currently commands action.

PLASTICITY: Strong practice-driven plasticity of the motor map, three-factor, dopamine-gated.

MODULATION: Dopamine, acetylcholine.

2.1.6 PREMOTOR CORTEX (PMC) AND SUPPLEMENTARY MOTOR AREA (SMA), BRODMANN AREA 6

Plain language: the planning offices of movement.

The premotor cortex plans movements guided by external events; the supplementary motor area plans movements generated from within.

INPUTS (SUBSCRIBE)

• PMC: from the posterior parietal cortex, by way of the parieto-frontal projections, carrying spatial and object information for reaching and grasping.

From the prefrontal cortex, carrying goals.

  https://www.linkedin.com/in/d-r-dison

From the ventral lateral thalamus, carrying cerebellar influence.

• SMA: from the basal ganglia by way of VA and VL thalamus, carrying the results of action selection.

From the prefrontal cortex.

OUTPUTS (PUBLISH)

• To M1, by way of dense corticocortical projections.

• To the spinal cord, directly, by way of a contribution to the corticospinal tract.

This is important and often forgotten: M1 is not the sole origin of the pyramidal tract.

• To the striatum and the subthalamic nucleus.

• To the contralateral hemisphere, by way of the corpus callosum, coordinating bimanual movement.

FUNCTION AND FACULTY

PMC: sensory-guided movement, motor preparation, and the mirror properties by which observing an action activates the circuits for performing it.

SMA: internally generated movement sequences, bimanual coordination, and the sense of volition.

Destruction of the SMA can produce alien hand syndrome, in which a limb performs coordinated purposeful actions the patient disowns.

SYSTEM MEMBERSHIP: The motor system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-INTEGRATOR issuing commands. State: activation per movement parameter. Rule: integrates plan from premotor/basal-ganglia/cerebellum, emits descending command.

DRIVES: Serves whichever drive currently commands action.

PLASTICITY: Strong practice-driven plasticity of the motor map, three-factor, dopamine-gated.

MODULATION: Dopamine, acetylcholine.

2.1.7 DORSOLATERAL PREFRONTAL CORTEX (DLPFC, BRODMANN AREAS 9 AND 46)

Plain language: the seat of working memory and of holding a plan in mind while distraction assails it.

If you have ever kept a phone number in your head while someone talked at you, this region did the work.

INPUTS (SUBSCRIBE)

  https://www.linkedin.com/in/d-r-dison

• From the mediodorsal nucleus (MD) of the thalamus, by way of the anterior thalamic radiation.

The MD is the thalamic gateway to the prefrontal cortex, and the DLPFC is defined anatomically, in the classical literature, as the cortex that receives MD projections.

• From the posterior parietal cortex, by way of the superior longitudinal fasciculus.

Cargo: spatial information.

• From the temporal association cortex, by way of long corticocortical association fibres.

Cargo: object identity.

• From the anterior cingulate cortex.

Cargo: conflict and error signals.

• From the ventral tegmental area, by way of the mesocortical dopamine projection.

This input is critical: dopamine in the DLPFC tunes the signal-to-noise ratio of working memory, and both too little and too much impair it, which is why the relationship between dopamine and cognitive performance is an inverted U.

OUTPUTS (PUBLISH)

• To the premotor cortex and the SMA.

Cargo: the plan, handed to the motor hierarchy for execution.

• To the caudate nucleus, by way of the corticostriatal projection into the associative territory of the striatum.

This is the entry point of the cognitive basal ganglia loop.

• To the mediodorsal thalamus, by way of corticothalamic feedback.

• To the hippocampus, by way of the uncinate fasciculus and the cingulum bundle.

Cargo: top-down control of memory retrieval.

FUNCTION AND FACULTY

Working memory, meaning the active maintenance of information over seconds in the absence of the stimulus; cognitive flexibility, meaning the ability to switch rules; abstract reasoning; and the inhibition of prepotent responses.

Destruction produces dysexecutive syndrome: perseveration (the inability to stop doing what one is doing), distractibility, and a failure of planning, with intelligence as measured by standard tests often intact.

SYSTEM MEMBERSHIP: The central executive network (Layer 10).

  https://www.linkedin.com/in/d-r-dison

SPECIFICATION BLOCK

DYNAMICS: ATTRACTOR (working memory: a point attractor holding a value across a delay). State: maintained activation pattern; gain. Rule: recurrent self-excitation sustains a representation absent the stimulus.

DRIVES: Serves the active goal by holding it online.

PLASTICITY: Slow; rule and strategy learning, dopamine-gated.

MODULATION: Dopamine, on an inverted-U; too little or too much impairs.

2.1.8 VENTROMEDIAL PREFRONTAL CORTEX (VMPFC) AND ORBITOFRONTAL CORTEX (OFC), BRODMANN AREAS 10, 11, 12, 25, 32

Plain language: the region that assigns value.

Not what a thing is, but what it is worth to you, and whether it will feel good or bad.

This is the region destroyed in Phineas Gage.

INPUTS (SUBSCRIBE)

• From all five sensory modalities, by way of convergent corticocortical projections.

The OFC is the only cortical region receiving direct input from every sense, including taste and smell, which is what allows it to compute the value of a whole multisensory experience.

• From the amygdala, by way of the ventral amygdalofugal pathway and the uncinate fasciculus.

Cargo: emotional salience.

• From the mediodorsal thalamus, medial portion.

• From the hippocampus, by way of the uncinate fasciculus.

Cargo: episodic context.

• From the ventral tegmental area, by way of the mesocortical dopamine projection.

Cargo: reward prediction error.

OUTPUTS (PUBLISH)

• To the amygdala, by way of the reciprocal amygdalofugal and uncinate projections.

Cargo: top-down regulation and extinction of fear.

This is the pathway that psychotherapy is thought to strengthen.

• To the hypothalamus and the periaqueductal gray, by way of descending projections.

  https://www.linkedin.com/in/d-r-dison

Cargo: autonomic and behavioural implementation of the valuation.

• To the ventral striatum and nucleus accumbens, by way of corticostriatal projections.

This is the entry to the limbic basal ganglia loop.

• To the DLPFC.

Cargo: the value signal on which the executive must act.

FUNCTION AND FACULTY

Valuation, reward and punishment representation, reversal learning (updating when the value of something changes), social and moral judgement, and the somatic marker function by which bodily states are used as evidence in decisions.

Destruction produces the Phineas Gage syndrome: preserved intelligence and language, with catastrophic impairment of judgement, social conduct, and the ability to make decisions in one's own long-term interest.

SYSTEM MEMBERSHIP: The valuation and limbic system, and the default mode network (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: COMPARATOR-and-INTEGRATOR (valuation). State: value activation per option. Rule: integrates multisensory and emotional input into a value estimate; compares options.

DRIVES: Computes value FOR the drives; represents how much each option serves current needs.

PLASTICITY: Rapid reversal learning, dopamine-gated (updates value when contingencies change).

MODULATION: Dopamine, serotonin.

2.1.9 ANTERIOR CINGULATE CORTEX (ACC, BRODMANN AREAS 24, 32, 33)

Plain language: the part of the brain that notices when something is going wrong.

It fires when you make an error, when you face a conflict, when you feel pain, and when you are socially rejected.

INPUTS (SUBSCRIBE)

• From the DLPFC and the OFC.

• From the amygdala and the insula.

Cargo: interoceptive and affective state.

• From the anterior and mediodorsal thalamic nuclei.

• From the locus coeruleus and the ventral tegmental area.

  https://www.linkedin.com/in/d-r-dison

OUTPUTS (PUBLISH)

• To the DLPFC.

Cargo: the conflict signal that recruits executive control.

The ACC detects that control is needed; the DLPFC supplies it.

• To the premotor cortex and the SMA.

Cargo: modulation of response.

• To the periaqueductal gray and the autonomic brainstem nuclei.

Cargo: the autonomic component of emotion.

• To the striatum.

FUNCTION AND FACULTY

Conflict monitoring, error detection, effort allocation, and the affective (as opposed to the sensory) dimension of pain.

It is the region that makes pain unpleasant rather than merely detected, and, remarkably, it is the same region that is activated by social rejection, which is the anatomical basis of the finding that a broken heart and a broken arm share circuitry.

SYSTEM MEMBERSHIP: The salience network (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: COMPARATOR (conflict/error). State: conflict and error activation. Rule: fires on response conflict, error, pain, and social rejection; output recruits control.

DRIVES: Serves safety and social drives (registers their violation as pain).

PLASTICITY: Learns to predict effort and error, dopamine-gated.

MODULATION: Dopamine, norepinephrine.

2.1.10 POSTERIOR CINGULATE CORTEX (PCC) AND PRECUNEUS, BRODMANN AREAS 23, 31, 7m

Plain language: one of the most metabolically expensive regions in the brain, and one that is busiest when you are doing nothing in particular.

It is central to remembering your own life and imagining your future.

INPUTS (SUBSCRIBE)

• From the medial prefrontal cortex, by way of the cingulum bundle.

• From the hippocampus and parahippocampal cortex, by way of the cingulum bundle and the retrosplenial cortex.

  https://www.linkedin.com/in/d-r-dison

• From the angular gyrus and the lateral parietal cortex.

• From the anterior thalamic nuclei.

OUTPUTS (PUBLISH)

• To the medial prefrontal cortex, by way of the cingulum bundle.

• To the hippocampal formation, by way of the retrosplenial cortex and the parahippocampal cortex.

This is the principal route by which the cortex queries memory.

• To the angular gyrus.

FUNCTION AND FACULTY

Autobiographical memory retrieval, self-referential processing, scene construction, and mental time travel, meaning the capacity to project oneself into the past or the future.

Its metabolic rate is among the highest in the brain at rest, and it is one of the earliest sites of amyloid deposition and hypometabolism in Alzheimer's disease, which is why the loss of autobiographical self is so early and so central a feature of that disease.

SYSTEM MEMBERSHIP: The default mode network (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: ATTRACTOR (self-model hub). State: activation of self-referential representation. Rule: sustains autobiographical and self-schema patterns; core of the default mode network.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it serves internally-directed cognition, rather than monitoring a bodily variable of its own.

PLASTICITY: Slow; consolidates autobiographical structure.

MODULATION: Acetylcholine.

2.1.11 THE DEFAULT MODE NETWORK CORE TRIAD (PUC)

Plain language: three cortical regions that consistently act together as the anatomical heart of the resting, self-reflecting brain.

Each is individually named; the trio has never been named as a single anatomical object.

DEFINITION

The Default Mode Network Core Triad is the set of three cortical regions constituting the anatomical core of the default mode network: the medial prefrontal cortex (mPFC), the posterior cingulate cortex and precuneus (PCC), and the angular gyrus (AG).

PROOF OF EXISTENCE

These three regions show the highest within-network functional connectivity of any default mode network (DMN) nodes, are the first to deactivate on the onset of goal-directed tasks, and are the most consistently implicated in self-referential cognition across the entire imaging literature.

They are described together as the core nodes of the DMN in every major review of the network.

PROOF OF NAMELESSNESS

The DMN is named.

Each of the three regions is named.

The phrase "core nodes" appears descriptively in reviews.

But the specific three-region set, considered as a named Layer 9 anatomical entity with a defined internal interface structure, has never been formally designated.

The DMN is habitually described as a diffuse network rather than as a network with a named anatomical centre of gravity.

INTERNAL INTERFACES

• mPFC to PCC, by way of the cingulum bundle.

• PCC to AG, by way of short parietal association fibres.

• mPFC to AG, by way of the inferior fronto-occipital fasciculus and long association fibres.

FUNCTION AND FACULTY

The anatomical hub of self-referential processing: resting cognition, mind-wandering, autobiographical recall, envisioning the future, and modelling the mental states of others.

Its disruption, particularly the preferential attack on mPFC and PCC in Alzheimer's disease, is associated with the loss of autobiographical self.

SYSTEM MEMBERSHIP: The default mode network (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: Assign per Chapter 8 archetype (RELAY / INTEGRATOR / OSCILLATOR / ATTRACTOR / GATE / COMPARATOR). State: activation; gain. Rule: transforms input per its function above.

DRIVES: Not Applicable, for the reason that this construct does not monitor a bodily variable and so defends no set-point; it would serve a drive only if it monitored such a variable, which it does not.

PLASTICITY: Weights change by the three-factor rule of Chapter 10 where the region learns; otherwise Not Applicable, for the reason that a region without modifiable connections has no weights to update.

MODULATION: Set by the neuromodulators named in its inputs.

2.1.12 INSULAR CORTEX (THE INSULA)

Plain language: the hidden lobe, buried inside the fold between the temporal and frontal lobes.

It is where the body is felt.

Hunger, nausea, a racing heart, the urge to urinate, disgust, and the feeling of being alive in a body all pass through here.

INPUTS (SUBSCRIBE)

• From the ventromedial posterior nucleus (VMpo) of the thalamus, by way of the interoceptive thalamocortical projection.

This nucleus receives lamina I spinothalamic fibres carrying the physiological condition of every tissue of the body.

• From the nucleus tractus solitarius, indirectly by way of the parabrachial nucleus and thalamus.

Cargo: visceral afferent information from the vagus nerve.

• From the amygdala, the ACC, and the OFC.

• From the olfactory and gustatory pathways.

The anterior insula is the primary gustatory (taste) cortex.

OUTPUTS (PUBLISH)

• To the anterior cingulate cortex.

The insula and the ACC together constitute the salience network: the insula detects the significant bodily event, the ACC mobilizes the response.

• To the amygdala.

• To the hypothalamus and brainstem autonomic nuclei, by way of descending visceromotor projections.

• To the DLPFC.

FUNCTION AND FACULTY

Interoception, meaning the perception of the internal state of the body; the subjective feeling of emotion; disgust; empathy for pain; risk and uncertainty representation; and, in the anterior insula, awareness itself of one's bodily state.

It is the region proposed to generate the felt sense of being a self in a body.

SYSTEM MEMBERSHIP: The salience network and the interoceptive system (Layer 10).

SPECIFICATION BLOCK

  https://www.linkedin.com/in/d-r-dison

DYNAMICS: INTEGRATOR-and-COMPARATOR (interoception). State: activation representing bodily state and its salience. Rule: integrates visceral afferents; compares felt state to expectation; anterior insula switches the triple-network.

DRIVES: Reports the state of EVERY homeostatic drive to consciousness; the felt sense of need.

PLASTICITY: Plastic; learns bodily-state predictions.

MODULATION: Serotonin, norepinephrine.

2.1.13 BROCA'S AREA (LEFT INFERIOR FRONTAL GYRUS, BRODMANN AREAS 44 AND 45)

Plain language: where speech is assembled before it is spoken.

Damage here leaves you understanding everything and able to say almost nothing.

INPUTS (SUBSCRIBE)

• From Wernicke's area, by way of the arcuate fasciculus.

Cargo: the phonological content to be produced.

• From the DLPFC.

Cargo: the intention to speak, and the content.

• From the premotor cortex and SMA.

OUTPUTS (PUBLISH)

• To the primary motor cortex, and thence by the corticobulbar tract to the cranial nerve nuclei that drive the larynx, tongue, and lips.

• To the basal ganglia, for the sequencing of speech motor programs.

• To Wernicke's area, by way of the arcuate fasciculus, reciprocally.

FUNCTION AND FACULTY

Speech production, syntactic processing, and the motor planning of articulation.

Destruction produces Broca's aphasia: effortful, halting, agrammatic speech, with comprehension largely preserved, and, crucially, with the patient painfully aware of the deficit.

SYSTEM MEMBERSHIP: The language system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR-and-GATE (speech assembly). State: activation of the phonological-motor plan. Rule: sequences phonology into an articulatory program handed to motor cortex.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or

  https://www.linkedin.com/in/d-r-dison

motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Language-learning plasticity, strongest in the critical period.

MODULATION: Dopamine.

2.1.14 WERNICKE'S AREA (LEFT POSTERIOR SUPERIOR TEMPORAL GYRUS, BRODMANN AREA 22)

Plain language: where sound becomes meaning.

Damage here leaves you speaking fluently and saying nothing that makes sense, and unable to understand a word said to you.

INPUTS (SUBSCRIBE)

• From the primary auditory cortex, by way of short corticocortical association fibres.

• From the angular gyrus, carrying written language input by way of the visual system.

OUTPUTS (PUBLISH)

• To Broca's area, by way of the arcuate fasciculus.

This is the classical interface of the language system, and its selective destruction produces conduction aphasia, in which comprehension and fluency are intact but the patient cannot repeat what they have just heard.

• To the temporal association cortex and the semantic network.

FUNCTION AND FACULTY

Language comprehension, lexical access, and the mapping of sound onto meaning.

Destruction produces Wernicke's aphasia: fluent, well-articulated, grammatically correct speech that is empty of meaning, combined with severe comprehension deficit, and, tragically, with the patient unaware of the deficit.

SYSTEM MEMBERSHIP: The language system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: ATTRACTOR (lexical-semantic). State: activation of word-meaning representations. Rule: maps sound patterns onto stored meanings by pattern completion.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Vocabulary learning, lifelong but critical-period-biased.

MODULATION: Acetylcholine.

  https://www.linkedin.com/in/d-r-dison

2.1.15 THE MEDIAL TEMPORAL INTERFACE CORTICES: ENTORHINAL, PERIRHINAL, AND PARAHIPPOCAMPAL

Plain language: the funnel.

Everything the cortex knows must pass through these regions to get into the hippocampus and become a memory.

ENTORHINAL CORTEX (BRODMANN AREA 28)

INPUTS (SUBSCRIBE): From the perirhinal and parahippocampal cortices, which in turn collect from the entire association cortex.

From the amygdala.

From the olfactory bulb, directly.

OUTPUTS (PUBLISH): To the dentate gyrus, CA3, and CA1 of the hippocampus, by way of the perforant path.

This is the principal gateway into the hippocampal formation.

To the neocortex, by way of the return projections that mediate memory retrieval.

FUNCTION: The obligatory interface between neocortex and hippocampus.

Contains grid cells, which fire in a hexagonal lattice across an environment and provide the metric of the brain's spatial map.

Layer II projects to dentate and CA3; layer III projects directly to CA1, forming the temporoammonic path.

PERIRHINAL CORTEX (BRODMANN AREAS 35 AND 36)

INPUTS (SUBSCRIBE): From the ventral visual stream and other unimodal association cortices.

Cargo: object identity.

OUTPUTS (PUBLISH): To the lateral entorhinal cortex, and thence to the hippocampus.

FUNCTION: Object recognition memory and familiarity judgement.

Carries the "what" of an experience.

PARAHIPPOCAMPAL CORTEX (BRODMANN AREA 36 posterior)

INPUTS (SUBSCRIBE): From the dorsal visual stream and the posterior parietal and retrosplenial cortices.

Cargo: spatial and scene information.

OUTPUTS (PUBLISH): To the medial entorhinal cortex, and thence to the hippocampus.

FUNCTION: Spatial context and scene processing.

  https://www.linkedin.com/in/d-r-dison

Carries the "where" of an experience.

Note the elegance of this architecture: the "what" and the "where" streams, separated since V1, are kept separate all the way to the entorhinal cortex, and are bound together for the first time inside the hippocampus.

This is the anatomical basis of episodic memory, which is precisely the binding of what happened to where and when it happened.

SYSTEM MEMBERSHIP: The declarative memory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: Assign per Chapter 8 archetype (RELAY / INTEGRATOR / OSCILLATOR / ATTRACTOR / GATE / COMPARATOR). State: activation; gain. Rule: transforms input per its function above.

DRIVES: Not Applicable, for the reason that this construct does not monitor a bodily variable and so defends no set-point; it would serve a drive only if it monitored such a variable, which it does not.

PLASTICITY: Weights change by the three-factor rule of Chapter 10 where the region learns; otherwise Not Applicable, for the reason that a region without modifiable connections has no weights to update.

MODULATION: Set by the neuromodulators named in its inputs.

2.1.16 RETROSPLENIAL CORTEX (BRODMANN AREAS 29 AND 30)

Plain language: the translator between the map in your head and the view from where you stand.

INPUTS (SUBSCRIBE): From the anterior thalamic nuclei, the subiculum, the parahippocampal cortex, and the posterior cingulate.

OUTPUTS (PUBLISH): To the parahippocampal cortex, the anterior thalamus, and the posterior cingulate.

FUNCTION AND FACULTY: The translation between egocentric spatial representation (where things are relative to me) and allocentric representation (where things are on the map).

Destruction produces topographical disorientation: the patient can recognize landmarks but cannot use them to navigate.

SYSTEM MEMBERSHIP: The spatial navigation system and the default mode network (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: COMPARATOR (reference-frame translation). State: activation coding heading and place. Rule: translates between egocentric and allocentric frames.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it serves spatial navigation, rather than monitoring a bodily variable of its own.

  https://www.linkedin.com/in/d-r-dison

PLASTICITY: Plastic map learning.

MODULATION: Acetylcholine.

2.1.17 POSTERIOR PARIETAL CORTEX (BRODMANN AREAS 5, 7, 39, 40)

Plain language: where the senses are stitched together into a single space in which action can be planned.

Damage on the right can make a patient deny that the left half of the world exists.

INPUTS (SUBSCRIBE)

• From S1, by way of feedforward corticocortical projections.

• From V5/MT and the dorsal visual stream.

• From the vestibular nuclei, by way of thalamic relay.

Cargo: head position and motion.

• From the pulvinar and the lateral posterior thalamic nucleus.

OUTPUTS (PUBLISH)

• To the premotor cortex and the frontal eye fields, by way of the superior longitudinal fasciculus.

Cargo: the spatial coordinates for reaching, grasping, and looking.

• To the DLPFC, by way of the superior longitudinal fasciculus.

• To the cerebellum, by way of the corticopontine projection.

• To the hippocampal formation, by way of the parahippocampal and retrosplenial cortices.

FUNCTION AND FACULTY

Multisensory spatial integration; the construction of body schema; visually guided reaching and grasping; and directed spatial attention.

The angular gyrus (BA 39) additionally serves semantic processing and is a core node of the default mode network.

Destruction of the right posterior parietal cortex produces hemispatial neglect, one of the strangest syndromes in neurology: the patient does not merely fail to see the left side of the world, but fails to conceive of it, eating only the right half of the plate and drawing only the right half of a clock, and denying that anything is missing.

SYSTEM MEMBERSHIP: The dorsal attention network and the visual system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (multisensory space). State: activation of spatial and body-schema representations. Rule: fuses senses into a common spatial frame for action and attention.

DRIVES: Serves action toward drive-relevant targets.

PLASTICITY: Plastic; tool use and spatial learning.

MODULATION: Acetylcholine, norepinephrine.

2.1.18 TEMPORAL ASSOCIATION CORTEX AND THE INFEROTEMPORAL CORTEX

Plain language: where the brain finally knows what it is looking at.

INPUTS (SUBSCRIBE): From V4 and the ventral visual stream.

From the auditory association cortex.

OUTPUTS (PUBLISH): To the perirhinal cortex, and thence to memory.

To the amygdala, by way of direct temporo-amygdalar projections.

To the OFC, by way of the uncinate fasciculus.

To the DLPFC.

FUNCTION AND FACULTY: Invariant object recognition, meaning the recognition of an object as the same object despite changes in size, angle, and lighting.

It contains the fusiform face area, whose destruction produces prosopagnosia, the inability to recognize faces, including one's own in a mirror, with other visual recognition intact.

It also houses semantic memory, the store of general knowledge about the world; its progressive degeneration produces semantic dementia.

SYSTEM MEMBERSHIP: The ventral visual stream and the semantic system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: ATTRACTOR (object recognition). State: activation of object/category representations. Rule: invariant recognition by pattern completion; holds semantic knowledge.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Slow perceptual and semantic learning; face/object expertise.

MODULATION: Acetylcholine.

2.1.19 FRONTAL EYE FIELDS (FEF, BRODMANN AREA 8)

  https://www.linkedin.com/in/d-r-dison

Plain language: the cortex's gaze-command centre, the region that decides where the eyes look next, and a core node of the brain's top-down attention network.

INPUTS (SUBSCRIBE): From the visual association areas, chiefly V4, V5/MT, and the posterior parietal cortex (especially the lateral intraparietal area), carrying candidate salient locations; from the dorsolateral prefrontal cortex, carrying goals; and from the mediodorsal and pulvinar thalamus.

OUTPUTS (PUBLISH): To the superior colliculus, the fast subcortical gaze map; to the brainstem gaze centres, including the paramedian pontine reticular formation for horizontal saccades; to the caudate nucleus; and reciprocally to the parietal and prefrontal cortex.

FUNCTION AND FACULTY: The voluntary control of saccadic eye movements and the top-down allocation of spatial attention.

It converts a chosen target, driven either by a goal from the prefrontal cortex or by salience from the parietal cortex, into a gaze command; and, even when the eyes do not move, it biases visual processing toward the attended location, which is the mechanism of covert attention.

It is a principal source of the top-down attention signal that the rest of the visual system subscribes to.

SYSTEM MEMBERSHIP: The oculomotor system and the dorsal attention network (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE-and-SELECTOR (a spatial priority map issuing a gaze command). State: a spatial priority and salience map; the selected saccade target. Rule: integrates goal from the dorsolateral prefrontal cortex and salience from the parietal cortex into a winner-take-all target, issues the gaze command, and biases sensory gain at the attended location.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception and action on which the drives act.

PLASTICITY: Learns target values and attentional priorities, reward-modulated.

MODULATION: Dopamine, acetylcholine for attentional gain, and norepinephrine.

2.1.20 SUPRAMARGINAL GYRUS (SMG, BRODMANN AREA 40)

Plain language: the phonological and praxis part of the inferior parietal lobule, on the left the store of speech sounds and on the right a hub of spatial attention.

INPUTS (SUBSCRIBE): somatosensory, auditory, and visual association cortex; the superior temporal cortex.

OUTPUTS (PUBLISH): Broca's area by way of the arcuate fasciculus; the premotor cortex.

FUNCTION AND FACULTY: phonological short-term storage, sound-based language processing, tool use and skilled action (praxis), and tactile-spatial attention.

SYSTEM MEMBERSHIP: the language and dorsal attention systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (phonological/praxis buffer). State: the active phonological or action representation. Rule: holds and manipulates sound and action codes.

MODULATION: dopamine, acetylcholine.

2.1.21 TEMPOROPARIETAL JUNCTION (TPJ)

Plain language: the crossroads of attention and social thought, where the brain reorients to the unexpected and reasons about other minds.

INPUTS (SUBSCRIBE): the ventral attention/salience network; the superior temporal sulcus; multisensory association cortex.

OUTPUTS (PUBLISH): the frontal cortex for mentalizing; the dorsal attention network, which it interrupts.

FUNCTION AND FACULTY: bottom-up reorienting of attention to salient events, and theory of mind, the attribution of beliefs and intentions to others.

SYSTEM MEMBERSHIP: the ventral attention and social-cognition systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: COMPARATOR-and-INTERRUPT. State: detected salience or social inference. Rule: flags the unexpected and computes others' mental states.

MODULATION: norepinephrine, acetylcholine.

2.1.22 SUPERIOR TEMPORAL SULCUS (STS)

Plain language: the social-perception strip, tuned to faces in motion, voices, and biological movement.

INPUTS (SUBSCRIBE): the fusiform and motion visual areas; the auditory cortex.

OUTPUTS (PUBLISH): the amygdala; the prefrontal mentalizing cortex; the temporoparietal junction.

FUNCTION AND FACULTY: perception of biological motion, voice, gaze, and dynamic faces, and the multisensory basis of social perception.

SYSTEM MEMBERSHIP: the social-perception system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (multisensory social percepts). State: the current social percept. Rule: binds moving faces, voices, and motion into social meaning.

  https://www.linkedin.com/in/d-r-dison

MODULATION: acetylcholine, oxytocin (social salience).

2.1.23 MIDDLE TEMPORAL GYRUS (MTG, BRODMANN AREA 21)

Plain language: a principal store and access point of word meaning and general knowledge.

INPUTS (SUBSCRIBE): the superior temporal auditory cortex; the ventral visual stream; the angular gyrus.

OUTPUTS (PUBLISH): the inferior frontal cortex; the anterior temporal semantic hub.

FUNCTION AND FACULTY: lexical and semantic processing, retrieval of concepts and word meanings, and a node of the default mode network.

SYSTEM MEMBERSHIP: the language and semantic memory systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (semantic access). State: the active concept. Rule: maps words and objects to stored meaning.

MODULATION: dopamine, acetylcholine.

2.1.24 FRONTOPOLAR / ANTERIOR PREFRONTAL CORTEX (BRODMANN AREA 10)

Plain language: the most anterior cortex, the seat of holding a plan while doing something else, and of thinking about one's own thoughts.

INPUTS (SUBSCRIBE): the dorsolateral and ventrolateral prefrontal cortex; the anterior temporal cortex; the cingulate.

OUTPUTS (PUBLISH): the dorsolateral and ventrolateral prefrontal cortex; the medial prefrontal cortex.

FUNCTION AND FACULTY: prospective memory, multitasking and branching, relational and analogical reasoning, metacognition, and the balance of exploration against exploitation.

SYSTEM MEMBERSHIP: the executive control system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (supervisory/branching control). State: the held superordinate goal. Rule: maintains a pending goal across an interruption and coordinates subgoals.

MODULATION: dopamine, norepinephrine.

2.1.25 PRE-SUPPLEMENTARY MOTOR AREA (PRE-SMA, ROSTRAL MEDIAL BRODMANN AREA 6)

Plain language: the switchboard of action, where a movement is chosen, sequenced, or cancelled before it is executed.

INPUTS (SUBSCRIBE): the dorsolateral prefrontal cortex; the anterior cingulate; the basal ganglia by way of the thalamus.

  https://www.linkedin.com/in/d-r-dison

OUTPUTS (PUBLISH): the supplementary motor area and primary motor cortex; the striatum and, by the hyperdirect route, the subthalamic nucleus.

FUNCTION AND FACULTY: action selection, motor-sequence assembly, set-switching, and response inhibition, the cortical origin of the stopping signal.

SYSTEM MEMBERSHIP: the action-control system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE-and-SELECTOR. State: the selected or withheld action. Rule: selects, sequences, or cancels an action before execution.

MODULATION: dopamine, norepinephrine.

2.1.26 SUPPLEMENTARY EYE FIELD (SEF, MEDIAL BRODMANN AREA 6/8)

Plain language: the medial gaze planner, for internally-guided and sequenced eye movements.

INPUTS (SUBSCRIBE): the frontal eye fields; the parietal cortex; the cingulate.

OUTPUTS (PUBLISH): the superior colliculus; the brainstem gaze centres.

FUNCTION AND FACULTY: internally-generated, learned, and sequential saccades, and the coordination of gaze with intended action.

SYSTEM MEMBERSHIP: the oculomotor system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: SELECTOR (internally-guided gaze). State: the planned gaze sequence. Rule: issues internally-timed saccade plans.

MODULATION: dopamine.

2.1.27 SUPERIOR PARIETAL LOBULE (SPL, BRODMANN AREAS 5 AND 7)

Plain language: the dorsal-stream workbench for space and reaching, where the body and the visual world are held in a common map.

INPUTS (SUBSCRIBE): the primary somatosensory cortex; the dorsal visual stream; the pulvinar.

OUTPUTS (PUBLISH): the premotor cortex and frontal eye fields; the dorsolateral prefrontal cortex; the intraparietal sulcus.

FUNCTION AND FACULTY: visuospatial attention, spatial working memory, sensorimotor integration for reaching, and coordinate transformation between the eye, head, and hand.

SYSTEM MEMBERSHIP: the dorsal attention and sensorimotor systems (Layer 10).

SPECIFICATION BLOCK

  https://www.linkedin.com/in/d-r-dison

DYNAMICS: INTEGRATOR (spatial coordinate frames). State: the current spatial map. Rule: transforms between egocentric reference frames for action.

MODULATION: acetylcholine, norepinephrine.

2.1.28 PARAHIPPOCAMPAL PLACE AREA (PPA)

Plain language: the scene detector, the patch of ventral visual cortex that recognizes places and spatial layouts.

INPUTS (SUBSCRIBE): the ventral visual stream; the retrosplenial cortex.

OUTPUTS (PUBLISH): the hippocampal formation; the entorhinal cortex; the retrosplenial cortex.

FUNCTION AND FACULTY: recognition of scenes, spatial layouts, and landmarks, the visual front end of navigation and context memory.

SYSTEM MEMBERSHIP: the navigation and scene-memory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-DETECTOR (scene selectivity). State: the recognized scene. Rule: signals the presence and identity of a place.

MODULATION: acetylcholine.

2.1.29 VISUAL WORD FORM AREA (VWFA, LEFT OCCIPITOTEMPORAL FUSIFORM)

Plain language: the reading patch, a left-hemisphere region that recognizes written words as visual objects.

INPUTS (SUBSCRIBE): the early visual cortex.

OUTPUTS (PUBLISH): the language network (the superior temporal and inferior frontal cortex).

FUNCTION AND FACULTY: fast, orientation-invariant recognition of letter strings and written words, the visual gateway to reading.

SYSTEM MEMBERSHIP: the language system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-DETECTOR (orthographic). State: the recognized word form. Rule: maps visual letter strings onto the language system.

MODULATION: acetylcholine.

2.1.30 THE LATERAL OCCIPITAL COMPLEX AND EXTRASTRIATE BODY AREA (LOC, EBA)

Plain language: the object-and-body recognition cortex on the lateral occipitotemporal surface.

INPUTS (SUBSCRIBE): the early visual areas and the motion complex.

OUTPUTS (PUBLISH): the ventral temporal recognition areas; the parietal and premotor cortex (for bodies).

FUNCTION AND FACULTY: recognition of object shape (the lateral occipital complex) and of body parts and postures (the extrastriate body area).

SYSTEM MEMBERSHIP: the visual object-recognition system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-DETECTOR (shape/body selectivity). State: the recognized object or body form. Rule: signals object and body identity.

MODULATION: acetylcholine.

2.1.31 SECONDARY SOMATOSENSORY CORTEX (S2, PARIETAL OPERCULUM)

Plain language: the higher touch area, for recognizing objects by feel and integrating touch across the two sides of the body.

INPUTS (SUBSCRIBE): the primary somatosensory cortex; the somatosensory thalamus.

OUTPUTS (PUBLISH): the insula; the posterior parietal association cortex.

FUNCTION AND FACULTY: tactile object recognition, sensorimotor integration, and the affective and attentional dimensions of touch.

SYSTEM MEMBERSHIP: the somatosensory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (tactile objects). State: the felt object. Rule: binds touch into recognizable objects and refers it to the insula.

MODULATION: acetylcholine, norepinephrine.

2.1.32 THE AUDITORY BELT AND PARABELT (BRODMANN AREA 42 AND BEYOND)

Plain language: the second and third stages of hearing, where simple tones become complex sounds and speech precursors.

INPUTS (SUBSCRIBE): the primary auditory cortex; the medial geniculate nucleus.

OUTPUTS (PUBLISH): the superior temporal gyrus and sulcus; the language and prefrontal cortex.

FUNCTION AND FACULTY: analysis of complex sounds, spectrotemporal patterns, and the acoustic precursors of speech and music.

SYSTEM MEMBERSHIP: the auditory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: hierarchical TRANSFORM (spectrotemporal). State: the complex-sound representation. Rule: builds complex sound features from tonotopic input.

MODULATION: acetylcholine, norepinephrine.

  https://www.linkedin.com/in/d-r-dison

2.1.33 THE INSULA, ANTERIOR AND POSTERIOR DIVISIONS

Plain language: a division note for the insula (see 2.1.12): the posterior insula is the primary interoceptive and pain cortex, the anterior insula the seat of subjective feeling and salience.

INPUTS (SUBSCRIBE): posterior insula: viscerosensory thalamus (the posterior ventral medial nucleus) and spinothalamic pathways. Anterior insula: the posterior insula, the anterior cingulate, and the amygdala.

OUTPUTS (PUBLISH): posterior to anterior insula; anterior insula to the anterior cingulate, prefrontal cortex, amygdala, and autonomic centres.

FUNCTION AND FACULTY: posterior: primary interoception and pain. Anterior: interoceptive awareness, subjective feeling states, empathy, and salience detection (a hub of the salience network).

SYSTEM MEMBERSHIP: the interoceptive and salience systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: posterior RELAY to anterior COMPARATOR (interoceptive salience). State: the felt bodily state and its salience. Rule: reads the body and computes its subjective and salient significance.

MODULATION: norepinephrine, serotonin, dopamine.

2.2 THE SUBCORTICAL TELENCEPHALON

2.2.1 THE HIPPOCAMPAL FORMATION (DENTATE GYRUS, CA3, CA2, CA1, SUBICULUM)

Plain language: the structure that turns experience into memory.

It does not store your memories.

It writes them, and hands them to the cortex for keeping.

Remove it, and the patient lives permanently in the present moment, unable to lay down a single new fact or event, while everything learned before the surgery remains intact.

This is the most thoroughly characterized circuit in the mammalian brain, and it is presented here in the detail it deserves.

INPUTS (SUBSCRIBE)

• From the entorhinal cortex layer II, by way of the perforant path, terminating on the granule cells of the dentate gyrus.

This is the principal input, and it carries the highly processed convergence of the entire association cortex.

The path is so named because its fibres perforate the subiculum on their way in.

• From the entorhinal cortex layer III, by way of the temporoammonic path, terminating directly on CA1 and bypassing the dentate gyrus and CA3 entirely.

  https://www.linkedin.com/in/d-r-dison

This is a genuine within-structure shortcut, and it appears to carry current sensory reality against which CA3's reconstructed prediction can be compared.

• From the medial septal nucleus and the diagonal band of Broca, by way of the fornix.

Cargo: acetylcholine, and the pacemaker signal that generates the hippocampal theta rhythm.

• From the locus coeruleus (norepinephrine), the raphe nuclei (serotonin), and the ventral tegmental area (dopamine), by way of ascending neuromodulatory projections.

Dopamine arriving from the VTA acts as a novelty signal and gates whether an experience is worth storing.

• From the amygdala, by way of direct amygdalo-hippocampal projections.

Cargo: emotional salience, which is why emotional events are remembered better.

• From the contralateral hippocampus, by way of the hippocampal commissure.

INTERNAL CIRCUIT: THE TRISYNAPTIC LOOP

• Synapse 1: entorhinal cortex to dentate gyrus, by the perforant path.

• Synapse 2: dentate gyrus to CA3, by the mossy fibres.

These are among the largest and most powerful synapses in the brain: a single mossy fibre can drive a CA3 pyramidal cell to fire on its own, which is why the dentate is called a detonator.

• Synapse 3: CA3 to CA1, by the Schaffer collaterals.

This is the synapse at which long-term potentiation was discovered and where most of what is known about the molecular biology of memory has been learned.

• CA3 additionally possesses dense recurrent collaterals onto itself, making it an autoassociative network capable of pattern completion: the reconstruction of a whole memory from a fragment of it.

• CA1 projects to the subiculum, which is the principal output stage.

OUTPUTS (PUBLISH)

• From the subiculum and CA1, by way of the fornix, to the mammillary bodies of the hypothalamus, the anterior thalamic nuclei, the septal nuclei, and the nucleus accumbens.

The fornix is the great efferent cable of the hippocampus, and it is the first leg of the Papez circuit.

• From CA1 and the subiculum, by way of return projections to the entorhinal cortex, and thence back to the neocortex.

  https://www.linkedin.com/in/d-r-dison

This is the retrieval route: the path by which a memory is returned to the cortex that will eventually own it.

• To the nucleus accumbens, by way of the fornix.

Cargo: contextual information used to gate motivated behaviour.

FUNCTION AND FACULTY

The encoding of declarative memory, both episodic (events) and semantic (facts); pattern separation, performed by the dentate gyrus, which makes similar experiences distinguishable; pattern completion, performed by CA3, which retrieves a whole from a part; spatial navigation, mediated by place cells, which fire when the animal occupies a particular location; and systems consolidation, the gradual transfer of memory from hippocampal dependence to cortical independence over years.

Bilateral destruction produces the syndrome of patient H.M.: dense anterograde amnesia, the permanent inability to form new declarative memories, with intact working memory, intact procedural learning, intact intelligence, intact personality, and intact remote memory.

THE TEMPORAL DIMENSION

The hippocampus is also the brain's clearest example of a structure whose function is inseparable from its rhythm, and the account above is incomplete without it.

During active exploration the hippocampus runs at theta, roughly 4 to 8 hertz, paced from the medial septum.

Place cells fire at systematically advancing phases of the theta cycle as the animal crosses their fields, a phenomenon called phase precession, and this converts a spatial sequence into a temporal firing sequence that the synapses of CA3 can learn.

Encoding and retrieval are segregated to opposite phases of the theta cycle, so that the hippocampus alternates, several times a second, between writing and reading.

During rest and sleep the rhythm changes entirely.

The hippocampus generates sharp-wave ripples: brief, extraordinarily synchronous bursts at 150 to 250 hertz, the most synchronous population events in the healthy brain.

During a ripple, the place-cell sequences of recent experience are REPLAYED, compressed roughly twentyfold, and sometimes run in reverse.

This replay is the physical act of consolidation: it reactivates the day's traces and, nested inside thalamic spindles and cortical slow oscillations (Chapter 5), transfers them to the cortex for permanent storage.

Suppress ripples experimentally and memory consolidation fails.

The hippocampus does not merely store the day; at night, it dictates the day back to the cortex, quickly and repeatedly, until the cortex has learned it.

CA3's dense recurrent connectivity, noted above as the substrate of pattern completion, is also what makes it an oscillatory engine: it is an autoassociative network that can ignite a whole stored pattern from a fragment, and the ripple is that ignition made visible.

SYSTEM MEMBERSHIP: The declarative memory system and the spatial navigation system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY, tonotopic. State: activation per frequency channel; gain. Rule: tracks medial geniculate input scaled by gain.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Adult plasticity of frequency tuning, gated by attention/acetylcholine.

MODULATION: Acetylcholine, norepinephrine.

2.2.1.1 PRESUBICULUM

Plain language: the head-direction relay of the hippocampal formation, the brain's internal compass reading.

INPUTS (SUBSCRIBE): the anterior thalamic nuclei (head-direction signal); the subiculum; the retrosplenial cortex.

OUTPUTS (PUBLISH): the medial entorhinal cortex, feeding the grid-cell map.

FUNCTION AND FACULTY: carrying the head-direction signal into the spatial-memory system for orientation and navigation.

SYSTEM MEMBERSHIP: the navigation system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY (head-direction). State: current heading. Rule: relays the compass signal to the grid system.

MODULATION: acetylcholine.

2.2.1.2 PARASUBICULUM

Plain language: a parahippocampal relay carrying grid, head-direction, and boundary signals into the entorhinal map.

INPUTS (SUBSCRIBE): the subiculum; the anterior thalamus.

OUTPUTS (PUBLISH): layer two of the entorhinal cortex.

FUNCTION AND FACULTY: gating spatial signals (grid, head-direction, border) into the entorhinal-hippocampal system.

SYSTEM MEMBERSHIP: the navigation system (Layer 10).

  https://www.linkedin.com/in/d-r-dison

SPECIFICATION BLOCK

DYNAMICS: RELAY (spatial signals). State: current spatial code. Rule: gates spatial signals into entorhinal cortex.

MODULATION: acetylcholine.

2.2.2 THE AMYGDALA (LATERAL, BASAL, ACCESSORY BASAL, CENTRAL, MEDIAL, AND CORTICAL NUCLEI)

Plain language: the threat detector, and more generally the assigner of emotional significance.

It decides, faster than you can think, whether the thing in front of you matters.

The amygdala is not one structure but a cluster of functionally distinct nuclei, and treating it as a unit is the commonest error made about it.

THE BASOLATERAL COMPLEX (LATERAL, BASAL, ACCESSORY BASAL NUCLEI): THE INPUT DIVISION

INPUTS (SUBSCRIBE)

• From all sensory association cortices, by way of corticoamygdalar projections.

Cargo: fully processed perceptual information.

• From the thalamus directly, by way of the thalamo-amygdalar projection, most notably from the medial geniculate and posterior thalamic nuclei.

This is the celebrated "low road": a fast, crude, subcortical route by which a threat can reach the amygdala before the cortex has finished identifying it.

It is why you jump back from the coiled shape on the path before you have consciously recognized it as a snake, or as a rope.

• From the hippocampus, by way of hippocampo-amygdalar projections.

Cargo: the context in which an event occurred.

• From the medial prefrontal cortex and OFC, by way of the uncinate fasciculus.

Cargo: top-down regulation.

• From the ventral pallidum and substantia innominata, by way of dense GABAergic basal forebrain projections.

INTERNAL INTERFACE

• The basolateral complex projects to the central nucleus, by way of intra-amygdalar projections passing through the intercalated cell masses, which are clusters of GABAergic neurons that gate the flow from input division to output division.

  https://www.linkedin.com/in/d-r-dison

These intercalated masses are the physical substrate of fear extinction: the prefrontal cortex excites them, and they in turn inhibit the central nucleus, suppressing the fear response without erasing the underlying memory.

THE CENTRAL NUCLEUS: THE OUTPUT DIVISION

OUTPUTS (PUBLISH)

• To the hypothalamus (lateral hypothalamic area), by way of the ventral amygdalofugal pathway.

Effect: sympathetic activation, meaning increased heart rate and blood pressure.

• To the periaqueductal gray, by way of the ventral amygdalofugal pathway.

Effect: freezing, and defensive behaviour.

• To the paraventricular nucleus of the hypothalamus.

Effect: activation of the HPA axis, and therefore cortisol release.

• To the parabrachial nucleus.

Effect: altered respiration.

• To the locus coeruleus.

Effect: arousal and vigilance.

• To the nucleus basalis of Meynert.

Effect: cortical acetylcholine release, and therefore attention.

• To the bed nucleus of the stria terminalis, by way of the stria terminalis.

Effect: sustained anxiety, as distinct from acute fear.

• To the mediodorsal thalamus and thence to the prefrontal cortex, by way of the ventral amygdalofugal pathway.

Effect: the conscious feeling of fear.

THE MEDIAL NUCLEUS

INPUTS (SUBSCRIBE): From the accessory olfactory bulb and the main olfactory system.

OUTPUTS (PUBLISH): To the medial hypothalamus, by way of the stria terminalis.

FUNCTION: Pheromonal and social-olfactory processing; reproductive and defensive social behaviour.

FUNCTION AND FACULTY OF THE AMYGDALA AS A WHOLE

Threat detection; fear conditioning, in which a neutral stimulus paired with an aversive one acquires the power to elicit fear; emotional modulation of memory, whereby emotionally arousing events are remembered more vividly; the recognition of fear in the faces of others; and the assignment of value to both aversive and appetitive stimuli.

Bilateral destruction produces Kluver-Bucy syndrome and, in the celebrated human case of patient S.M., a complete absence of felt fear: the patient could not be frightened by snakes, haunted houses, or a knife held to her throat, and could not recognize fear in a photographed face, though every other emotion was recognized normally.

SYSTEM MEMBERSHIP: The limbic and salience systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: Assign per Chapter 8 archetype (RELAY / INTEGRATOR / OSCILLATOR / ATTRACTOR / GATE / COMPARATOR). State: activation; gain. Rule: transforms input per its function above.

DRIVES: Not Applicable, for the reason that this construct does not monitor a bodily variable and so defends no set-point; it would serve a drive only if it monitored such a variable, which it does not.

PLASTICITY: Weights change by the three-factor rule of Chapter 10 where the region learns; otherwise Not Applicable, for the reason that a region without modifiable connections has no weights to update.

MODULATION: Set by the neuromodulators named in its inputs.

2.2.3 THE BED NUCLEUS OF THE STRIA TERMINALIS (BNST)

Plain language: the amygdala's slow cousin.

Where the amygdala handles a threat that is here now, the BNST handles the dread of a threat that might come.

INPUTS (SUBSCRIBE): From the basolateral and central amygdala, by way of the stria terminalis.

From the hippocampus and the medial prefrontal cortex.

From the paraventricular nucleus.

OUTPUTS (PUBLISH): To the paraventricular nucleus of the hypothalamus, driving the HPA axis.

To the periaqueductal gray.

To the ventral tegmental area.

To the locus coeruleus.

FUNCTION AND FACULTY: Sustained anxiety, as distinct from phasic fear; the anticipation of unpredictable threat; and the integration of stress with the HPA axis.

It is sometimes called the extended amygdala, together with the central nucleus and the nucleus accumbens shell.

  https://www.linkedin.com/in/d-r-dison

It is a principal target of research into generalized anxiety disorder, precisely because it mediates anxiety that has no object.

SYSTEM MEMBERSHIP: The stress and anxiety system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (sustained anxiety). State: slow threat-anticipation activation. Rule: integrates diffuse/unpredictable threat over long timescales (vs the amygdala's phasic fear).

DRIVES: Serves the SAFETY drive, its sustained-anxiety arm.

PLASTICITY: Slow stress-related plasticity; set-point can shift with chronic stress.

MODULATION: CRH (corticotropin-releasing hormone), norepinephrine.

2.2.4 THE BASAL GANGLIA

The basal ganglia are a set of interconnected subcortical nuclei that together perform action selection: the choosing of one behaviour and the suppression of all competing alternatives.

They are best understood not as a motor structure but as a general-purpose selector, applied to movements in the motor loop, to thoughts in the associative loop, and to motivations in the limbic loop.

The architecture is a loop: cortex to striatum to pallidum to thalamus and back to cortex.

The default state of the loop is inhibition.

The basal ganglia hold every possible action suppressed, and selection consists of releasing one.

This inhibitory-by-default design is worth pausing on, because it explains the diseases.

The output nuclei tonically inhibit the thalamus.

To act is to remove that inhibition from one channel.

Therefore too little basal ganglia output produces involuntary movement (Huntington's chorea), and too much produces the inability to move at all (Parkinsonian akinesia).

The pathology is a failure of the brake, in one direction or the other.

2.2.4.1 THE STRIATUM (CAUDATE NUCLEUS, PUTAMEN, AND NUCLEUS ACCUMBENS)

Plain language: the input gate of the basal ganglia.

Everything the cortex wants to do is proposed here first.

INPUTS (SUBSCRIBE)

  https://www.linkedin.com/in/d-r-dison

• From essentially the entire cerebral cortex, by way of the corticostriatal projection arising from lamina 5.

The topography is strict and functionally decisive: sensorimotor cortex projects to the putamen; association cortex projects to the caudate; limbic cortex, the amygdala, and the hippocampus project to the nucleus accumbens.

This topography is what creates the three parallel loops.

• From the substantia nigra pars compacta, by way of the nigrostriatal dopamine pathway, to the caudate and putamen.

This is the pathway that dies in Parkinson's disease.

• From the ventral tegmental area, by way of the mesolimbic dopamine pathway, to the nucleus accumbens.

This is the pathway implicated in reward and in addiction.

• From the intralaminar thalamic nuclei (centromedian and parafascicular), by way of the thalamostriatal projection.

Cargo: attentional and arousal signals.

OUTPUTS (PUBLISH)

• Direct pathway: from D1-receptor-expressing medium spiny neurons, GABAergic and inhibitory, to the internal globus pallidus (GPi) and the substantia nigra pars reticulata (SNr).

Effect: inhibits the inhibitor, thereby disinhibiting the thalamus, thereby permitting movement.

The direct pathway is the "go" pathway.

• Indirect pathway: from D2-receptor-expressing medium spiny neurons to the external globus pallidus (GPe), thence to the subthalamic nucleus, thence to GPi and SNr.

Effect: increases inhibition of the thalamus, thereby suppressing movement.

The indirect pathway is the "no-go" pathway.

Dopamine is the switch that biases between them: it excites D1 cells (promoting go) and inhibits D2 cells (removing no-go).

Both effects push in the same direction, which is why dopamine loss is so catastrophic for movement.

FUNCTION AND FACULTY: Action selection; habit formation and procedural learning (putamen); goal-directed and cognitive action selection (caudate); reward-based learning and motivation (nucleus accumbens).

Degeneration of striatal medium spiny neurons produces Huntington's disease, in which the indirect pathway fails first, releasing unwanted movements as chorea.

SYSTEM MEMBERSHIP: The motor, cognitive, and limbic basal ganglia loops (Layer 10).

  https://www.linkedin.com/in/d-r-dison

SPECIFICATION BLOCK

DYNAMICS: GATE-and-COMPARATOR (limbic action gate). State: incentive-salience activation. Rule: gates motivated behavior by value; assigns wanting.

DRIVES: Translates drive-relevant value into approach; the wanting stage of every drive.

PLASTICITY: Strong dopamine-gated three-factor plasticity; the site hijacked in addiction.

MODULATION: Dopamine (mesolimbic), opioids.

2.2.4.2 THE GLOBUS PALLIDUS, EXTERNAL SEGMENT (GPe)

INPUTS (SUBSCRIBE): From the striatum, by the indirect pathway (GABAergic).

From the subthalamic nucleus (glutamatergic).

OUTPUTS (PUBLISH): To the subthalamic nucleus (GABAergic).

To the GPi and SNr.

Also, by way of recently characterized arkypallidal neurons, back to the striatum.

FUNCTION: The central relay of the indirect pathway, and, through its reciprocal loop with the STN, a generator of the pathological beta oscillations that characterize Parkinson's disease.

2.2.4.3 THE SUBTHALAMIC NUCLEUS (STN)

Plain language: the emergency brake.

The only excitatory nucleus inside the basal ganglia, and the target of deep brain stimulation for Parkinson's disease.

INPUTS (SUBSCRIBE)

• From the GPe, by the indirect pathway (GABAergic).

• From the motor and premotor cortex, directly, by way of the hyperdirect pathway, bypassing the striatum entirely.

This is a monosynaptic glutamatergic projection and it is fast.

OUTPUTS (PUBLISH): To the GPi and SNr (glutamatergic, excitatory).

Effect: increases inhibitory output, suppressing movement globally.

FUNCTION AND FACULTY: Global response inhibition, the fast "hold your horses" signal that arrests an action already underway.

Its destruction produces hemiballismus: violent, uncontrollable flinging movements of the contralateral limbs, which is precisely what one expects when the brake is destroyed.

2.2.4.4 THE GLOBUS PALLIDUS, INTERNAL SEGMENT (GPi), AND THE SUBSTANTIA NIGRA PARS RETICULATA (SNr)

Plain language: the output nozzle of the basal ganglia.

These two structures are functionally one, separated only by a passing fibre bundle.

INPUTS (SUBSCRIBE): From the striatum (direct pathway, inhibitory).

From the STN (excitatory).

From the GPe (inhibitory).

OUTPUTS (PUBLISH)

• To the ventral anterior (VA) and ventral lateral (VL) nuclei of the thalamus, by way of the pallidothalamic tract, which comprises the ansa lenticularis and the lenticular fasciculus, converging as the thalamic fasciculus.

GABAergic and tonically inhibitory.

• To the pedunculopontine nucleus and the superior colliculus (from SNr), controlling locomotion and eye movements respectively.

FUNCTION: Tonic inhibition of the thalamus, released selectively to permit a chosen action.

This is the structure lesioned in a pallidotomy, and stimulated in deep brain stimulation for dystonia.

2.2.4.5 THE SUBSTANTIA NIGRA PARS COMPACTA (SNc)

INPUTS (SUBSCRIBE): From the striatum, the STN, and the pedunculopontine nucleus.

From the lateral habenula, an inhibitory input carrying the negative reward signal.

OUTPUTS (PUBLISH): To the caudate and putamen, by way of the nigrostriatal dopamine pathway.

FUNCTION AND FACULTY: Supplies the dopamine that biases action selection toward "go." Its neurons encode reward prediction error: they fire above baseline when an outcome is better than expected, and pause when it is worse.

Progressive degeneration of these neurons produces Parkinson's disease, whose cardinal signs (bradykinesia, rigidity, resting tremor) do not appear until roughly sixty to eighty percent of them are already dead.

2.2.5 THE NUCLEUS ACCUMBENS (SHELL AND CORE) AND THE VENTRAL PALLIDUM

Plain language: where wanting happens.

The final common pathway of motivation, and the structure hijacked by every drug of abuse.

  https://www.linkedin.com/in/d-r-dison

INPUTS (SUBSCRIBE)

• From the ventral tegmental area, by way of the mesolimbic dopamine pathway.

This is the pathway on which cocaine, amphetamine, nicotine, opioids, and alcohol all converge.

• From the basolateral amygdala.

Cargo: emotional value.

• From the hippocampus, by way of the fornix.

Cargo: context.

• From the medial prefrontal cortex and OFC.

Cargo: goals and valuation.

OUTPUTS (PUBLISH): To the ventral pallidum, and thence to the mediodorsal thalamus and back to the prefrontal cortex, completing the limbic loop.

To the lateral hypothalamus.

To the VTA.

FUNCTION AND FACULTY: Incentive salience, meaning the attribution of "wanting" to a stimulus, which the literature carefully distinguishes from "liking"; reward-based learning; and the translation of motivation into action.

The shell mediates unconditioned reward; the core mediates learned, cue-driven reward-seeking, which is the substrate of relapse in addiction.

SYSTEM MEMBERSHIP: The mesolimbic reward system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE-and-COMPARATOR (limbic action gate). State: incentive-salience activation. Rule: gates motivated behavior by value; assigns wanting.

DRIVES: Translates drive-relevant value into approach; the wanting stage of every drive.

PLASTICITY: Strong dopamine-gated three-factor plasticity; the site hijacked in addiction.

MODULATION: Dopamine (mesolimbic), opioids.

2.2.6 THE CLAUSTRUM

Plain language: a thin sheet of neurons beneath the cortex, connected to almost everything, whose function remains genuinely unknown.

Francis Crick spent the last day of his life working on a paper about it.

INPUTS (SUBSCRIBE): From nearly every cortical region, reciprocally.

  https://www.linkedin.com/in/d-r-dison

OUTPUTS (PUBLISH): To nearly every cortical region, reciprocally.

Recent connectomic work finds that claustral input to a given cortical area can equal or exceed the thalamic input to that area, which is an extraordinary and underappreciated fact.

FUNCTION AND FACULTY: Unresolved.

The leading hypotheses are the binding of information across cortical modalities into a unified percept, and the coordination of cortex-wide slow-wave activity.

Crick and Koch proposed it as a candidate conductor of the orchestra of consciousness.

Focal electrical stimulation of the claustrum in one human patient reversibly abolished consciousness, which is suggestive but is a single case.

SYSTEM MEMBERSHIP: Undetermined; possibly a cortex-wide integrative hub (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: Hypothesized ATTRACTOR/BROADCAST integrator (function unresolved). State: undetermined. Rule: candidate cortex-wide binder/synchronizer; modeled provisionally as a global coherence hub.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Unknown; modeled as fixed pending evidence.

MODULATION: Unknown.

2.2.7 THE BASAL FOREBRAIN CHOLINERGIC COMPLEX (NUCLEUS BASALIS OF MEYNERT, MEDIAL SEPTAL NUCLEUS, DIAGONAL BAND OF BROCA)

Plain language: the brain's attention supply.

It releases acetylcholine across the entire cortex, and its death is a major event in Alzheimer's disease.

INPUTS (SUBSCRIBE): From the amygdala, the hypothalamus (including orexin neurons), the VTA, the brainstem cholinergic and monoaminergic nuclei, and the prefrontal cortex.

OUTPUTS (PUBLISH)

• Nucleus basalis of Meynert: to the entire cerebral cortex, by way of diffuse cholinergic projections.

This is a volume-transmission projection: it is a broadcast, not an address.

• Medial septal nucleus and diagonal band: to the hippocampus, by way of the fornix.

Cargo: acetylcholine and the theta rhythm pacemaker signal.

FUNCTION AND FACULTY: Cortical arousal, attention, signal-to-noise enhancement, and the gating of cortical plasticity.

Acetylcholine essentially tells the cortex "this is worth learning." Degeneration of the nucleus basalis is an early and severe feature of Alzheimer's disease, and it is the rationale for the cholinesterase-inhibitor drugs, which attempt to prop up a failing broadcast.

SYSTEM MEMBERSHIP: The ascending arousal system and the attention system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR-and-GATE (speech assembly). State: activation of the phonological-motor plan. Rule: sequences phonology into an articulatory program handed to motor cortex.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Language-learning plasticity, strongest in the critical period.

MODULATION: Dopamine.

2.2.8 THE OLFACTORY BULB

Plain language: the nose's piece of brain, and the one sensory structure that does not report to the thalamus first.

INPUTS (SUBSCRIBE): From the olfactory receptor neurons of the nasal epithelium, by way of the olfactory nerve (cranial nerve I), whose axons pass through the cribriform plate and terminate in the glomeruli.

OUTPUTS (PUBLISH): To the piriform cortex (the primary olfactory cortex), the amygdala, and the entorhinal cortex, directly, by way of the lateral olfactory tract.

No thalamic relay intervenes.

FUNCTION AND FACULTY: Odour detection and discrimination.

Its anatomical peculiarity, the absence of an obligatory thalamic relay, is the reason smell has such an unusually direct and involuntary access to emotion and to autobiographical memory: the olfactory tract reaches the amygdala and the entorhinal cortex in two synapses from the outside world.

SYSTEM MEMBERSHIP: The olfactory system (Layer 10).

SPECIFICATION BLOCK

  https://www.linkedin.com/in/d-r-dison

DYNAMICS: RELAY-and-OSCILLATOR (odor coding), no thalamic relay. State: activation per glomerular channel; gamma phase. Rule: transforms receptor input into odor codes, projecting directly to cortex and amygdala.

DRIVES: Feeds appetite and threat drives directly (food and danger smells).

PLASTICITY: Ongoing adult plasticity (rare neurogenesis).

MODULATION: Acetylcholine, norepinephrine.

2.2.9 THE PIRIFORM CORTEX (PRIMARY OLFACTORY CORTEX)

Plain language: the brain's smell cortex, the first cortical stage of olfaction, and the one sensory cortex that does not receive its input by way of the thalamus.

INPUTS (SUBSCRIBE): From the olfactory bulb directly, by way of the lateral olfactory tract, with no thalamic relay; from dense recurrent associative connections within the piriform itself; and from the orbitofrontal cortex and amygdala as feedback.

OUTPUTS (PUBLISH): To the orbitofrontal cortex, both directly and by way of the mediodorsal thalamus; to the amygdala; to the entorhinal cortex; and to the hypothalamus.

FUNCTION AND FACULTY: The formation of odour objects.

It takes the combinatorial glomerular code published by the olfactory bulb and, through its dense recurrent auto-associative connectivity, binds a set of activated receptor channels into a single recognizable smell, and completes a partial or noisy input to a stored odour, a content-addressable memory for scent.

Its direct, thalamus-free access to the amygdala and entorhinal cortex is the anatomical reason odour reaches emotion and autobiographical memory with so little mediation.

SYSTEM MEMBERSHIP: The olfactory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: ATTRACTOR (an auto-associative odour memory). State: activation across odour-object assemblies. Rule: binds the bulb's combinatorial code into an odour object and completes partial patterns through recurrent connectivity, an attractor network of the same family as hippocampal CA3.

DRIVES: Feeds the appetite and threat drives directly, through food and danger odours.

PLASTICITY: Strong experience-dependent plasticity; odour learning and discrimination.

MODULATION: Acetylcholine for encoding, and norepinephrine.

2.2.10 THE OLFACTORY TUBERCLE (OT)

Plain language: a hybrid of striatum and olfactory cortex that ties smell to reward and motivated approach.

  https://www.linkedin.com/in/d-r-dison

INPUTS (SUBSCRIBE): the olfactory bulb and piriform cortex; dopamine from the ventral tegmental area and substantia nigra.

OUTPUTS (PUBLISH): the ventral pallidum; reciprocal olfactory areas.

FUNCTION AND FACULTY: odour-guided motivated behaviour and reward learning, a ventral-striatal node for smell.

SYSTEM MEMBERSHIP: the reward and olfactory systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE (odour-reward). State: odour incentive value. Rule: gates approach to rewarding smells.

MODULATION: dopamine.

2.2.11 THE ANTERIOR OLFACTORY NUCLEUS (AON)

Plain language: the first associative olfactory relay, and the bridge that shares smell between the two hemispheres.

INPUTS (SUBSCRIBE): the olfactory bulb.

OUTPUTS (PUBLISH): the piriform cortex, olfactory tubercle, and entorhinal cortex; the contralateral bulb by the anterior commissure.

FUNCTION AND FACULTY: associative and interhemispheric integration of odour, and state-dependent gating of smell.

SYSTEM MEMBERSHIP: the olfactory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-GATE (odour). State: odour context. Rule: integrates and gates olfactory signals across hemispheres.

MODULATION: acetylcholine, norepinephrine.

2.2.12 THE ENDOPIRIFORM NUCLEUS

Plain language: a deep neighbour of the claustrum that binds olfactory and limbic information and can seed seizures.

INPUTS (SUBSCRIBE): the piriform and olfactory cortices; limbic cortex.

OUTPUTS (PUBLISH): the ventral and limbic cortices; the ventral hippocampus.

FUNCTION AND FACULTY: multisensory and limbic integration, and a recognition-memory node; highly epileptogenic.

SYSTEM MEMBERSHIP: the olfactory-limbic system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (olfactory-limbic). State: bound percept. Rule: integrates smell with limbic context.

  https://www.linkedin.com/in/d-r-dison

MODULATION: acetylcholine.

2.2.13 THE TAENIA TECTA

Plain language: an anterior olfactory-limbic cortex continuous with the hippocampal formation.

INPUTS (SUBSCRIBE): the olfactory bulb and anterior olfactory nucleus.

OUTPUTS (PUBLISH): olfactory-limbic and hippocampal-continuation regions.

FUNCTION AND FACULTY: olfactory-limbic integration at the front of the limbic system.

SYSTEM MEMBERSHIP: the olfactory-limbic system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (olfactory-limbic). State: olfactory-limbic activation. Rule: relays smell into the limbic continuum.

MODULATION: acetylcholine.

2.2.14 THE ISLANDS OF CALLEJA

Plain language: clusters of tiny dopamine-sensitive granule cells in the ventral striatum tied to grooming and mood.

INPUTS (SUBSCRIBE): dopamine from the ventral tegmental area and substantia nigra; local cholinergic and serotonergic input.

OUTPUTS (PUBLISH): local ventral-striatal circuits.

FUNCTION AND FACULTY: modulation of grooming, repetitive behaviour, and mood state.

SYSTEM MEMBERSHIP: the reward and affect systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: MODULATOR (ventral-striatal). State: local dopaminergic tone. Rule: biases repetitive motivated behaviour.

MODULATION: dopamine, serotonin.

2.2.15 THE LATERAL SEPTAL NUCLEUS (LS)

Plain language: the relay that carries the hippocampus's word to the hypothalamus, joining memory and space to emotion and drive.

INPUTS (SUBSCRIBE): the hippocampus; the amygdala; the hypothalamus; midline thalamus.

OUTPUTS (PUBLISH): the hypothalamus (preoptic, paraventricular, lateral areas); the nucleus accumbens; the bed nucleus of the stria terminalis.

FUNCTION AND FACULTY: integrating spatial and emotional information to regulate social behaviour, feeding, stress, and reward.

SYSTEM MEMBERSHIP: the limbic system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-GATE (limbic). State: integrated limbic signal. Rule: relays hippocampal output to drive centres.

MODULATION: serotonin, dopamine, vasopressin, oxytocin.

2.3 THE DIENCEPHALON

2.3.1 THE THALAMUS

Plain language: the gateway to consciousness.

With the sole exception of smell, nothing reaches the cortex without passing through here first.

But calling it a relay is a serious understatement, and the modern literature has spent twenty years correcting that word.

The correction matters, and it is best stated as a fact about counting.

If the thalamus were merely a relay, the fibres running from thalamus to cortex would outnumber those running back.

They do not.

The corticothalamic feedback projection outnumbers the thalamocortical forward projection by roughly ten to one.

The cortex is talking to the thalamus far more than the thalamus is talking to the cortex, and any account of the thalamus as a passive waystation cannot accommodate that ratio.

THE DRIVER AND MODULATOR DISTINCTION

The single most useful organizing principle for the thalamus is the distinction between drivers and modulators.

A driver is the input that carries the message: it determines what the relay cell transmits.

A modulator is an input that adjusts how faithfully, how strongly, or whether at all that message is passed on.

Drivers are few and powerful; modulators are numerous and adjust the gain.

FIRST-ORDER AND HIGHER-ORDER NUCLEI

A first-order nucleus receives its driver input from a subcortical source and delivers a message to cortex for the first time.

The lateral geniculate nucleus is the exemplar: its driver is the retina.

A higher-order nucleus receives its driver input from cortical lamina 5, and relays it onward to a different cortical area.

The pulvinar is the exemplar.

  https://www.linkedin.com/in/d-r-dison

This means a substantial portion of cortex-to-cortex communication does not travel directly from cortex to cortex at all: it goes down to the thalamus and back up.

This transthalamic corticocortical route is one of the most important and least popularized facts in modern systems neuroscience.

2.3.1.1 LATERAL GENICULATE NUCLEUS (LGN)

INPUTS (SUBSCRIBE): Driver: the retina, by way of the optic tract.

Modulators: V1 lamina 6, by way of the corticogeniculate feedback projection; the thalamic reticular nucleus; the brainstem cholinergic nuclei.

OUTPUTS (PUBLISH): To V1, by way of the optic radiation, terminating in lamina 4C.

FUNCTION: The visual relay, organized into magnocellular laminae (motion, low detail, fast), parvocellular laminae (fine detail, colour, slow), and koniocellular laminae.

First-order.

2.3.1.2 MEDIAL GENICULATE NUCLEUS (MGN)

INPUTS (SUBSCRIBE): Driver, ventral division: the inferior colliculus, by way of the brachium of the inferior colliculus.

Driver, dorsal division: cortical lamina 5 (higher-order).

OUTPUTS (PUBLISH): To A1, by way of the acoustic radiation.

Also, from the medial division, directly to the amygdala, forming the fast subcortical fear route.

FUNCTION: The auditory relay.

First-order (ventral) and higher-order (dorsal).

2.3.1.3 VENTRAL POSTEROLATERAL (VPL) AND VENTRAL POSTEROMEDIAL (VPM) NUCLEI

INPUTS (SUBSCRIBE): VPL driver: the medial lemniscus (fine touch, proprioception) and the spinothalamic tract (pain, temperature) from the body.

VPM driver: the trigeminal lemniscus, from the face.

OUTPUTS (PUBLISH): To S1, by way of the thalamocortical somatosensory radiation.

FUNCTION: The somatosensory relay, somatotopically organized.

First-order.

2.3.1.4 VENTRAL ANTERIOR (VA) AND VENTRAL LATERAL (VL) NUCLEI

  https://www.linkedin.com/in/d-r-dison

INPUTS (SUBSCRIBE): From the GPi and SNr, by way of the pallidothalamic tract (inhibitory).

From the deep cerebellar nuclei, by way of the dentato-thalamic tract (excitatory).

OUTPUTS (PUBLISH): To the primary motor cortex, premotor cortex, and SMA.

FUNCTION: The motor thalamus.

This is the single most consequential convergence point in the motor system: it is the one place where the basal ganglia and the cerebellum, two entirely separate motor learning machines, deliver their outputs onto the same target.

VL is a principal target of deep brain stimulation for essential tremor.

2.3.1.5 MEDIODORSAL NUCLEUS (MD)

INPUTS (SUBSCRIBE): From the amygdala, the OFC, the olfactory cortex, the SNr, and the ventral pallidum.

OUTPUTS (PUBLISH): To the prefrontal cortex, by way of the anterior thalamic radiation.

The prefrontal cortex is classically defined as that cortex which receives MD projections.

FUNCTION: The prefrontal relay.

Working memory, executive function, and, notably, the maintenance of cognitive representations across delays.

It is severely damaged in Korsakoff's syndrome, and its destruction contributes to the confabulation seen in that disorder.

2.3.1.6 ANTERIOR THALAMIC NUCLEI

INPUTS (SUBSCRIBE): From the mammillary bodies, by way of the mammillothalamic tract.

From the subiculum, by way of the fornix.

OUTPUTS (PUBLISH): To the cingulate cortex, particularly the posterior cingulate and retrosplenial cortex.

FUNCTION: A core relay of the Papez circuit, and therefore of episodic memory and spatial orientation.

Contains head-direction cells.

Its damage, along with the mammillary bodies, is the anatomical basis of the amnesia of Korsakoff's syndrome, which is caused by thiamine deficiency, most commonly in chronic alcoholism.

2.3.1.7 PULVINAR AND LATERAL POSTERIOR NUCLEUS

INPUTS (SUBSCRIBE): Driver: cortical lamina 5, from visual and parietal cortex (higher-order).

  https://www.linkedin.com/in/d-r-dison

Also from the superior colliculus.

OUTPUTS (PUBLISH): To the visual, parietal, and temporal association cortices.

FUNCTION: The largest nucleus of the human thalamus, and the paradigmatic higher-order nucleus.

It mediates transthalamic corticocortical communication, coordinates the timing of information transfer between cortical areas, and is central to visual attention.

Damage produces deficits in attending to salient stimuli.

It is also a component of the surviving pathway that mediates blindsight.

2.3.1.8 INTRALAMINAR NUCLEI (CENTROMEDIAN AND PARAFASCICULAR)

INPUTS (SUBSCRIBE): From the reticular formation, the spinothalamic tract, and the GPi.

OUTPUTS (PUBLISH): Diffusely to the cortex, and densely to the striatum by way of the thalamostriatal projection.

FUNCTION: Arousal, consciousness, and attentional orienting.

These nuclei are components of the ascending reticular activating system, and bilateral damage to them produces coma.

2.3.1.9 THE THALAMIC RETICULAR NUCLEUS (TRN)

Plain language: a thin sheet of purely inhibitory cells wrapped around the thalamus like a shell.

Every signal passing between thalamus and cortex passes through it, and it is the switchboard operator of attention.

The TRN deserves separate and emphatic treatment, because it is the only thalamic nucleus that sends no output to the cortex at all.

INPUTS (SUBSCRIBE): From collaterals of every thalamocortical fibre passing outward, and from every corticothalamic fibre passing inward.

From the basal forebrain, the brainstem cholinergic nuclei, and, importantly, from the GPe and the prefrontal cortex.

OUTPUTS (PUBLISH): Exclusively to the thalamic relay nuclei.

GABAergic and inhibitory.

It projects to the thalamus and nowhere else.

FUNCTION AND FACULTY: Selective attention, implemented by inhibitory gating.

Crick called it the "guardian of the gateway": if the thalamus is the gate to the cortex, the TRN decides what passes.

  https://www.linkedin.com/in/d-r-dison

The prefrontal cortex projects to the TRN to suppress the thalamic relay of distracting inputs, which is the anatomical mechanism of top-down attentional selection.

The TRN also generates sleep spindles, and its dysfunction is implicated in the attentional deficits of schizophrenia and in absence epilepsy.

SYSTEM MEMBERSHIP: The attention system and the sleep-wake system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: OSCILLATOR-and-GATE. State: inhibitory activation per thalamic channel; spindle phase. Rule: rhythmic inhibition gates which relays reach cortex; prefrontal input biases toward relevant channels.

DRIVES: Serves attention (whichever drive dominates).

PLASTICITY: Limited; gating set top-down rather than by local learning.

MODULATION: Acetylcholine switches spindle (sleep) vs relay-gating (wake) modes.

2.3.1.10 THE NUCLEUS REUNIENS (RE) AND THE MIDLINE THALAMUS

Plain language: the switchboard of the hippocampal-prefrontal conversation, the midline thalamic hub memory and planning depend on.

INPUTS (SUBSCRIBE): reciprocal connections with the medial prefrontal cortex and the hippocampus; brainstem and hypothalamic arousal input.

OUTPUTS (PUBLISH): reciprocal to the hippocampus and medial prefrontal cortex.

FUNCTION AND FACULTY: coordinating the dialogue between hippocampus and prefrontal cortex for spatial working memory, memory consolidation, and fear extinction.

SYSTEM MEMBERSHIP: the memory and executive systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-COORDINATOR (hippocampal-prefrontal). State: the routed signal. Rule: synchronizes hippocampus and prefrontal cortex.

MODULATION: acetylcholine, serotonin.

2.3.1.11 THE PARAVENTRICULAR THALAMUS (PVT)

Plain language: the midline thalamic alarm-and-motivation node, tying arousal and stress to what the animal wants.

INPUTS (SUBSCRIBE): the hypothalamus (orexin from the lateral area, the paraventricular and suprachiasmatic nuclei); the brainstem.

OUTPUTS (PUBLISH): the nucleus accumbens shell, the amygdala, the bed nucleus of the stria terminalis, and the prefrontal and cingulate cortex.

  https://www.linkedin.com/in/d-r-dison

FUNCTION AND FACULTY: integrating arousal, circadian state, and stress with motivation, salience, reward-seeking, and threat avoidance.

SYSTEM MEMBERSHIP: the salience and motivation systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE-and-BROADCAST (state to motivation). State: arousal-and-stress level. Rule: gates motivated behaviour by internal state.

MODULATION: orexin, dopamine, serotonin.

2.3.1.12 THE LATERAL DORSAL NUCLEUS (LD)

Plain language: a limbic thalamic nucleus carrying head-direction and spatial signals for navigation.

INPUTS (SUBSCRIBE): the pretectum and visual system; the hippocampal formation; the cingulate.

OUTPUTS (PUBLISH): the cingulate, retrosplenial, and parahippocampal cortex; the hippocampus.

FUNCTION AND FACULTY: spatial navigation and the limbic head-direction system, an adjunct to the anterior thalamic memory circuit.

SYSTEM MEMBERSHIP: the navigation and memory systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY (spatial-limbic). State: heading and place signal. Rule: relays spatial signals to limbic cortex.

MODULATION: acetylcholine.

2.3.1.13 THE POSTERIOR VENTRAL MEDIAL NUCLEUS (VMpo)

Plain language: the dedicated pain-and-body-state relay, the thalamic gateway to feeling the condition of the body.

INPUTS (SUBSCRIBE): lamina-one neurons of the spinal and trigeminal dorsal horn.

OUTPUTS (PUBLISH): the dorsal posterior insula (the interoceptive cortex).

FUNCTION AND FACULTY: relaying pain, temperature, and interoceptive signals that become the felt state of the body.

SYSTEM MEMBERSHIP: the interoceptive system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY (interoceptive). State: the body-state signal. Rule: relays felt bodily condition to the insula.

MODULATION: serotonin, norepinephrine.

2.3.2 THE HYPOTHALAMUS

  https://www.linkedin.com/in/d-r-dison

Plain language: a structure the size of an almond that runs your body.

Temperature, hunger, thirst, sleep, sex, stress, and the entire endocrine system are all governed from here.

It is the single point at which the nervous system and the endocrine system are not two systems but one.

It is essential to the framework advanced in the companion volume that the hypothalamus be recognized for what it is: the physical location of the neuroendocrine interface.

Its magnocellular neurons are simultaneously Layer 6 neurons and Layer 10 endocrine organs.

Any account that puts the nervous system inside the hierarchy and the endocrine system outside it must cut these cells in half.

GENERAL INPUTS (SUBSCRIBE)

• From the retina, directly, by way of the retinohypothalamic tract, to the suprachiasmatic nucleus.

Cargo: ambient light.

This is a dedicated, private line from the eye to the clock.

• From the hippocampus, by way of the fornix, to the mammillary bodies.

• From the amygdala, by way of the stria terminalis and the ventral amygdalofugal pathway.

• From the brainstem, by way of the dorsal longitudinal fasciculus and the medial forebrain bundle.

• From the nucleus tractus solitarius, by way of ascending visceral afferents.

Cargo: the state of the viscera.

• From the bloodstream, directly, at the circumventricular organs, which lack a blood-brain barrier.

Cargo: leptin, ghrelin, glucose, osmolality, cortisol.

This is a Layer-Skipping Channel in the sense of the companion volume.

GENERAL OUTPUTS (PUBLISH)

• To the posterior pituitary, by way of the hypothalamo-hypophyseal tract (the supraopticohypophyseal and paraventriculohypophyseal tracts).

Cargo: oxytocin and vasopressin, released directly into the systemic bloodstream.

• To the anterior pituitary, by way of the hypophyseal portal system, a dedicated private capillary bed.

Cargo: the releasing and inhibiting hormones.

  https://www.linkedin.com/in/d-r-dison

• To the brainstem and spinal cord autonomic nuclei, by way of the dorsal longitudinal fasciculus.

Cargo: autonomic commands.

• To the thalamus and cortex, by way of the mammillothalamic tract and diffuse projections.

THE PRINCIPAL HYPOTHALAMIC NUCLEI

• Suprachiasmatic nucleus (SCN).

Inputs (SUBSCRIBE): the retina, directly, by the retinohypothalamic tract.

Outputs (PUBLISH): to the pineal gland by a multisynaptic route through the paraventricular nucleus and the superior cervical ganglion; to other hypothalamic nuclei.

Function: the master circadian clock.

It is the reference oscillator to which every peripheral clock in the body is entrained.

• Paraventricular nucleus (PVN).

Inputs (SUBSCRIBE): the arcuate nucleus, the SCN, the amygdala, the BNST, the brainstem.

Outputs (PUBLISH): parvocellular neurons to the median eminence and thence the anterior pituitary, releasing corticotropin-releasing hormone; magnocellular neurons to the posterior pituitary, releasing oxytocin and vasopressin; direct projections to the brainstem and spinal cord.

Function: the apex of the HPA stress axis, and the principal integrator of stress, fluid balance, and autonomic outflow.

• Supraoptic nucleus (SON).

Outputs (PUBLISH): magnocellular projections to the posterior pituitary.

Function: the principal source of circulating vasopressin, governing water retention.

• Arcuate nucleus (ARC).

Inputs (SUBSCRIBE): the bloodstream directly, by way of the adjacent median eminence, a circumventricular organ.

Cargo: leptin from fat, ghrelin from the stomach, insulin, glucose.

Outputs (PUBLISH): to the PVN and the lateral hypothalamus.

Function: the energy sensor of the brain.

It contains two opposed populations: the AgRP/NPY neurons, which drive hunger, and the POMC neurons, which drive satiety.

• Lateral hypothalamic area (LHA).

  https://www.linkedin.com/in/d-r-dison

Function: the classical "feeding centre." Contains the orexin (hypocretin) neurons that stabilize wakefulness.

Loss of orexin neurons causes narcolepsy with cataplexy.

• Ventromedial nucleus (VMH).

Function: the classical "satiety centre." Also involved in defensive and reproductive behaviour.

• Mammillary bodies.

Inputs (SUBSCRIBE): the hippocampus, by way of the fornix.

Outputs (PUBLISH): the anterior thalamic nuclei, by way of the mammillothalamic tract; the brainstem, by way of the mammillotegmental tract.

Function: a relay in the Papez circuit; essential to episodic memory.

Their degeneration is the classical lesion of Korsakoff's syndrome.

• Preoptic area (including the ventrolateral preoptic nucleus, VLPO).

Function: thermoregulation, and sleep initiation.

The VLPO is the sleep switch: it inhibits the ascending arousal system, and the arousal system inhibits it back, producing a bistable flip-flop that makes wake and sleep two stable states with a fast transition between them.

• Tuberomammillary nucleus (TMN).

Outputs (PUBLISH): histaminergic projections to the entire cortex.

Function: wakefulness.

Antihistamines cause drowsiness by blocking this system.

FUNCTION AND FACULTY OF THE HYPOTHALAMUS AS A WHOLE

Homeostasis in its entirety: thermoregulation, osmoregulation, energy balance, circadian timing, sleep and wakefulness, the stress response, reproduction, and defensive and appetitive behaviour.

It is the interface at which Layer 12 needs become Layer 3 chemistry.

SYSTEM MEMBERSHIP: All neuroendocrine axes; the autonomic system; the circadian system; the sleep-wake system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: Assign per Chapter 8 archetype (RELAY / INTEGRATOR / OSCILLATOR / ATTRACTOR / GATE / COMPARATOR). State: activation; gain. Rule: transforms input per its function above.

DRIVES: Not Applicable, for the reason that this construct does not monitor a bodily variable and so defends no set-point; it would serve a drive only if it monitored such a variable, which it does not.

  https://www.linkedin.com/in/d-r-dison

PLASTICITY: Weights change by the three-factor rule of Chapter 10 where the region learns; otherwise Not Applicable, for the reason that a region without modifiable connections has no weights to update.

MODULATION: Set by the neuromodulators named in its inputs.

THE PUBLISHER NUCLEI OF THE HYPOTHALAMUS, IN FULL

The roster above lists the principal hypothalamic nuclei in brief. Five of them are not merely sub-regions but PUBLISHERS in the publish-and-subscribe sense of Chapter 7: each is the source of a distinct chemical broadcast that regions elsewhere subscribe to, and each will therefore be built as a distinct named module with its own identity and behaviour. Because their molecules are catalogued in Chapter 7, these five are given the full construct treatment here, as their cousins receive elsewhere in this chapter.

2.3.2.1 PARAVENTRICULAR NUCLEUS (PVN)

Plain language: the hypothalamus's stress-and-release command post, the apex of the body's stress axis and the source of three major peptide broadcasts at once.

INPUTS (SUBSCRIBE): From the arcuate nucleus, carrying energy state; from the suprachiasmatic nucleus, carrying time of day; from the amygdala and the bed nucleus of the stria terminalis, carrying threat; and from the brainstem, carrying visceral state.

OUTPUTS (PUBLISH): Parvocellular neurons to the median eminence and thence the anterior pituitary, publishing corticotropin-releasing hormone, the apex trigger of the stress axis; magnocellular neurons to the posterior pituitary, publishing oxytocin and vasopressin directly into the bloodstream; and direct projections to the brainstem and spinal-cord autonomic nuclei.

FUNCTION AND FACULTY: The apex of the hypothalamic-pituitary-adrenal stress axis, and the principal integrator of stress, fluid balance, and autonomic outflow.

In the publish-and-subscribe terms of Chapter 7, the paraventricular nucleus is the publisher of corticotropin-releasing hormone, oxytocin, and vasopressin, and its dendritic release within the brain and its endocrine release into the blood are independently regulated, so it publishes on two channels at once.

SYSTEM MEMBERSHIP: The stress axis, the neuroendocrine system, and the autonomic system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR-and-PUBLISHER. State: stress-mobilization level; fluid and osmotic set-point error. Rule: integrates threat, energy, time, and visceral state, and publishes corticotropin-releasing hormone, oxytocin, and vasopressin accordingly.

DRIVES: Serves the safety/stress drive and the fluid-balance drive; the apex publisher of the stress response.

  https://www.linkedin.com/in/d-r-dison

PLASTICITY: Chronic stress shifts its corticotropin-releasing-hormone set-point (allostasis).

MODULATION: Set by the neuromodulators named in its inputs; it is itself a publisher of modulators.

2.3.2.2 SUPRAOPTIC NUCLEUS (SON)

Plain language: the body's water-balance publisher, the principal source of the vasopressin that circulates in the blood.

INPUTS (SUBSCRIBE): From osmoreceptors that sense blood concentration, by way of the lamina terminalis circumventricular organs; and from the brainstem.

OUTPUTS (PUBLISH): Magnocellular projections to the posterior pituitary, publishing vasopressin, and some oxytocin, directly into the systemic bloodstream.

FUNCTION AND FACULTY: The governance of water retention.

It publishes vasopressin when the blood grows too concentrated, driving the kidney to conserve water; in publish-and-subscribe terms it is a nearly single-channel publisher, dedicated to the vasopressin broadcast that defends osmotic balance.

SYSTEM MEMBERSHIP: The neuroendocrine osmoregulatory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: COMPARATOR-and-PUBLISHER. State: plasma osmolality error. Rule: publishes vasopressin in proportion to the deviation of blood concentration from its set-point.

DRIVES: Serves the thirst and fluid-balance drive; monitors a bodily variable (osmolality).

PLASTICITY: Limited; the set-point can shift under chronic osmotic challenge.

MODULATION: Angiotensin and brainstem afferents.

2.3.2.3 ARCUATE NUCLEUS (ARC)

Plain language: the brain's fuel gauge, the energy sensor that reads the blood directly and publishes hunger or satiety.

INPUTS (SUBSCRIBE): From the bloodstream directly, at the adjacent median eminence, a circumventricular organ lacking a blood-brain barrier. Cargo: leptin from fat, ghrelin from the stomach, insulin, and glucose. This is a Layer-Skipping Channel in the sense of the companion volume, a bodily chemical reaching a Layer 9 sensor with no intervening neural relay.

OUTPUTS (PUBLISH): To the paraventricular nucleus and the lateral hypothalamic area; and, by way of the tuberoinfundibular pathway, dopamine to the anterior pituitary that holds prolactin release in check.

FUNCTION AND FACULTY: The energy sensor of the brain.

  https://www.linkedin.com/in/d-r-dison

It contains two opposed publisher populations, the AgRP/NPY neurons, which publish a hunger signal, and the POMC neurons, which publish a satiety signal; it is also the tuberoinfundibular publisher of dopamine in a pure endocrine role.

SYSTEM MEMBERSHIP: The energy-homeostasis and neuroendocrine systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: COMPARATOR-and-PUBLISHER (opposed populations). State: energy-balance error. Rule: reads circulating energy signals and publishes hunger (AgRP/NPY) or satiety (POMC), plus tonic prolactin-inhibiting dopamine.

DRIVES: Serves the hunger and energy drive directly; monitors a bodily variable (energy stores).

PLASTICITY: Set-points shift with chronic over- or under-nutrition, a source of the body's defended weight.

MODULATION: Leptin, ghrelin, and insulin as first messengers.

2.3.2.4 LATERAL HYPOTHALAMIC AREA (LHA)

Plain language: the wake-and-want centre, the publisher of orexin, the peptide that holds the brain awake.

INPUTS (SUBSCRIBE): From the limbic system, the prefrontal cortex, the arcuate nucleus carrying energy state, and the brainstem.

OUTPUTS (PUBLISH): Orexin (hypocretin) and melanin-concentrating-hormone projections among the most widespread in the brain, reaching the whole cortex and, decisively, the ascending arousal nuclei, the locus coeruleus, raphe, tuberomammillary nucleus, and basal forebrain.

FUNCTION AND FACULTY: The classical feeding centre, and the stabilizer of wakefulness.

Its orexin neurons publish the signal that holds the sleep-wake flip-flop switch in the waking state and prevents unwanted transitions; loss of these neurons causes narcolepsy with cataplexy.

In publish-and-subscribe terms, the lateral hypothalamic area is the publisher whose orexin broadcast the arousal nuclei subscribe to, adding hysteresis to a bistable switch.

SYSTEM MEMBERSHIP: The sleep-wake and appetitive-motivation systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: PUBLISHER-and-STABILIZER. State: wake-drive and arousal-stabilization level. Rule: publishes orexin to hold the wake state and adds hysteresis to the sleep-wake gate.

DRIVES: Stabilizes the sleep/wake drive and feeds the appetitive drive.

PLASTICITY: Not Applicable, for the reason that this construct is a fixed signalling source, not a site whose weights change with experience.

  https://www.linkedin.com/in/d-r-dison

MODULATION: Energy state, by way of the arcuate nucleus (leptin and ghrelin), and glucose.

2.3.2.5 TUBEROMAMMILLARY NUCLEUS (TMN)

Plain language: the sole histamine publisher of the brain, a few thousand cells that help hold the entire cortex awake.

INPUTS (SUBSCRIBE): From the ventrolateral preoptic nucleus, the sleep switch, which inhibits it; and from the orexin neurons of the lateral hypothalamic area, which excite it.

OUTPUTS (PUBLISH): Histaminergic projections to the entire cortex and to the other ascending arousal nuclei.

FUNCTION AND FACULTY: The maintenance of wakefulness.

It is the brain's only source of histamine, and it publishes that broadcast across the cortex to promote arousal; antihistamines that cross the blood-brain barrier block the H1 subscribers of this broadcast and cause profound drowsiness, a vivid demonstration of the reach-not-size principle, a few thousand cells silenced and consciousness dims.

SYSTEM MEMBERSHIP: The ascending arousal system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: BROADCAST PUBLISHER. State: wakefulness-promoting histamine level. Rule: publishes histamine to the cortex to sustain arousal; silenced by the sleep switch, sustained by orexin.

DRIVES: Serves the sleep/wake drive (wakefulness).

PLASTICITY: Not Applicable, for the reason that this construct is a fixed signalling source, not a site whose weights change with experience.

MODULATION: Inhibited by the ventrolateral preoptic nucleus; excited by orexin.

2.3.3 THE HABENULA (LATERAL AND MEDIAL)

Plain language: the disappointment centre.

It fires when you expected a reward and did not get one, and it punishes the brain's reward systems accordingly.

2.3.3.1 THE HABENULA-MONOAMINE FEEDBACK HUB (PUC)

DEFINITION

The lateral habenula (LHb) considered specifically in its role as the coordinating anatomical node of negative monoamine feedback: the single structure that simultaneously downregulates both dopamine and serotonin when outcomes are worse than expected.

INPUTS (SUBSCRIBE): From the internal globus pallidus, by way of the stria medullaris.

  https://www.linkedin.com/in/d-r-dison

From the medial prefrontal cortex, the lateral hypothalamus, and the limbic system.

OUTPUTS (PUBLISH): To the ventral tegmental area, by way of the fasciculus retroflexus, acting through the rostromedial tegmental nucleus (RMTg) to inhibit dopamine neurons.

To the raphe nuclei, inhibiting serotonin neurons.

PROOF OF EXISTENCE: The LHb is activated by reward omission, by punishment, and by negative prediction errors, and it inhibits both major monoamine systems in consequence.

Its hyperactivity is a well-replicated finding in models of depression, and it is an emerging target for deep brain stimulation in treatment-resistant depression.

PROOF OF NAMELESSNESS: The lateral habenula is named.

The VTA and raphe are named.

Reward prediction error is named.

But the LHb specifically in its role as the unified negative-feedback hub for both monoamine systems at once has never been named as a distinct Layer 9 functional construct.

FUNCTION AND FACULTY: Encoding of negative reward prediction error; behavioural suppression following disappointment; and, when chronically overactive, the anhedonia and behavioural despair characteristic of depression.

SYSTEM MEMBERSHIP: The reward and mood-regulation systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: Assign per Chapter 8 archetype (RELAY / INTEGRATOR / OSCILLATOR / ATTRACTOR / GATE / COMPARATOR). State: activation; gain. Rule: transforms input per its function above.

DRIVES: Not Applicable, for the reason that this construct does not monitor a bodily variable and so defends no set-point; it would serve a drive only if it monitored such a variable, which it does not.

PLASTICITY: Weights change by the three-factor rule of Chapter 10 where the region learns; otherwise Not Applicable, for the reason that a region without modifiable connections has no weights to update.

MODULATION: Set by the neuromodulators named in its inputs.

2.3.4 THE PINEAL GLAND

INPUTS (SUBSCRIBE): From the SCN, by a remarkable and circuitous multisynaptic route: SCN to paraventricular nucleus to the intermediolateral cell column of the spinal cord to the superior cervical ganglion, and thence by sympathetic fibres back up into the skull to the pineal.

  https://www.linkedin.com/in/d-r-dison

OUTPUTS (PUBLISH): Melatonin, released directly into the bloodstream and the cerebrospinal fluid.

FUNCTION AND FACULTY: The transduction of the neural circadian signal into an endocrine one.

Melatonin is the hormone of darkness: it is the body's chemical announcement that night has fallen.

Note that the pineal's output is hormonal, not neural, which makes it a Layer 9 structure whose upward interface is entirely endocrine.

SYSTEM MEMBERSHIP: The circadian system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY transducing time-of-day into hormone. State: melatonin output. Rule: converts the suprachiasmatic clock signal into nocturnal melatonin release.

DRIVES: Serves the sleep/circadian drive.

PLASTICITY: Not Applicable, for the reason that this construct is a fixed transducer whose conversion function does not change with experience; it has no learnable weights.

MODULATION: Norepinephrine (sympathetic drive from the clock).

2.3.5 THE SUBTHALAMUS AND THE ZONA INCERTA THE ZONA INCERTA (ZI)

Plain language: an under-named integrative sheet beneath the thalamus that gates arousal, attention, and movement.

INPUTS (SUBSCRIBE): the cortex, basal ganglia, brainstem, and sensory systems.

OUTPUTS (PUBLISH): the higher-order thalamus, the paraventricular thalamus, the ventral tegmental area, the periaqueductal gray, the superior colliculus, and the cortex.

FUNCTION AND FACULTY: a wide-reaching integrator of sensory processing, arousal, attention, locomotion, feeding, defence, and sleep, gating thalamocortical and thalamostriatal transmission.

SYSTEM MEMBERSHIP: the arousal and attention systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE (thalamocortical/behavioural). State: gating tone. Rule: sets the gain of thalamic and behavioural circuits.

MODULATION: dopamine, acetylcholine, gamma-aminobutyric acid.

2.4 THE BRAINSTEM

The brainstem is the stalk connecting the brain to the spinal cord, comprising the midbrain, the pons, and the medulla.

  https://www.linkedin.com/in/d-r-dison

It contains the nuclei of ten of the twelve cranial nerves, the machinery of breathing and circulation, and the sources of every ascending neuromodulatory system.

It is small, it is crowded, and it is the one part of the nervous system in which a lesion of a few cubic millimetres can be immediately fatal.

2.4.1 THE MIDBRAIN

2.4.1.1 SUPERIOR COLLICULUS

Plain language: the reflex orienting machine.

It turns your eyes and head toward a sudden event before you know what the event was.

INPUTS (SUBSCRIBE): From the retina, directly, by way of a branch of the optic tract.

From V1, by the corticotectal projection.

From the auditory and somatosensory systems.

OUTPUTS (PUBLISH): To the brainstem gaze centres and the oculomotor nuclei.

To the spinal cord, by way of the tectospinal tract, turning the head.

To the pulvinar, and thence to the cortex, forming the surviving route that mediates blindsight.

FUNCTION AND FACULTY: Saccadic eye movement generation and visual orienting.

It contains superimposed retinotopic, auditory, and somatotopic maps in spatial register, so that a sight, a sound, and a touch from the same location activate the same column.

This is one of the most elegant multisensory architectures in the nervous system.

2.4.1.2 INFERIOR COLLICULUS

INPUTS (SUBSCRIBE): From the cochlear nuclei, the superior olivary complex, and the lateral lemniscus.

From A1, by descending corticocollicular projections.

OUTPUTS (PUBLISH): To the medial geniculate nucleus, by way of the brachium of the inferior colliculus.

To the superior colliculus.

FUNCTION: The obligatory midbrain relay of the auditory pathway; integration of binaural cues for sound localization; the acoustic startle reflex.

2.4.1.3 PERIAQUEDUCTAL GRAY (PAG)

2.4.1.3.1 THE PERIAQUEDUCTAL GRAY PAIN MODULATION HUB (PUC)

  https://www.linkedin.com/in/d-r-dison

Plain language: the brain's own morphine dispensary, and the reason a soldier can be wounded and not feel it until the battle ends.

DEFINITION: The PAG considered specifically as the organized hub of the descending pain modulation system.

INPUTS (SUBSCRIBE): From the medial prefrontal cortex, the ACC, and the insula.

From the amygdala (central nucleus).

From the hypothalamus.

OUTPUTS (PUBLISH): To the rostral ventromedial medulla (RVM), and thence, by way of the descending pain modulatory pathway in the dorsolateral funiculus, to the dorsal horn of the spinal cord, where it gates the very first synapse of the pain pathway.

Also to the locus coeruleus.

PROOF OF EXISTENCE: Electrical stimulation of the PAG produces profound analgesia sufficient for surgery.

It is dense in opioid receptors, and it is the principal site of action of morphine's analgesic effect.

The descending pathway is fully characterized.

PROOF OF NAMELESSNESS: The PAG is named.

Descending pain modulation is described.

But the PAG in its specific role as the organizing hub of that system, distinct from its other functions in defensive behaviour and vocalization, has never been named as a Layer 9 functional hub.

FUNCTION AND FACULTY: Descending analgesia; defensive behaviour (freezing, fight, flight); vocalization; and autonomic control of the emotional response.

It is the structure that implements the amygdala's fear command as behaviour.

SYSTEM MEMBERSHIP: The pain modulation system and the defensive behaviour system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE (descending pain and defense). State: analgesia/defense activation. Rule: gates descending pain modulation and selects freeze/flight/fight; opioid-sensitive.

DRIVES: Serves the SAFETY and pain drives; executes the amygdala's threat command.

PLASTICITY: Opioid-gated plasticity of pain thresholds.

MODULATION: Endogenous opioids, serotonin, norepinephrine.

  https://www.linkedin.com/in/d-r-dison

2.4.1.4 VENTRAL TEGMENTAL AREA (VTA)

INPUTS (SUBSCRIBE): From the lateral habenula, by way of the RMTg (inhibitory).

From the prefrontal cortex, the lateral hypothalamus (orexin), the pedunculopontine nucleus, and the nucleus accumbens.

OUTPUTS (PUBLISH)

• To the nucleus accumbens, by the mesolimbic dopamine pathway. Cargo: incentive salience, "wanting."

• To the prefrontal cortex, by the mesocortical dopamine pathway.

Cargo: working memory gain control.

• To the amygdala and hippocampus.

Cargo: a novelty signal that gates memory encoding.

FUNCTION AND FACULTY: Reward prediction error signalling; motivation; and the reinforcement of behaviour.

Every drug of abuse, without exception, increases dopamine in the nucleus accumbens, either directly or by disinhibiting the VTA.

Two refinements to the classic picture are now established and belong in any working model.

First, dopamine speaks in two tenses.

Its fast PHASIC bursts carry the reward prediction error, the moment-to-moment surprise signal: a burst when an outcome beats expectation, a pause when it falls short, and, tellingly, silence when a fully predicted reward arrives, because a predicted reward is not news.

Its slow TONIC level tracks something different, the ongoing value of the current situation, rising steadily as an animal draws near a reward in space or time.

The same molecule thus encodes both the sudden surprise and the slow background sense of how well things are going.

Second, the dopamine neurons are not one population but at least two.

Medial VTA neurons carry the classic value and prediction-error signal.

Lateral VTA neurons and much of the substantia nigra carry a SALIENCE signal that fires for any important event, rewarding or aversive alike, and it is this salience population, not the reward population, that responds to a painful or threatening stimulus.

The single phrase "the dopamine reward signal" therefore conceals two computations in two anatomies, and a model that represents only the first will mispredict every response to threat.

2.4.1.5 RED NUCLEUS

  https://www.linkedin.com/in/d-r-dison

INPUTS (SUBSCRIBE): From the motor cortex, by the corticorubral projection.

From the cerebellum (interposed nuclei), by the superior cerebellar peduncle.

OUTPUTS (PUBLISH): To the spinal cord, by the rubrospinal tract.

To the inferior olive, by the central tegmental tract.

FUNCTION: Flexor facilitation of the upper limb.

In humans this pathway is much reduced relative to other mammals, having been largely superseded by the corticospinal tract, but it is clinically important: decorticate posturing (arms flexed) reflects a lesion above the red nucleus with the rubrospinal tract intact, while decerebrate posturing (arms extended) reflects a lesion below it.

The distinction is a bedside localization of catastrophic injury.

2.4.1.6 THE ROSTROMEDIAL TEGMENTAL NUCLEUS (RMTg)

Plain language: the master brake on dopamine, the inhibitory nucleus that turns disappointment into a stop signal.

INPUTS (SUBSCRIBE): the lateral habenula (its main driver); the extended amygdala; the lateral septum; the periaqueductal gray.

OUTPUTS (PUBLISH): dense inhibition of the dopamine neurons of the ventral tegmental area and substantia nigra; the raphe and cholinergic nuclei.

FUNCTION AND FACULTY: encoding aversion and negative outcomes and braking dopamine, the effector of the anti-reward pathway.

SYSTEM MEMBERSHIP: the reward and aversion system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE (dopamine brake). State: the aversion signal. Rule: inhibits dopamine when outcomes are worse than expected.

MODULATION: gamma-aminobutyric acid output; habenular glutamate input.

2.4.1.7 THE INTERPEDUNCULAR NUCLEUS (IPN)

Plain language: the terminus of the medial habenula's private line, tied to fear, anxiety, and nicotine.

INPUTS (SUBSCRIBE): the medial habenula by the fasciculus retroflexus; the ventral tegmental area.

OUTPUTS (PUBLISH): the raphe, the tegmentum, and the laterodorsal tegmental nucleus.

FUNCTION AND FACULTY: aversive learning, anxiety, mood, and the aversive component of nicotine and addiction.

SYSTEM MEMBERSHIP: the aversion and mood system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY-and-GATE (aversive). State: aversive state. Rule: relays medial-habenular aversion to midbrain nuclei.

MODULATION: acetylcholine, serotonin.

2.4.1.8 THE CUNEIFORM NUCLEUS AND THE MESENCEPHALIC LOCOMOTOR REGION

Plain language: the midbrain go-fast centre for urgent, escape, and defensive locomotion.

INPUTS (SUBSCRIBE): the superior colliculus, periaqueductal gray, substantia nigra, and hypothalamus.

OUTPUTS (PUBLISH): the reticulospinal system driving spinal locomotor circuits; diffuse arousal.

FUNCTION AND FACULTY: commanding high-speed and defensive locomotion and coupling arousal to movement.

SYSTEM MEMBERSHIP: the locomotor and defence systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE (locomotor drive). State: locomotor set-point. Rule: sets locomotor speed and urgency.

MODULATION: glutamate; noradrenergic and cholinergic modulation.

2.4.1.9 THE EDINGER-WESTPHAL NUCLEUS

Plain language: the parasympathetic pupil-and-lens controller, and, in its peptidergic part, a stress-and-consumption centre.

INPUTS (SUBSCRIBE): the olivary pretectal nucleus (light reflex); the periaqueductal gray; accessory oculomotor input.

OUTPUTS (PUBLISH): the ciliary ganglion (pupil constriction, accommodation); central peptidergic projections.

FUNCTION AND FACULTY: the pupillary light reflex and lens accommodation, and, via urocortin, aspects of stress and consummatory behaviour.

SYSTEM MEMBERSHIP: the autonomic and stress systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY (pupillo-autonomic). State: pupil and stress signal. Rule: constricts the pupil and signals stress.

MODULATION: acetylcholine; urocortin.

2.4.1.10 THE CENTRAL MESENCEPHALIC RETICULAR FORMATION (cMRF)

Plain language: the midbrain relay that turns a chosen target into a horizontal-gaze command.

INPUTS (SUBSCRIBE): the superior colliculus; the frontal eye fields.

  https://www.linkedin.com/in/d-r-dison

OUTPUTS (PUBLISH): the paramedian pontine reticular formation (the horizontal gaze centre).

FUNCTION AND FACULTY: conducting saccade commands from the colliculus and frontal eye fields to the pontine gaze generator.

SYSTEM MEMBERSHIP: the oculomotor system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY (saccade routing). State: the gaze command. Rule: routes saccade signals to the gaze centre.

MODULATION: acetylcholine.

2.4.2 THE PONS

2.4.2.1 LOCUS COERULEUS (LC)

Plain language: a tiny blue speck of about fifteen thousand cells per side, whose axons reach essentially the entire brain.

It is the alarm bell.

INPUTS (SUBSCRIBE): From the prefrontal cortex, the amygdala (central nucleus), the nucleus paragigantocellularis, and the orexin neurons of the lateral hypothalamus.

OUTPUTS (PUBLISH): To virtually the whole brain, by way of the ascending noradrenergic bundle: the entire cortex, hippocampus, amygdala, thalamus, cerebellum, and spinal cord.

This is a volume-transmission projection, and it is the paradigm case of a broadcast.

FUNCTION AND FACULTY: Arousal, vigilance, and the reorienting of attention to unexpected and salient events.

Its phasic firing marks a surprising stimulus; its tonic firing sets the overall level of arousal.

It is silent during rapid eye movement sleep.

Its degeneration is an extremely early event in both Alzheimer's and Parkinson's disease, possibly the earliest detectable lesion in either.

2.4.2.2 RAPHE NUCLEI

INPUTS (SUBSCRIBE): From the prefrontal cortex, the lateral habenula (inhibitory), and the hypothalamus.

OUTPUTS (PUBLISH): Serotonergic projections to essentially the entire brain, and descending projections to the dorsal horn of the spinal cord, where they participate in pain gating.

FUNCTION AND FACULTY: Modulation of mood, aggression, impulsivity, appetite, and sleep architecture.

  https://www.linkedin.com/in/d-r-dison

The dorsal and median raphe supply the forebrain; the caudal raphe supply the spinal cord.

This is the system targeted by the selective serotonin reuptake inhibitors.

2.4.2.3 PONTINE NUCLEI

INPUTS (SUBSCRIBE): From the entire cerebral cortex, by way of the massive corticopontine projection.

This projection is enormous: roughly twenty million fibres, dwarfing the one million of the corticospinal tract.

OUTPUTS (PUBLISH): To the contralateral cerebellar cortex, by way of the middle cerebellar peduncle, as mossy fibres.

FUNCTION: The obligatory relay of the cortico-ponto-cerebellar pathway.

It is the structure by which the cortex informs the cerebellum of what it intends to do.

Its sheer size is the anatomical statement of how much the cortex depends on the cerebellum.

2.4.2.4 PEDUNCULOPONTINE AND LATERODORSAL TEGMENTAL NUCLEI (PPN, LDT)

OUTPUTS (PUBLISH): Cholinergic projections to the thalamus (especially the intralaminar nuclei), the basal forebrain, and the basal ganglia.

FUNCTION: Components of the ascending arousal system.

They are active in waking and in rapid eye movement sleep, and they generate the ponto-geniculo-occipital waves of REM sleep.

The PPN also participates in the initiation of locomotion, and is an experimental target for deep brain stimulation in Parkinsonian gait freezing.

2.4.2.5 PARABRACHIAL NUCLEUS

INPUTS (SUBSCRIBE): From the nucleus tractus solitarius.

Cargo: all visceral and taste afferent information.

OUTPUTS (PUBLISH): To the thalamus (VMpo), the hypothalamus, the amygdala, and the insula.

FUNCTION: The obligatory relay of interoceptive information from the brainstem to the forebrain.

It is the structure by which the body's internal state reaches consciousness, and it mediates the aversive quality of nausea, dyspnoea, and pain.

2.4.2.6 TRIGEMINAL NUCLEI

INPUTS (SUBSCRIBE): From the trigeminal ganglion, by way of cranial nerve V.

Cargo: all sensation from the face and the anterior scalp.

OUTPUTS (PUBLISH): To the VPM of the thalamus, by way of the trigeminal lemniscus.

FUNCTION: Facial sensation, and the motor control of mastication.

The spinal trigeminal nucleus carries pain and temperature; the principal sensory nucleus carries fine touch.

2.4.2.6 THE SUPERIOR OLIVARY COMPLEX (SOC)

Plain language: the first place the two ears are compared, the origin of knowing where a sound is.

INPUTS (SUBSCRIBE): both cochlear nuclei.

OUTPUTS (PUBLISH): the inferior colliculus by the lateral lemniscus; the cochlea, by the olivocochlear efferents.

FUNCTION AND FACULTY: binaural sound localization by interaural time and level differences, and efferent control of the cochlea.

SYSTEM MEMBERSHIP: the auditory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: COMPARATOR (binaural). State: interaural difference. Rule: computes sound location from the two ears.

MODULATION: acetylcholine (olivocochlear).

2.4.2.7 THE VESTIBULAR NUCLEI

Plain language: the balance computer of the brainstem, holding the world still as the head moves.

INPUTS (SUBSCRIBE): the vestibular nerve; the cerebellum; visual and proprioceptive signals.

OUTPUTS (PUBLISH): the oculomotor nuclei (the vestibulo-ocular reflex); the spinal cord (posture); the thalamus and cortex; the cerebellum.

FUNCTION AND FACULTY: balance, gaze stabilization, posture, and the sense of spatial orientation and self-motion.

SYSTEM MEMBERSHIP: the vestibular and spatial systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (self-motion). State: head-motion estimate. Rule: stabilizes gaze and posture against head movement.

MODULATION: acetylcholine, histamine.

2.4.2.8 THE NUCLEUS INCERTUS

Plain language: a pontine stress-and-theta node that paces the hippocampus and signals behavioural arousal.

  https://www.linkedin.com/in/d-r-dison

INPUTS (SUBSCRIBE): corticotropin-releasing-hormone and orexin input; the interpeduncular nucleus.

OUTPUTS (PUBLISH): the septum, hippocampus, hypothalamus, amygdala, and prefrontal cortex.

FUNCTION AND FACULTY: modulating arousal, stress, hippocampal theta rhythm, memory, and locomotor speed by way of the peptide relaxin-three.

SYSTEM MEMBERSHIP: the arousal and stress systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: MODULATOR (theta/arousal). State: arousal-stress level. Rule: paces hippocampal theta and signals stress.

MODULATION: relaxin-three, gamma-aminobutyric acid.

2.4.2.9 THE NUCLEUS PREPOSITUS HYPOGLOSSI

Plain language: the neural integrator for horizontal gaze, the circuit that turns a velocity command into a held eye position.

INPUTS (SUBSCRIBE): the gaze centres and vestibular nuclei.

OUTPUTS (PUBLISH): the oculomotor and abducens nuclei; the vestibular nuclei and cerebellum.

FUNCTION AND FACULTY: mathematically integrating eye-velocity signals so the eyes can hold an eccentric position, the gaze-holding memory.

SYSTEM MEMBERSHIP: the oculomotor system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (gaze-holding). State: held eye position. Rule: integrates velocity into position.

MODULATION: acetylcholine.

2.4.3 THE MEDULLA

2.4.3.1 THE VAGAL COMPLEX: THE NUCLEUS TRACTUS SOLITARIUS (NTS) AND THE DORSAL MOTOR NUCLEUS OF THE VAGUS

Plain language: the body's inbox and outbox.

Every report from your heart, lungs, gut, and blood vessels arrives here.

NUCLEUS TRACTUS SOLITARIUS (NTS)

INPUTS (SUBSCRIBE): From the vagus nerve (cranial nerve X), the glossopharyngeal nerve (IX), and the facial nerve (VII), by way of the solitary tract.

Cargo: baroreceptor signals (blood pressure), chemoreceptor signals (blood gases), pulmonary stretch, gastrointestinal distension and nutrient content, and taste.

  https://www.linkedin.com/in/d-r-dison

Roughly eighty percent of vagal fibres are afferent, which is to say that the vagus is overwhelmingly a sensory nerve, a fact that surprises almost everyone who learns it.

OUTPUTS (PUBLISH): To the dorsal motor nucleus of the vagus and the nucleus ambiguus, forming the reflex arcs.

To the parabrachial nucleus, and thence to the hypothalamus, amygdala, and insula.

To the ventrolateral medulla, controlling blood pressure.

FUNCTION: The first central relay of all visceral afference.

It is the entry port of the vagal Layer-Skipping Channel described in the companion volume: a chemical event in the gut lumen, sensed by a neuropod cell, reaches this nucleus in milliseconds.

DORSAL MOTOR NUCLEUS OF THE VAGUS

OUTPUTS (PUBLISH): Parasympathetic preganglionic fibres, by way of the vagus nerve, to the heart, lungs, and the entire gastrointestinal tract to the level of the transverse colon.

FUNCTION: Parasympathetic visceromotor control: "rest and digest."

2.4.3.2 AREA POSTREMA

INPUTS (SUBSCRIBE): From the bloodstream, directly.

The area postrema is a circumventricular organ: it lacks a blood-brain barrier.

OUTPUTS (PUBLISH): To the NTS, and thence to the vomiting centre of the reticular formation.

FUNCTION AND FACULTY: The chemoreceptor trigger zone.

It samples the blood for circulating toxins and initiates vomiting.

It is the anatomical reason that chemotherapy causes nausea, and it is a textbook example of a deliberate hole in the blood-brain barrier: a port, not a leak.

2.4.3.3 INFERIOR OLIVARY NUCLEUS

INPUTS (SUBSCRIBE): From the spinal cord, the red nucleus (by the central tegmental tract), and the cerebral cortex.

OUTPUTS (PUBLISH): To the contralateral cerebellar cortex, by way of the inferior cerebellar peduncle, as the climbing fibres.

Each Purkinje cell receives exactly one climbing fibre, and that single input is so powerful it cannot be ignored.

FUNCTION: The generator of the cerebellar teaching signal.

The climbing fibre carries an error signal: the message that what happened was not what was predicted.

This is the physical substrate of supervised learning in the cerebellum.

2.4.3.4 VENTRAL RESPIRATORY GROUP AND THE PRE-BOTZINGER COMPLEX

INPUTS (SUBSCRIBE): From the NTS (chemoreceptor and stretch information), and from the pons.

OUTPUTS (PUBLISH): To the phrenic motor neurons of the cervical spinal cord, and to the intercostal motor neurons.

FUNCTION AND FACULTY: The generation of the respiratory rhythm.

The pre-Botzinger complex is a central pattern generator: an isolated slice of it, in a dish, continues to produce a breathing rhythm.

Every breath you have ever taken originated here, and its destruction is immediately fatal.

2.4.3.5 RETICULAR FORMATION

INPUTS (SUBSCRIBE): Collaterals from every ascending sensory pathway.

OUTPUTS (PUBLISH): Diffusely to the thalamus (intralaminar nuclei) and cortex, forming the ascending reticular activating system; and to the spinal cord, by the reticulospinal tracts.

FUNCTION AND FACULTY: Consciousness itself, in the sense of the level of arousal.

The demonstration by Moruzzi and Magoun that stimulation of the reticular formation awakens a sleeping animal, and that its destruction produces irreversible coma, is one of the foundational experiments of modern neuroscience.

It also contains the vomiting, cardiovascular, and micturition centres, and it mediates postural tone.

2.4.3.6 THE ROSTRAL VENTROLATERAL MEDULLA (RVLM)

Plain language: the master switch of blood pressure and sympathetic tone.

INPUTS (SUBSCRIBE): the nucleus of the solitary tract (baroreflex); the hypothalamus; the periaqueductal gray.

OUTPUTS (PUBLISH): the sympathetic preganglionic neurons of the spinal cord.

FUNCTION AND FACULTY: setting resting blood pressure and driving the sympathetic fight-or-flight response.

SYSTEM MEMBERSHIP: the autonomic system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: COMPARATOR-and-DRIVE (sympathetic). State: pressor set-point. Rule: sets sympathetic outflow.

MODULATION: norepinephrine, glutamate.

2.4.3.7 THE NUCLEUS AMBIGUUS

Plain language: the motor nucleus of swallow and voice, and the brake on the heart.

INPUTS (SUBSCRIBE): the nucleus of the solitary tract; the pre-Botzinger complex; motor and autonomic centres.

OUTPUTS (PUBLISH): the pharyngeal and laryngeal muscles (by the vagus and glossopharyngeal nerves); the heart (parasympathetic).

FUNCTION AND FACULTY: swallowing and phonation, and the parasympathetic slowing of the heart.

SYSTEM MEMBERSHIP: the motor and autonomic systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: RELAY (branchiomotor/cardiac). State: motor command. Rule: drives swallow, voice, and cardiac vagal tone.

MODULATION: acetylcholine.

2.4.3.8 THE PRE-BOTZINGER COMPLEX

Plain language: the pacemaker of breathing, the small knot of neurons that makes each breath begin.

INPUTS (SUBSCRIBE): chemosensory signals; the pontine respiratory group; arousal systems.

OUTPUTS (PUBLISH): the phrenic and airway motor neurons; the cardiovascular and arousal centres.

FUNCTION AND FACULTY: generating the inspiratory rhythm and coupling breathing to arousal and emotional state.

SYSTEM MEMBERSHIP: the respiratory and arousal systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: OSCILLATOR (respiratory rhythm). State: breathing phase. Rule: generates the inspiratory rhythm.

MODULATION: acetylcholine, substance P, opioids.

2.4.4 THE ASCENDING AROUSAL SYSTEM NUCLEI (PUC)

DEFINITION: The set of brainstem, hypothalamic, and basal forebrain nuclei that collectively constitute the ascending arousal system, considered as a single named Layer 9 anatomical ensemble.

THE MEMBERS

- Locus coeruleus (norepinephrine).

  https://www.linkedin.com/in/d-r-dison

- Raphe nuclei (serotonin).
- Tuberomammillary nucleus (histamine).
- Ventral tegmental area (dopamine).
- Pedunculopontine and laterodorsal tegmental nuclei (acetylcholine).
- Basal forebrain cholinergic complex (acetylcholine).
- Lateral hypothalamic orexin neurons (orexin/hypocretin).
- Parabrachial nucleus (glutamate).

PROOF OF EXISTENCE: These nuclei are the complete set of sources of the ascending projections that maintain cortical arousal.

They are jointly inhibited by the ventrolateral preoptic nucleus during sleep, and they jointly inhibit it in return, forming the sleep-wake flip-flop switch.

Damage to them collectively produces coma.

PROOF OF NAMELESSNESS: Each nucleus is individually named.

The "ascending reticular activating system" is a named functional concept dating to 1949, but it is a functional and physiological term, not an anatomical one, and it long predates the identification of the specific nuclei involved.

The precise set of nuclei constituting the system has never been named as a unified Layer 9 anatomical ensemble.

FUNCTION AND FACULTY: The maintenance of cortical arousal and the level of consciousness; the transitions between waking, non-REM sleep, and REM sleep.

Orexin is the stabilizer of the whole switch, and its loss produces narcolepsy: a switch that flips uncontrollably.

SYSTEM MEMBERSHIP: The sleep-wake and arousal systems (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: Assign per Chapter 8 archetype (RELAY / INTEGRATOR / OSCILLATOR / ATTRACTOR / GATE / COMPARATOR). State: activation; gain. Rule: transforms input per its function above.

DRIVES: Not Applicable, for the reason that this construct does not monitor a bodily variable and so defends no set-point; it would serve a drive only if it monitored such a variable, which it does not.

PLASTICITY: Weights change by the three-factor rule of Chapter 10 where the region learns; otherwise Not Applicable, for the reason that a region without modifiable connections has no weights to update.

MODULATION: Set by the neuromodulators named in its inputs.

2.5 THE CEREBELLUM

  https://www.linkedin.com/in/d-r-dison

Plain language: the little brain at the back.

It contains more neurons than all the rest of the brain combined, and it does not initiate anything.

It corrects.

It is the organ of skill, of timing, and of the difference between the movement you intended and the movement you made.

STRUCTURE

The cerebellar cortex has three laminae and an astonishingly uniform circuit repeated everywhere: granule cells, whose parallel fibres run in sheets; Purkinje cells, the sole output of the cerebellar cortex, which are GABAergic and therefore inhibitory; and the interneurons.

The uniformity is the point: the cerebellum applies one algorithm, and what differs from region to region is only what is plugged into it.

INPUTS (SUBSCRIBE)

• Mossy fibres, arriving by the middle and inferior cerebellar peduncles, from: the pontine nuclei (carrying the entire cerebral cortex, by way of the corticopontine projection); the spinal cord (by the spinocerebellar tracts, carrying proprioception); the vestibular nuclei; and the reticular formation.

Cargo: context, intention, and the current state of the body.

• Climbing fibres, arriving by the inferior cerebellar peduncle, exclusively from the inferior olivary nucleus.

Cargo: the error signal.

One climbing fibre per Purkinje cell, and its input is overwhelming.

INTERNAL CIRCUIT

Mossy fibre to granule cell to parallel fibre to Purkinje cell.

The parallel fibre-Purkinje cell synapse undergoes long-term depression when it is active at the same moment as a climbing fibre input.

This is the learning rule: the climbing fibre says "that was wrong," and the parallel fibre synapses that were active at that moment are weakened.

Over many trials, the movement is corrected.

This is supervised learning, implemented in wetware, and it was worked out in detail by Marr and Albus before it was confirmed experimentally.

OUTPUTS (PUBLISH)

The Purkinje cells project, inhibitorily, to the deep cerebellar nuclei, which are the sole output of the cerebellum.

  https://www.linkedin.com/in/d-r-dison

• Dentate nucleus, to the thalamus (VL), by way of the superior cerebellar peduncle and the dentato-thalamo-cortical pathway, and thence to the motor and prefrontal cortex.

Cargo: motor planning and, in the lateral cerebellum, cognitive timing.

• Interposed nuclei (globose and emboliform), to the red nucleus and the thalamus.

Cargo: limb movement correction.

• Fastigial nucleus, to the vestibular nuclei and the reticular formation.

Cargo: balance and posture.

FUNCTION AND FACULTY

Motor coordination, timing, and error correction; the acquisition of skilled movement; the calibration of reflexes; and, as is now firmly established but still under-taught, a substantial contribution to cognition, language, and emotion by way of the cerebro-cerebellar loops with the prefrontal cortex.

Damage produces ataxia: not weakness, but incoordination.

The patient can move; they cannot move accurately.

The cardinal signs are dysmetria (overshooting a target), intention tremor (a tremor that worsens as the hand approaches its goal), and dysdiadochokinesia (the inability to perform rapid alternating movements).

Damage to the posterior cerebellum can also produce the cerebellar cognitive affective syndrome, with executive and emotional disturbance and no motor sign at all.

SYSTEM MEMBERSHIP: The motor system, and the cerebro-cerebellar cognitive loops (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: Assign per Chapter 8 archetype (RELAY / INTEGRATOR / OSCILLATOR / ATTRACTOR / GATE / COMPARATOR). State: activation; gain. Rule: transforms input per its function above.

DRIVES: Not Applicable, for the reason that this construct does not monitor a bodily variable and so defends no set-point; it would serve a drive only if it monitored such a variable, which it does not.

PLASTICITY: Weights change by the three-factor rule of Chapter 10 where the region learns; otherwise Not Applicable, for the reason that a region without modifiable connections has no weights to update.

MODULATION: Set by the neuromodulators named in its inputs.

2.6 THE SPINAL CORD

Plain language: not merely a cable.

It is a computer that knows how to walk, and it will keep on walking without any instruction from the brain at all.

STRUCTURE: Grey matter arranged in a butterfly, with the dorsal horn receiving sensation, the ventral horn issuing motor commands, and the intermediolateral cell column (in thoracic and upper lumbar segments) issuing sympathetic outflow.

White matter surrounds it in ascending and descending tracts.

INPUTS (SUBSCRIBE)

• From the dorsal root ganglia, by way of the dorsal roots.

Cargo: all somatic sensation from the body.

• From the motor cortex, by way of the corticospinal tract.

• From the brainstem, by way of the rubrospinal, reticulospinal, vestibulospinal, and tectospinal tracts.

• From the raphe nuclei and locus coeruleus, by descending monoaminergic projections that gate pain at the dorsal horn.

OUTPUTS (PUBLISH)

• To the muscles, by way of the ventral roots and the alpha motor neurons, which constitute the final common path of Sherrington: every movement, whatever its origin in the brain, must pass through these cells.

• To the brain, by way of the dorsal column-medial lemniscal pathway (fine touch, vibration, proprioception, decussating in the medulla), the spinothalamic tract (pain, temperature, decussating at the segment of entry), and the spinocerebellar tracts (unconscious proprioception).

• To the sympathetic chain, by way of the intermediolateral cell column.

FUNCTION AND FACULTY: Reflexes, including the monosynaptic stretch reflex, which is the simplest circuit in the nervous system and the one tested by the tendon hammer; central pattern generation for locomotion; the gate control of pain at the dorsal horn; and the relay of all sensation and motor command between body and brain.

2.7 THE PERIPHERAL GANGLIA

A ganglion is a cluster of neuronal cell bodies lying outside the central nervous system.

They are Layer 9 constructs by every criterion, and they are routinely omitted from region atlases, which is an error.

SENSORY GANGLIA

• Dorsal root ganglia (DRG).

Pseudounipolar sensory neurons.

Inputs (SUBSCRIBE): from the skin, muscle, joints, and viscera.

Outputs (PUBLISH): to the spinal dorsal horn and the dorsal columns.

Function: the cell bodies of all somatic sensory afferents.

They perform no synaptic integration; each cell is a single unbranched relay from the periphery to the cord.

• Trigeminal ganglion.

The DRG equivalent for the face.

• Spiral (cochlear) ganglion and vestibular (Scarpa's) ganglion.

Hearing and balance.

• Nodose (inferior vagal) ganglion.

The cell bodies of the vagal afferents.

This is the ganglion whose cells are contacted by the gut's neuropod cells, and it is therefore the peripheral origin of the gut-brain Layer-Skipping Channel.

AUTONOMIC GANGLIA

• Sympathetic chain (paravertebral) ganglia.

Inputs (SUBSCRIBE): preganglionic fibres from the intermediolateral cell column.

Outputs (PUBLISH): postganglionic noradrenergic fibres to the target organs.

Function: sympathetic distribution.

The superior cervical ganglion is a notable member, supplying the eye and, importantly, the pineal gland.

• Prevertebral ganglia: the celiac, superior mesenteric, and inferior mesenteric ganglia.

• Parasympathetic ganglia: the ciliary, pterygopalatine, submandibular, and otic ganglia in the head, and the intramural ganglia of the viscera.

2.7.1 THE PREVERTEBRAL GANGLION INTEGRATIVE CLUSTER (PUC)

DEFINITION: The celiac, superior mesenteric, and inferior mesenteric ganglia considered as a single functional integrative unit.

INPUTS (SUBSCRIBE): From spinal preganglionic sympathetic neurons.

And, critically, from the enteric nervous system itself, by way of afferent fibres running from the gut back to the ganglia.

OUTPUTS (PUBLISH): Postganglionic sympathetic fibres to the abdominal and pelvic viscera.

  https://www.linkedin.com/in/d-r-dison

PROOF OF EXISTENCE: These ganglia receive convergent input from both the central nervous system and the enteric nervous system, and they contain interneurons.

They therefore perform genuine integration, not mere relay: they compute a decision about visceral sympathetic output from two independent information sources.

PROOF OF NAMELESSNESS: Each ganglion is individually named.

The concept of a peripheral integrative centre exists in the autonomic literature.

But the three prevertebral ganglia as a unified Layer 9 peripheral processing unit have never been named.

FUNCTION AND FACULTY: Peripheral integration of central and enteric signals; regulation of gut motility, secretion, and splanchnic blood flow.

This construct is important to the framework because it demonstrates that genuine neural computation occurs outside the skull.

2.8 THE ENTERIC NERVOUS SYSTEM

• Myenteric (Auerbach's) plexus.

Between the muscle layers of the gut.

Function: control of gut motility.

• Submucosal (Meissner's) plexus.

Function: control of secretion and local blood flow.

Together these contain roughly five hundred million neurons, more than the spinal cord, and they can generate coordinated peristalsis with the vagus completely severed.

This is why the enteric nervous system is legitimately called the second brain.

2.9 THE ENDOCRINE GLANDS AS LAYER 9 STRUCTURES

The companion volume argues that the endocrine glands are Layer 9 anatomical structures on exactly the same footing as the brain nuclei, and that a framework that admits the amygdala but excludes the adrenal gland is drawing an arbitrary line.

They are therefore listed here.

• The pituitary gland.

Anterior lobe: true glandular tissue, receiving hypothalamic releasing hormones by way of the hypophyseal portal system, and secreting ACTH, TSH, GH, LH, FSH, and prolactin.

Posterior lobe: not a gland at all, but a bundle of hypothalamic axon terminals releasing oxytocin and vasopressin directly into the blood.

  https://www.linkedin.com/in/d-r-dison

• The pineal gland.

Melatonin.

Treated above at 2.3.4.

• The thyroid gland.

Thyroid hormone, which sets the metabolic rate of every cell and is indispensable to brain development.

• The parathyroid glands.

Calcium regulation.

• The adrenal cortex.

Cortisol and aldosterone.

• The adrenal medulla.

Adrenaline.

Note that its chromaffin cells are modified sympathetic neurons derived from the neural crest, receiving direct preganglionic synaptic input.

They are neurons that have traded their axons for access to the bloodstream, and they are the strongest single argument for the unified neuroendocrine treatment.

• The pancreatic islets.

Insulin and glucagon.

• The gonads.

Testosterone, oestrogen, progesterone, which exert both organizational effects on the developing brain and activational effects on the adult one.

• The diffuse enteroendocrine system.

Collectively the largest endocrine organ in the body, and the source of the gut hormones that report on nutritional state to the hypothalamus.

2.10 THE LARGE-SCALE COGNITIVE NETWORKS (LAYER 10)

The regions catalogued above are Layer 9 constructs. When several regions couple into a coherent, repeatedly-observed functional whole, that whole is a Layer 10 construct: a large-scale network. Modern cognitive neuroscience describes much of cognition at this level, so a complete catalogue must name the networks as carefully as it names the regions. Each network is given here with its principal nodes and, in the publish-and-subscribe vocabulary of this volume, what it subscribes to and what it publishes.

2.10.1 THE DEFAULT MODE NETWORK (DMN)

Plain language: the brain's inward-facing network, most active when attention turns away from the outside world toward memory, self, and imagination.

NODES: medial prefrontal cortex (ventromedial and dorsomedial), posterior cingulate cortex and precuneus, the angular gyrus of the inferior parietal lobule bilaterally, the lateral temporal cortex (middle temporal gyrus), and the hippocampal formation.

INPUTS (SUBSCRIBE): the medial temporal memory system; self-referential and valuation signals from medial prefrontal cortex; and a release signal, since the network engages when the task-positive networks disengage.

OUTPUTS (PUBLISH): a broadcast of internally-generated content to prefrontal and parietal cortex; drive to the hippocampus for retrieval; and an anticorrelation signal that trades off against the dorsal attention and executive networks.

FUNCTION AND FACULTY: autobiographical and prospective memory, self-referential thought, mentalizing about others, moral reasoning, and mind-wandering.

SYSTEM MEMBERSHIP: the internally-directed cognition system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: ATTRACTOR/BROADCAST (an internally-sustained state). State: activation of self-and-memory content. Rule: sustains and broadcasts internally-generated content; anticorrelates with task-positive networks.

DRIVES: Not Applicable, for the reason that this network serves cognition on which drives act rather than monitoring a bodily variable of its own.

PLASTICITY: Its coupling strengthens with experience and is disrupted in ageing and disease.

MODULATION: Dopamine, acetylcholine, and the arousal systems set its engagement.

2.10.2 THE SALIENCE NETWORK / VENTRAL ATTENTION NETWORK (SN / VAN)

Plain language: the network that detects what suddenly matters and switches the brain between its inward and outward modes.

NODES: the anterior insula (bilateral, right-dominant), the dorsal anterior cingulate cortex, and, for the ventral attention component, the temporoparietal junction and ventral inferior frontal cortex.

INPUTS (SUBSCRIBE): interoceptive signals by way of the insula; salient and unexpected sensory events; and reward and threat signals from the amygdala and brainstem.

OUTPUTS (PUBLISH): a switching signal that interrupts the default mode network and recruits the executive network; attentional reorienting commands; and autonomic set-points to brainstem centres.

FUNCTION AND FACULTY: salience detection, bottom-up reorienting of attention, interoceptive awareness, and network switching.

SYSTEM MEMBERSHIP: the salience and reorienting system (Layer 10).

SPECIFICATION BLOCK

  https://www.linkedin.com/in/d-r-dison

DYNAMICS: COMPARATOR-and-SWITCH. State: current salience estimate. Rule: detects the most salient event and toggles control between the default mode and executive networks.

DRIVES: Serves the safety and homeostatic drives by flagging their violations.

PLASTICITY: Learns what is salient by experience.

MODULATION: Norepinephrine (surprise), dopamine, acetylcholine.

2.10.3 THE DORSAL ATTENTION NETWORK (DAN)

Plain language: the network that aims attention deliberately, on purpose, at a chosen place or feature.

NODES: the frontal eye fields bilaterally, the intraparietal sulcus and superior parietal lobule, and the motion area MT-plus.

INPUTS (SUBSCRIBE): goal and priority signals from the executive network; the current visual scene from visual cortex.

OUTPUTS (PUBLISH): top-down gain applied to visual cortex at the attended location or feature; gaze commands to the superior colliculus and oculomotor system.

FUNCTION AND FACULTY: voluntary, goal-directed orienting of visuospatial attention and the selection of targets for the eyes and hands.

SYSTEM MEMBERSHIP: the top-down attention system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE (a spatial priority map applying gain). State: the attended location or feature. Rule: publishes attentional gain to sensory cortex and gaze commands per the current goal.

DRIVES: Not Applicable, for the reason that it serves perception and action on which drives act.

PLASTICITY: Attentional templates are learned.

MODULATION: Acetylcholine (attentional gain), norepinephrine, dopamine.

2.10.4 THE FRONTOPARIETAL CONTROL NETWORK / CENTRAL EXECUTIVE NETWORK (FPN / CEN)

Plain language: the brain's flexible controller, the network that holds goals, plans, and manipulates information in mind.

NODES: the dorsolateral prefrontal cortex, the frontopolar and rostrolateral prefrontal cortex, the anterior inferior parietal lobule around the intraparietal sulcus, the dorsal precuneus, the midcingulate cortex, the head of the caudate, and the mediodorsal thalamus.

INPUTS (SUBSCRIBE): task-engagement signals from the salience network; sensory evidence from association cortex; and value signals from the limbic and reward systems.

  https://www.linkedin.com/in/d-r-dison

OUTPUTS (PUBLISH): top-down control biases to the dorsal attention network and to sensory and motor systems; and a suppression signal to the default mode network during external tasks.

FUNCTION AND FACULTY: working memory, cognitive control, goal maintenance, adaptive task-switching, and decision-making, the core of fluid intelligence.

SYSTEM MEMBERSHIP: the executive control system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR-and-CONTROLLER (persistent working-memory states). State: the active goal and the contents of working memory. Rule: maintains goals and publishes control biases that reconfigure the other networks per task.

DRIVES: Not Applicable, for the reason that it serves cognition on which drives act.

PLASTICITY: Training strengthens its control; it is the substrate of skill in reasoning.

MODULATION: Dopamine (the inverted-U of prefrontal gain), norepinephrine, acetylcholine.

2.10.5 THE CINGULO-OPERCULAR NETWORK (CON)

Plain language: the network that holds a task-set steady across time, keeping the brain on task between moment-to-moment decisions.

NODES: the dorsal anterior cingulate cortex and adjacent medial superior frontal cortex, the anterior insula and frontal operculum, the anterior prefrontal cortex, and the thalamus.

INPUTS (SUBSCRIBE): performance feedback and error signals; arousal signals from the ascending systems; task instructions from the executive network.

OUTPUTS (PUBLISH): a tonic maintenance signal that sustains a stable task-set; error and feedback signals that adjust control; and initiation signals for action.

FUNCTION AND FACULTY: tonic alertness, stable task-set maintenance, error and feedback monitoring, and sustained control across trials.

SYSTEM MEMBERSHIP: the sustained-control system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR (tonic set-maintenance). State: the maintained task-set and alertness level. Rule: sustains a task-set across time and publishes error and feedback corrections.

DRIVES: Not Applicable.

PLASTICITY: Learns the value of maintaining versus switching sets.

MODULATION: Norepinephrine (tonic arousal), acetylcholine, dopamine.

  https://www.linkedin.com/in/d-r-dison

2.10.6 THE LANGUAGE NETWORK

Plain language: the network that turns sound and text into meaning and meaning back into speech.

NODES: Broca's area (the inferior frontal gyrus, areas 44 and 45), Wernicke's area (the posterior superior temporal gyrus and sulcus), the middle temporal gyrus, and the angular and supramarginal gyri, connected by the arcuate fasciculus and the extreme capsule, with the frontal aslant tract linking to the pre-supplementary motor area.

INPUTS (SUBSCRIBE): speech from the auditory cortex, and written words from the visual word form area.

OUTPUTS (PUBLISH): articulatory programs to the motor and premotor cortex; and semantic content to the memory and executive systems.

FUNCTION AND FACULTY: comprehension and production of language, syntax, semantics, and phonology.

SYSTEM MEMBERSHIP: the language system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR-and-SEQUENCER. State: the current linguistic representation. Rule: maps sound and text to meaning and meaning to articulation.

DRIVES: Not Applicable.

PLASTICITY: Language is learned; strongest plasticity in the critical period.

MODULATION: Dopamine, acetylcholine.

2.10.7 THE SENSORIMOTOR NETWORK (SMN)

Plain language: the perception-and-action network, the loop that senses the body and moves it.

NODES: the primary motor cortex, the primary somatosensory cortex, the supplementary motor area, and the premotor cortex, often divided into hand, foot, and mouth sub-networks.

INPUTS (SUBSCRIBE): somatosensory feedback by way of the ventral posterior thalamus; motor plans from the cerebellum and basal ganglia by way of the ventral thalamus.

OUTPUTS (PUBLISH): motor commands to the corticospinal tract and brainstem; efference copies to the cerebellum and parietal cortex.

FUNCTION AND FACULTY: motor execution and somatosensory feedback, the closed loop of embodied action.

SYSTEM MEMBERSHIP: the perception-action system (Layer 10).

SPECIFICATION BLOCK

  https://www.linkedin.com/in/d-r-dison

DYNAMICS: RELAY-and-CONTROLLER (closed sensorimotor loop). State: the current motor command and sensory feedback. Rule: issues motor commands and corrects them by feedback.

DRIVES: Serves every drive that requires action upon the world.

PLASTICITY: Motor learning reshapes it continuously.

MODULATION: Dopamine (vigor), acetylcholine, norepinephrine.

2.10.8 THE VISUAL NETWORK

Plain language: the seeing network, the visual hierarchy from first cortical stage to object and scene recognition.

NODES: the primary visual cortex, the extrastriate areas (secondary visual cortex through the fourth visual area), the motion complex MT-plus, and the ventral (object and scene) and dorsal (spatial) stream extensions.

INPUTS (SUBSCRIBE): the retinal image by way of the lateral geniculate nucleus, and attentional selection by way of the pulvinar.

OUTPUTS (PUBLISH): a spatial stream to the parietal cortex and dorsal attention network, and an object stream to the temporal recognition areas.

FUNCTION AND FACULTY: visual perception across the whole hierarchy, from edges to objects, scenes, and motion.

SYSTEM MEMBERSHIP: the visual system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: hierarchical RELAY-and-TRANSFORM. State: activation across the visual hierarchy. Rule: transforms the retinal image into progressively more abstract representations along two streams.

DRIVES: Not Applicable.

PLASTICITY: Perceptual learning tunes it; strong developmental critical period.

MODULATION: Acetylcholine, norepinephrine, attention from the dorsal attention network.

2.10.9 THE AUDITORY NETWORK

Plain language: the hearing network, from tonotopic sound analysis to speech and music.

NODES: the primary auditory cortex (Heschl's gyrus), the auditory belt and parabelt, and the superior temporal gyrus and sulcus.

INPUTS (SUBSCRIBE): sound by way of the medial geniculate nucleus.

OUTPUTS (PUBLISH): spectrotemporal patterns to the language network and to the multisensory superior temporal sulcus, and to prefrontal cortex.

FUNCTION AND FACULTY: auditory perception, sound-pattern and speech-sound analysis, and the precursors of language and music.

  https://www.linkedin.com/in/d-r-dison

SYSTEM MEMBERSHIP: the auditory system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: hierarchical RELAY-and-TRANSFORM (tonotopic). State: spectrotemporal activation. Rule: analyses sound into frequency and pattern and passes it to language and multisensory areas.

DRIVES: Not Applicable.

PLASTICITY: Auditory and speech learning tune it.

MODULATION: Acetylcholine, norepinephrine, attention.

2.10.10 THE LIMBIC NETWORK

Plain language: the network of value and feeling, where emotion, reward, and memory are joined to cognition.

NODES: the orbitofrontal cortex, the ventromedial prefrontal cortex, and the temporal pole, coupled with the subcortical amygdala, hippocampus, and ventral striatum.

INPUTS (SUBSCRIBE): sensory and mnemonic content from association cortex; emotional value from the amygdala; and interoceptive state from the insula and brainstem.

OUTPUTS (PUBLISH): value and affect signals to the executive and default networks; and drive to the hypothalamus, brainstem autonomic centres, and striatum.

FUNCTION AND FACULTY: valuation, emotion, reward, and the integration of affect with memory and decision.

SYSTEM MEMBERSHIP: the limbic and valuation system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: INTEGRATOR-and-APPRAISER. State: the current value and affective appraisal. Rule: integrates sensory, mnemonic, and interoceptive input into a value signal published to cognition and to the body.

DRIVES: Serves and reports on every homeostatic and social drive.

PLASTICITY: Reward and emotional learning reshape it (dopamine-gated).

MODULATION: Dopamine, serotonin, norepinephrine, the neuropeptides, and cortisol.

2.11 THE COGNITIVE SYSTEMS AND PATHWAYS (LAYER 10)

Some named constructs are neither single regions nor whole networks but recurring PATHWAYS and SYSTEMS that thread through several regions. A complete catalogue must name them, because a model will have to instantiate them as loops and rules, not as boxes.

2.11.1 THE CORTICO-BASAL-GANGLIA-THALAMO-CORTICAL LOOPS (CBGTC)

  https://www.linkedin.com/in/d-r-dison

Plain language: the family of parallel loops by which the cortex asks the basal ganglia to select an action, a thought, or a goal, and receives the selection back through the thalamus.

INPUTS (SUBSCRIBE): a cortical territory publishes to the striatum; dopamine from the midbrain sets the learning signal.

OUTPUTS (PUBLISH): through the pallidum or substantia nigra reticulata to a specific thalamic nucleus and back to the originating cortex.

FUNCTION AND FACULTY: selection and gating, running in segregated parallel loops, a motor loop, an oculomotor loop, an executive or cognitive loop (dorsolateral prefrontal to dorsomedial striatum), and a limbic or motivational loop (orbitofrontal and cingulate to ventral striatum).

SYSTEM MEMBERSHIP: the action- and thought-selection system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE (parallel selection loops). State: the selected channel per loop. Rule: releases one channel from tonic inhibition to permit an action or thought.

MODULATION: dopamine (the teaching and vigor signal).

2.11.2 THE DIRECT, INDIRECT, AND HYPERDIRECT PATHWAYS

Plain language: the three routes through the basal ganglia that say go, no-go, and stop.

INPUTS (SUBSCRIBE): cortical glutamate to the striatum (direct via dopamine-one cells, indirect via dopamine-two cells) and to the subthalamic nucleus (hyperdirect).

OUTPUTS (PUBLISH): the direct pathway disinhibits the thalamus (go); the indirect pathway, through the external pallidum and subthalamic nucleus, further inhibits the thalamus (no-go); the hyperdirect pathway applies a fast global brake (stop).

FUNCTION AND FACULTY: the facilitation, suppression, and rapid cancellation of actions, the microcircuit basis of selection and inhibition.

SYSTEM MEMBERSHIP: the action-selection system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: GATE (go/no-go/stop). State: the balance of the three pathways. Rule: their sum decides whether an action is released, withheld, or cancelled.

MODULATION: dopamine (opposite signs on the direct and indirect pathways).

2.11.3 THE PAPEZ CIRCUIT

Plain language: the closed limbic loop that binds memory to emotion, the ring whose breakage anywhere causes amnesia.

INPUTS (SUBSCRIBE): experience entering the hippocampus from the entorhinal cortex.

  https://www.linkedin.com/in/d-r-dison

OUTPUTS (PUBLISH): the loop hippocampus to fornix to mammillary bodies to mammillothalamic tract to anterior thalamus to cingulate to cingulum to parahippocampal and entorhinal cortex and back to the hippocampus.

FUNCTION AND FACULTY: episodic and declarative memory and its emotional colouring; damage at any node produces diencephalic or limbic amnesia (as in Korsakoff syndrome).

SYSTEM MEMBERSHIP: the memory and emotion system (Layer 10).

SPECIFICATION BLOCK

DYNAMICS: closed LOOP (reentrant). State: the circulating memory trace. Rule: recirculates memory content through the limbic ring for consolidation.

MODULATION: acetylcholine, dopamine, the theta rhythm.

2.11.4 THE TRIPARTITE SYNAPSE

Plain language: the synapse as three partners, not two, the presynaptic terminal, the postsynaptic cell, and the astrocyte that listens and answers.

INPUTS (SUBSCRIBE): the perisynaptic astrocyte subscribes to neurotransmitter spillover through its own receptors.

OUTPUTS (PUBLISH): the astrocyte publishes gliotransmitters, glutamate, D-serine, and adenosine triphosphate or adenosine, back onto the synapse.

FUNCTION AND FACULTY: glial gating of synaptic transmission and plasticity, the reason a model of neurons alone omits a partner that gates learning (see D-serine in Chapter 7).

SYSTEM MEMBERSHIP: the synaptic-plasticity system (Layer 3 and Layer 4, named here for completeness).

SPECIFICATION BLOCK

DYNAMICS: MODULATOR (glial gating). State: astrocytic calcium and gliotransmitter level. Rule: the astrocyte reads and adjusts the synapse.

MODULATION: neuronal activity; the neuromodulators.

2.11.5 THE GLIA AS COGNITIVE PARTICIPANTS (ASTROCYTES, MICROGLIA, OLIGODENDROCYTES)

Plain language: the non-neuronal cells that shape cognition, one supplying and gating chemistry, one pruning connections, one tuning the speed of thought.

INPUTS (SUBSCRIBE): astrocytes subscribe to neurotransmitter spillover; microglia to neural-activity and complement signals; oligodendrocytes to axonal activity.

OUTPUTS (PUBLISH): astrocytes publish gliotransmitters and metabolic support; microglia publish synaptic pruning, phagocytosing weak synapses; oligodendrocytes publish adaptive myelin that sets conduction speed.

  https://www.linkedin.com/in/d-r-dison

FUNCTION AND FACULTY: astrocytes gate plasticity and homeostasis; microglia sculpt circuits by activity-dependent pruning (shaping learning and forgetting); oligodendrocytes tune conduction speed by experience-dependent myelination.

SYSTEM MEMBERSHIP: the plasticity and connectivity systems (Layer 4, named here for completeness).

SPECIFICATION BLOCK

DYNAMICS: MODULATOR/STRUCTURAL. State: glial-mediated synapse and myelin state. Rule: glia adjust which synapses persist and how fast axons conduct.

MODULATION: neural activity, complement and fractalkine signals, neurotrophins.

CHAPTER 3. THE INTERFACES

3.1 The Central Claim of This Chapter

In a layered architecture, the interfaces are not the plumbing between the interesting parts.

The interfaces are the interesting part.

This claim can be defended precisely, and the defence is the argument of this chapter.

Consider what happens when you attempt to define a brain region without reference to its connections.

You cannot.

The hippocampus is not defined by its cell types, because CA1 pyramidal cells are unremarkable pyramidal cells found in similar form elsewhere.

It is not defined by its microcircuit motifs, because autoassociation and pattern separation are generic Layer 8 motifs.

It is defined by the fact that it receives the perforant path from the entorhinal cortex and sends the fornix to the mammillary bodies.

Change those two connections, and you no longer have a hippocampus.

You have a piece of tissue that happens to be shaped like one.

This is the formal statement of the principle: at Layer 9, a region's identity is its connectional fingerprint, which is to say that a region is the sum of its interfaces.

It follows that a complete Layer 9 taxonomy must name the interfaces as carefully as it names the regions, and that a science which has named the regions but left many of the interfaces nameless is a science working with half a vocabulary.

3.2 A Formal Specification of a Neural Interface

  https://www.linkedin.com/in/d-r-dison

Every named interface in this chapter is specified by seven properties, borrowed from the engineering discipline of interface design.

• The origin.

Which region sends.

From where.

• The destination.

Which region receives.

To where.

• The cargo.

What information is carried, stated in terms of what it represents rather than what molecules carry it.

• The format.

The encoding: excitatory or inhibitory, driving or modulating, topographic or diffuse, point-to-point or broadcast.

• The direction.

Feedforward, feedback, or reciprocal.

• The faculty.

What the interface itself does to the signal, if anything.

This property is discussed at length in Section 3.3, and it is the one that most previous treatments omit entirely.

• The consequence of failure.

What is lost when the interface is cut.

A NOTE ON HOW ORIGIN AND DESTINATION ARE LABELLED: SEND, RECEIVE, PUBLISH, SUBSCRIBE

The origin and destination of each interface catalogued below carry a parenthetical borrowed from message-oriented software, so that a coding-assistant artificial-intelligence system building the digital model can read an interface's messaging mode at a glance.

The brain uses two, and only two, modes, and the format field distinguishes them.

A point-to-point interface is ADDRESSED: a dedicated tract runs from one region to one destination, and no other region is on the line.

Its origin is marked (SEND) and its destination (RECEIVE), the directed-channel primitive, because the sender knows its receiver, being wired to it.

  https://www.linkedin.com/in/d-r-dison

The corticospinal tract, the optic radiation, and the great sensory and corticocortical cables are all of this kind.

A broadcast interface is NOT addressed: it releases its cargo diffusely and acts on whatever bears the receptor, exactly as Grade 3 above describes.

Its origin is marked (PUBLISH) and its destination (SUBSCRIBE), the broadcast primitive, because the publisher does not know its subscribers, which are defined by the receptor they carry.

The three dopamine projections, the nigrostriatal, mesolimbic, and mesocortical, and the hypothalamo-hypophyseal release of oxytocin and vasopressin into the systemic bloodstream, are of this kind, and are the only broadcast interfaces in the catalogue below.

This is the same publish-and-subscribe pattern used for the molecules in Chapter 7, now joined to the send-and-receive pattern of directed wiring, so that the whole system speaks one communication language with exactly two primitives: directed channels and broadcast topics.

3.2.1 The Faculty Field, and Why It Is Not Optional

It is tempting to specify an interface by origin, destination, and cargo alone, and to leave it there.

That vocabulary carries a hidden and false premise: that an interface transports a message without altering it, and that all computation happens in the regions at either end.

An interface, on that view, is a pipe.

Some interfaces are indeed pipes.

Many are not.

The distinction is not one of degree, and it is diagnosed by a single test, which we shall call the Perfect Wire Test.

THE PERFECT WIRE TEST

Ask what would be lost if the interface were replaced by a perfect wire: a hypothetical connection of infinite fidelity that delivered the origin's output to the destination instantly and unchanged.

If nothing would be lost, the interface is transmissive.

It is a pipe, and the perfect wire is a perfect substitute.

Its faculty is fidelity, and fidelity is the whole of its contribution.

If something would be lost, the interface is computational.

It was doing work upon the signal in transit, and that work is its faculty.

The corticospinal tract passes as a pipe: a perfect wire from motor cortex to spinal motor neuron would serve as well, and its job is speed and fidelity and nothing more.

The thalamic reticular nucleus does not pass.

  https://www.linkedin.com/in/d-r-dison

Replace it with a perfect wire and every thalamic signal would reach the cortex with flawless fidelity, which sounds like an improvement and would be a catastrophe, because you would lose the capacity to ignore anything.

Attention is not a property of the thalamus and not a property of the cortex.

It is a property of the gate between them.

Every interface catalogued below therefore carries a FACULTY line, and each is classified as transmissive or computational.

3.3 THE COMPUTATIONAL INTERFACE (PUC), AND THE THREE GRADES

3.3.1 The Computational Interface (PUC)

DEFINITION: A Computational Interface is an interface between two Layer 9 regions that performs a transformation upon the signal in transit, or that possesses a faculty resident in neither the origin nor the destination, such that its replacement by a perfect wire would destroy a capacity of the organism.

PROOF OF EXISTENCE: The thalamic reticular nucleus, the optic chiasm, the intercalated cell masses, the dentate gyrus, the inferior olive, and the superior olivary complex all demonstrably transform signals in transit, and each possesses a faculty found in neither the structure that sends to it nor the structure it sends to.

PROOF OF NAMELESSNESS: Neuroscience possesses the driver and modulator distinction, which concerns how strongly an input determines its target's output.

It possesses the relay concept, and it has spent two decades arguing that the thalamus is "more than a relay." But it has no term for the general category of interfaces that compute rather than convey, and consequently no way to state the generalization that unites them.

CONSEQUENCE: Region and interface are not two kinds of object at Layer 9.

They are two roles, and a single structure may occupy both.

The thalamic reticular nucleus is a region by every anatomical criterion and a pure interface by every functional one.

3.3.2 The Three Grades

Independently of whether an interface computes, neural interfaces fall into three grades by the strength with which they determine their target.

GRADE 1: DRIVING INTERFACES

A driving interface determines what the target region transmits.

It carries the message itself.

The retinogeniculate projection is a driver: what the LGN sends to the cortex is, in essence, what the retina told it.

  https://www.linkedin.com/in/d-r-dison

Drivers are typically few, powerful, and terminate on proximal dendrites where they cannot be ignored.

The climbing fibre is the extreme case: a single climbing fibre input to a Purkinje cell is so powerful that it constitutes a command rather than a suggestion.

GRADE 2: MODULATING INTERFACES

A modulating interface adjusts how a message is transmitted without determining its content.

The corticogeniculate feedback projection is a modulator: it does not tell the LGN what the retina saw, but it powerfully adjusts the gain, the timing, and the reliability with which the LGN passes that message on.

Modulators are typically numerous, individually weak, and terminate on distal dendrites.

The driver-modulator distinction is not a subtlety.

It is the reason a projection ten times larger than another can nonetheless be the less important of the two for determining content.

Numerical dominance and functional dominance are different things, and conflating them has produced a great deal of confused writing about the thalamus.

GRADE 3: BROADCAST INTERFACES

A broadcast interface is not addressed.

It releases its cargo diffusely into a large volume of tissue, and it acts on whatever bears the receptor.

The locus coeruleus projection is a broadcast interface.

So is the nucleus basalis cholinergic projection, and so is every ascending monoamine system.

Broadcast interfaces are the Layer 9 instantiation of the Broadcast Principle set out in the companion volume: they exist precisely where the message is global rather than addressed, and where the cost of building a private line to every recipient would be prohibitive.

A broadcast interface is what you use when the message is "this matters, everybody attend," and there is no one to address it to because the answer is everyone.

3.4 THE ESTABLISHED AND NAMED INTERFACES

What follows is a catalogue of the major named interfaces of the human brain, specified according to Section 3.2.

3.4.1 SENSORY INPUT INTERFACES

• The optic nerve, chiasm, and tract.

FACULTY: COMPUTATIONAL.

  https://www.linkedin.com/in/d-r-dison

The chiasm performs a decussation, and the decussation is a computation: nasal fibres cross and temporal fibres do not, so that each hemisphere receives the contralateral half of the visual world rather than the contralateral eye.

Replace it with a perfect wire, each eye reporting to its own hemisphere, and stereoscopic depth perception is lost.

The sorting is the function.

Origin (SEND): retinal ganglion cells.

Destination (RECEIVE): LGN, superior colliculus, suprachiasmatic nucleus, pretectum.

Cargo: the visual image, already substantially processed by the retina.

Format: driving, topographic.

At the chiasm, nasal fibres decussate, so that each hemisphere receives the contralateral half of the visual world, not the contralateral eye.

Failure: blindness, with the pattern of the field defect localizing the lesion precisely, which is why the visual pathway is the neurologist's favourite teaching tool.

• The optic radiation (geniculocalcarine tract).

FACULTY: TRANSMISSIVE.

It carries the retinotopic image and does not transform it.

A perfect wire would serve.

Origin (SEND): LGN.

Destination (RECEIVE): V1.

Cargo: the retinotopic image.

Format: driving, terminating in lamina 4C.

Failure: homonymous hemianopia.

Its inferior fibres loop forward through the temporal lobe as Meyer's loop, which is why temporal lobectomy for epilepsy characteristically produces a superior quadrantanopia, a "pie in the sky" field defect.

• The acoustic radiation.

FACULTY: TRANSMISSIVE.

Origin (SEND): MGN.

Destination (RECEIVE): A1.

Cargo: tonotopic auditory information.

  https://www.linkedin.com/in/d-r-dison

• The medial lemniscus.

FACULTY: TRANSMISSIVE.

Fidelity is its whole contribution.

Origin (SEND): the dorsal column nuclei (gracile and cuneate) of the medulla.

Destination (RECEIVE): VPL of the thalamus.

Cargo: fine touch, vibration, proprioception.

Format: driving, topographic, and already decussated.

Failure: contralateral loss of fine touch and position sense, with pain and temperature preserved, a dissociation that is diagnostic.

• The spinothalamic tract.

FACULTY: TRANSMISSIVE in transit, though its origin at the dorsal horn is itself a computational gate, where descending pathways modulate whether the signal is transmitted at all.

Origin (SEND): the dorsal horn of the spinal cord.

Destination (RECEIVE): VPL, and the intralaminar and VMpo nuclei.

Cargo: pain, temperature, crude touch.

It decussates at the level of entry, not in the medulla, which is why a hemisection of the cord produces the celebrated Brown-Sequard syndrome: loss of pain and temperature on the opposite side, and loss of touch and proprioception on the same side.

• The lateral olfactory tract.

FACULTY: TRANSMISSIVE, but architecturally exceptional: by omitting the thalamic relay it removes an entire gating stage, which is why olfaction reaches emotion and memory with so little mediation.

Origin (SEND): the olfactory bulb.

Destination (RECEIVE): piriform cortex, amygdala, entorhinal cortex.

Cargo: odour.

Format: driving, and uniquely without thalamic relay.

3.4.2 MOTOR OUTPUT INTERFACES

• The corticospinal (pyramidal) tract.

FACULTY: TRANSMISSIVE, with one exception: the pyramidal decussation is computational, since the crossing rule is what establishes contralateral motor control, and it is why a stroke on one side paralyses the other.

Origin (SEND): M1, premotor, SMA, and S1, laminae 5.

  https://www.linkedin.com/in/d-r-dison

Destination (RECEIVE): spinal motor neurons and interneurons.

Cargo: voluntary motor commands, especially for fine, fractionated finger movement.

Format: driving, topographic, decussating at the medullary pyramids.

Failure: contralateral spastic paresis, with the Babinski sign.

• The corticobulbar tract.

FACULTY: COMPUTATIONAL in its routing.

Its bilateral innervation of the forehead and contralateral innervation of the lower face is a routing rule implemented in the interface, and it is what allows a clinician to distinguish a stroke from a facial nerve palsy at the bedside.

Origin (SEND): motor cortex.

Destination (RECEIVE): cranial nerve motor nuclei.

Cargo: control of face, jaw, tongue, larynx.

Note the clinically vital asymmetry: the forehead receives bilateral input, which is why a stroke spares the forehead while a facial nerve palsy does not.

• The rubrospinal, reticulospinal, vestibulospinal, and tectospinal tracts.

The extrapyramidal descending system, governing posture, tone, balance, and reflexive orienting.

3.4.3 THE BASAL GANGLIA INTERFACES

• The corticostriatal projection.

Origin (SEND): the entire cortex, lamina 5.

Destination (RECEIVE): caudate, putamen, nucleus accumbens, topographically.

Cargo: candidate actions, thoughts, and motivations awaiting selection.

Format: driving, glutamatergic.

• The hyperdirect pathway (cortico-subthalamic projection).

FACULTY: COMPUTATIONAL.

By bypassing the striatum it converts a deliberative selection process into a fast global veto.

The shortcut is the function.

Origin (SEND): motor and premotor cortex.

Destination (RECEIVE): subthalamic nucleus.

  https://www.linkedin.com/in/d-r-dison

Cargo: a global stop signal.

Format: driving, glutamatergic, monosynaptic, and fast.

This is a within-system shortcut that bypasses the striatum entirely, and it exists because stopping cannot afford the latency of deliberation.

• The nigrostriatal pathway.

FACULTY: COMPUTATIONAL.

It does not carry a movement command.

It carries a scalar teaching signal, reward prediction error, which reweights the selection process itself.

Origin (PUBLISH): substantia nigra pars compacta.

Destination (SUBSCRIBE): caudate and putamen.

Cargo: dopamine, encoding reward prediction error, biasing selection toward "go." Format: broadcast within the striatum.

Failure: Parkinson's disease.

• The mesolimbic pathway.

Origin (PUBLISH): VTA.

Destination (SUBSCRIBE): nucleus accumbens.

Cargo: incentive salience.

Failure or hijack: addiction.

• The mesocortical pathway.

Origin (PUBLISH): VTA.

Destination (SUBSCRIBE): prefrontal cortex.

Cargo: gain control of working memory.

• The pallidothalamic tract (ansa lenticularis and lenticular fasciculus, converging as the thalamic fasciculus).

FACULTY: COMPUTATIONAL.

It implements selection by disinhibition: the default is suppression of all actions, and to act is to release one channel.

The inversion of sign is the computation.

Origin (SEND): GPi.

Destination (RECEIVE): VA and VL thalamus.

Cargo: tonic inhibition, selectively released to permit action.

  https://www.linkedin.com/in/d-r-dison

Format: inhibitory, GABAergic.

• The stria medullaris and the fasciculus retroflexus.

Origin (SEND): GPi and limbic forebrain, then the habenula.

Destination (RECEIVE): habenula, then the VTA and raphe.

Cargo: the negative reward signal.

3.4.4 THE MEMORY AND LIMBIC INTERFACES

• The perforant path.

FACULTY: TRANSMISSIVE into the dentate, but the dentate gyrus that receives it is a computational interface stage: it re-encodes a dense cortical input into a sparse, decorrelated representation.

This is pattern separation, and it is performed in transit.

Replace the dentate with a perfect wire and similar experiences arrive at CA3 as similar patterns, and CA3, being autoassociative, confuses them.

You would lose the ability to distinguish where you parked today from where you parked yesterday.

Origin (SEND): entorhinal cortex, layer II.

Destination (RECEIVE): dentate gyrus and CA3.

Cargo: the convergent output of the entire association cortex, meaning the content of experience.

Format: driving, glutamatergic.

Failure: the inability to encode new declarative memory.

This is arguably the single most important interface in the study of memory.

• The temporoammonic path.

Origin (SEND): entorhinal cortex, layer III.

Destination (RECEIVE): CA1 directly.

Cargo: current sensory reality, against which CA3's reconstruction is compared.

A genuine within-structure shortcut.

• The mossy fibre pathway.

FACULTY: COMPUTATIONAL.

Its extreme synaptic strength enforces sparse, high-fidelity transfer, and it is why a single fibre can detonate its target.

Origin (SEND): dentate granule cells.

  https://www.linkedin.com/in/d-r-dison

Destination (RECEIVE): CA3.

Cargo: sparse, pattern-separated representations.

Format: driving, and so powerful that a single fibre can detonate its target.

• The Schaffer collaterals.

FACULTY: COMPUTATIONAL.

It is a plastic interface: its weights change with use, which means the interface stores information.

This is the synapse of long-term potentiation, and an interface that learns is by definition not a pipe.

Origin (SEND): CA3.

Destination (RECEIVE): CA1.

Cargo: the pattern-completed reconstruction.

This is the synapse of long-term potentiation.

• The fornix.

FACULTY: TRANSMISSIVE.

Origin (SEND): subiculum and CA1.

Destination (RECEIVE): mammillary bodies, anterior thalamus, septal nuclei, nucleus accumbens.

Cargo: the hippocampal output.

Format: the great efferent cable of the hippocampus.

Failure: anterograde amnesia.

• The mammillothalamic tract.

Origin (SEND): mammillary bodies.

Destination (RECEIVE): anterior thalamic nuclei.

Cargo: the Papez circuit signal.

Failure: Korsakoff amnesia.

• The cingulum bundle.

Origin and destination: reciprocally connects the cingulate cortex, the medial prefrontal cortex, the parahippocampal cortex, and the hippocampus.

Cargo: the traffic of the default mode network and the Papez circuit.

• The stria terminalis.

  https://www.linkedin.com/in/d-r-dison

Origin (SEND): amygdala.

Destination (RECEIVE): BNST, hypothalamus, septal nuclei.

Cargo: emotional signals to subcortical targets.

It takes a long arcing course, following the caudate.

• The ventral amygdalofugal pathway.

Origin (SEND): central and basolateral amygdala.

Destination (RECEIVE): mediodorsal thalamus, hypothalamus, brainstem, nucleus accumbens, prefrontal and cingulate cortex.

Cargo: emotional signals to both cortical and subcortical targets.

The division of labour is precise: the stria terminalis reaches subcortical structures, and the amygdalofugal pathway reaches cortical ones.

• The uncinate fasciculus.

FACULTY: TRANSMISSIVE as a bundle, but it carries the fibres of a computational interface: the fronto-amygdalar regulatory channel described at 3.5.8, whose sign inversion occurs at the intercalated cell masses.

Origin and destination: reciprocally connects the orbitofrontal cortex and the anterior temporal lobe and amygdala.

Cargo: the top-down regulation of emotion, and the retrieval of emotionally valenced memory.

Its integrity correlates with the capacity to extinguish fear, and it is the tract implicated in psychopathy.

3.4.5 THE CEREBELLAR INTERFACES

• The corticopontine projection.

FACULTY: TRANSMISSIVE.

Origin (SEND): the entire cerebral cortex.

Destination (RECEIVE): pontine nuclei.

Cargo: the intended action, and its context.

Roughly twenty million fibres.

• The pontocerebellar projection (middle cerebellar peduncle).

Origin (SEND): pontine nuclei.

Destination (RECEIVE): contralateral cerebellar cortex, as mossy fibres.

• The olivocerebellar projection (inferior cerebellar peduncle).

FACULTY: COMPUTATIONAL.

  https://www.linkedin.com/in/d-r-dison

The inferior olive manufactures the error signal by comparing an efference copy against sensory feedback.

The error exists in neither the cortex, which knows only what it intended, nor the periphery, which knows only what occurred.

It is created in the comparison, and the comparison is the interface.

Origin (SEND): inferior olive.

Destination (RECEIVE): cerebellar cortex, as climbing fibres.

Cargo: the error signal.

Format: driving, one fibre per Purkinje cell, overwhelming.

• The dentato-thalamo-cortical pathway (superior cerebellar peduncle).

FACULTY: TRANSMISSIVE.

Origin (SEND): dentate nucleus.

Destination (RECEIVE): VL thalamus, then motor and prefrontal cortex.

Cargo: the correction.

Failure: ataxia.

3.4.6 THE CORTICOCORTICAL INTERFACES

• The corpus callosum.

FACULTY: MIXED.

Largely transmissive, but its interhemispheric inhibition is computational: each hemisphere suppresses its counterpart, and the rivalry is resolved in the interface.

Roughly two hundred million fibres connecting homologous cortical areas across the hemispheres.

Failure or surgical section: the split-brain syndrome, in which each hemisphere can be shown to possess information the other lacks and cannot report.

• The arcuate fasciculus.

FACULTY: TRANSMISSIVE.

Its severance produces conduction aphasia precisely because it is a pipe: comprehension and fluency survive intact at either end, and only the transfer between them fails.

Origin and destination: Wernicke's area and Broca's area.

Cargo: phonological representations for repetition.

  https://www.linkedin.com/in/d-r-dison

Failure: conduction aphasia, in which comprehension and fluency survive but repetition fails.

• The superior longitudinal fasciculus.

Connects frontal, parietal, and temporal cortex.

Cargo: spatial attention and working memory traffic.

• The inferior fronto-occipital fasciculus and the inferior longitudinal fasciculus.

Carry the ventral visual stream forward to frontal and temporal targets.

3.4.7 THE NEUROENDOCRINE INTERFACES

These are interfaces in the full sense, and they are omitted from every standard white matter atlas because they are not made of white matter.

The framework of this volume insists on their inclusion.

• The hypophyseal portal system.

FACULTY: COMPUTATIONAL, and of a kind found nowhere else: it performs format conversion.

It transduces a neural code, expressed as action potentials, into an endocrine code, expressed as hormone concentration in blood.

A perfect wire cannot substitute, because the destination does not speak the language of the origin.

Origin (SEND): hypothalamic parvocellular neurons, terminating in the median eminence.

Destination (RECEIVE): the anterior pituitary.

Cargo: the releasing hormones (CRH, TRH, GnRH, GHRH) and inhibiting hormones.

Format: a dedicated private capillary bed, not an axon.

It exists so that hypothalamic hormones can reach the pituitary in high concentration without dilution in the systemic circulation.

It is an anatomical acknowledgement that even the endocrine system sometimes needs a private line rather than a broadcast.

• The hypothalamo-hypophyseal tract.

FACULTY: COMPUTATIONAL.

Same transduction, performed axonally: its terminals empty into capillaries rather than onto a synapse, converting spikes into circulating hormone.

Origin (PUBLISH): magnocellular neurons of the PVN and SON.

Destination (SUBSCRIBE): the posterior pituitary, and thence the systemic bloodstream.

  https://www.linkedin.com/in/d-r-dison

Cargo: oxytocin and vasopressin.

Format: an axonal tract whose terminals empty into capillaries rather than onto a synapse.

This is the single clearest anatomical demonstration that the nervous and endocrine systems are one system.

• The retinohypothalamic tract.

FACULTY: COMPUTATIONAL.

It discards essentially all the information in the visual scene and extracts a single scalar: ambient luminance.

Aggressive, deliberate information destruction is its function, because the clock needs to know only whether it is day.

Origin (SEND): intrinsically photosensitive retinal ganglion cells.

Destination (RECEIVE): the suprachiasmatic nucleus.

Cargo: ambient light level, and nothing else.

Format: a dedicated private line, entirely separate from the image-forming visual system.

Failure: the loss of circadian entrainment; blind individuals lacking this tract are free-running, while blind individuals who retain it are not.

• The sympathetic pineal pathway.

Origin (SEND): SCN, by way of the PVN, the intermediolateral cell column, and the superior cervical ganglion.

Destination (RECEIVE): the pineal gland.

Cargo: the command to secrete melatonin.

Format: a strikingly circuitous multisynaptic route that exits the skull and re-enters it.

3.4.8 FURTHER NAMED TRACTS AND INTERFACES

The following named connections complete the catalogue of major white-matter interfaces. Each carries the send-and-receive or publish-and-subscribe label of its mode.

• The superior longitudinal fasciculus, divisions one, two, and three (SLF I/II/III).

FACULTY: TRANSMISSIVE, three parallel dorsal association channels.

Origin (SEND): the superior parietal lobule (division one), the angular gyrus (division two), and the supramarginal gyrus (division three).

  https://www.linkedin.com/in/d-r-dison

Destination (RECEIVE): the superior frontal and cingulate cortex (one), the dorsolateral prefrontal cortex (two), and the ventral premotor and Broca's area (three).

Cargo: spatial and body-centred control (one), spatial attention (two), and articulatory language and gesture (three).

• The frontal aslant tract (FAT).

FACULTY: TRANSMISSIVE, the frontal link of the language and action-control systems.

Origin (SEND): the posterior inferior frontal gyrus (Broca's region).

Destination (RECEIVE): the pre-supplementary and supplementary motor area and the medial superior frontal cortex.

Cargo: speech initiation, verbal fluency, and inhibitory and executive control.

• The extreme capsule (the temporofrontal extreme-capsule fasciculus).

FACULTY: TRANSMISSIVE, the ventral counterpart to the arcuate fasciculus.

Origin (SEND): the middle and posterior superior temporal gyrus and the insula.

Destination (RECEIVE): the inferior frontal and prefrontal cortex.

Cargo: the ventral language stream: comprehension, syntax, and verbal working memory.

• The internal capsule and the corona radiata.

FACULTY: TRANSMISSIVE, the great projection-fibre bottleneck.

Origin (SEND): the entire cortex (descending) and the thalamus (ascending).

Destination (RECEIVE): the brainstem and spinal cord (descending) and the cortex (ascending), fanning through the corona radiata.

Cargo: the master conduit of ascending sensory and descending motor projection fibres.

• The anterior commissure.

FACULTY: TRANSMISSIVE, a commissure older than the corpus callosum.

Origin (SEND): the temporal lobes, amygdala, and olfactory structures of each hemisphere.

Destination (RECEIVE): their mirror structures in the opposite hemisphere.

Cargo: interhemispheric transfer for the temporal lobe, olfaction, and emotion.

• The medial forebrain bundle (MFB).

  https://www.linkedin.com/in/d-r-dison

FACULTY: TRANSMISSIVE as a bundle, but it carries broadcast neuromodulatory fibres, hence the publish-and-subscribe label.

Origin (PUBLISH): the ventral tegmental area and brainstem monoaminergic nuclei, running through the lateral hypothalamus.

Destination (SUBSCRIBE): the lateral hypothalamus, basal forebrain, nucleus accumbens, and prefrontal and limbic cortex.

Cargo: the reward-and-motivation superhighway, the diffuse dopaminergic and monoaminergic conduit of drive and mood.

3.5 THE PREVIOUSLY UNNAMED INTERFACES

The following interfaces are real, are documented in the literature, are functionally necessary, and have never been given accepted proper names as unified constructs.

Each is marked (PUC), and a name is proposed.

The justification for naming is the same in every case, and it is worth stating once.

Neuroscience has named the tracts that anatomists could see in a dissected brain, which means it has named the compact, myelinated, macroscopic bundles.

It has systematically failed to name the interfaces that are diffuse, that are defined functionally rather than macroscopically, or that were characterized after the classical naming period ended.

This is a historical accident of method, not a fact about the brain, and it leaves working scientists without words for things they need to discuss.

3.5.1 The Temporo-Frontal Semantic Integration Pathway (PUC)

Origin (SEND): the temporal association cortex and inferotemporal cortex.

Destination (RECEIVE): the prefrontal cortex, particularly the DLPFC and inferior frontal gyrus.

Cargo: object identity and lexical meaning, delivered forward for use in reasoning and decision.

Format: feedforward corticocortical, travelling largely within the inferior fronto-occipital fasciculus and the uncinate fasciculus.

Consequence of failure: the dissociation of knowing what a thing is from being able to reason about it.

The fibre bundles that carry this traffic are named.

The functional interface they jointly constitute, meaning the delivery of semantic content from the recognition system to the executive system, is not.

3.5.2 The Temporo-Hippocampal Encoding Channel (PUC)

  https://www.linkedin.com/in/d-r-dison

Origin (SEND): the temporal association cortex, by way of the perirhinal and parahippocampal cortices.

Destination (RECEIVE): the hippocampal formation, by way of the entorhinal cortex.

Cargo: recognized objects and scenes, submitted for storage.

Consequence of failure: the failure to encode what has been perceived, with perception itself intact.

The individual relays are named.

The channel as a functional whole, meaning the route by which a recognized experience is submitted to the memory system, is not.

3.5.3 The Parieto-Frontal Spatial Planning Pathway (PUC)

Origin (SEND): the posterior parietal cortex.

Destination (RECEIVE): the premotor cortex, frontal eye fields, and DLPFC.

Cargo: spatial coordinates, body schema, and target locations, delivered for action planning.

Format: feedforward, travelling in the superior longitudinal fasciculus.

Consequence of failure: optic ataxia, in which the patient sees an object perfectly well and cannot reach for it accurately.

The superior longitudinal fasciculus is named.

The specific functional interface, meaning the delivery of "where" to the system that must act on it, is not.

3.5.4 The Cortico-Cerebellar Refinement Loop (PUC)

Origin (SEND): the cerebral cortex, by the corticopontine projection.

Destination (RECEIVE): the cerebellum, and returning by the dentato-thalamo-cortical pathway to the cortex.

Cargo: outgoing, the intended movement; returning, the correction.

Format: a closed loop.

Consequence of failure: ataxia, and the loss of skill acquisition.

Every leg of this loop is named.

The loop itself, considered as a single closed interface with an outbound and an inbound arm, is not.

3.5.5 The Cerebello-Cortical Error Correction Pathway (PUC)

Origin (SEND): the deep cerebellar nuclei.

Destination (RECEIVE): the motor cortex and the parietal cortex, by way of the VL thalamus.

  https://www.linkedin.com/in/d-r-dison

Cargo: the discrepancy between intended and actual movement, expressed as a correction.

This is the return arm of the loop above, named separately because it carries a distinct cargo and can fail independently.

3.5.6 The Hippocampo-Cortical Retrieval Network (PUC)

Origin (SEND): CA1 and the subiculum.

Destination (RECEIVE): the neocortex, by way of the entorhinal, perirhinal, and parahippocampal cortices, and by way of the retrosplenial and posterior cingulate cortex.

Cargo: the reinstatement of a stored pattern into the cortex that will experience it as a memory.

Consequence of failure: the inability to retrieve, as distinct from the inability to encode.

This is the reverse direction of the perforant path system, and it is the physical route of remembering.

That it has no name, while the forward route has one, is a striking gap.

3.5.7 The Amygdalo-Frontal Threat Assessment Channel (PUC)

Origin (SEND): the basolateral and central amygdala.

Destination (RECEIVE): the OFC, VMPFC, and ACC, by way of the ventral amygdalofugal pathway and the uncinate fasciculus.

Cargo: the flag that something is emotionally significant, delivered so that executive systems may prioritize it.

Consequence of failure: the loss of emotional input into decision-making, which is precisely the deficit of the Phineas Gage syndrome.

3.5.8 The Fronto-Amygdalar Regulatory Channel (PUC)

Origin (SEND): the VMPFC and OFC.

Destination (RECEIVE): the intercalated cell masses of the amygdala, and thence the central nucleus.

Cargo: the top-down instruction to suppress a fear response.

Format: excitatory onto inhibitory interneurons, therefore net inhibitory.

Consequence of failure: the inability to extinguish conditioned fear, which is the core deficit of post-traumatic stress disorder.

This is the return arm of 3.5.7, and it is arguably the single most clinically important nameless interface in the brain: it is the pathway that exposure therapy is understood to strengthen.

3.5.9 The Thalamic Sensory Relay Ensemble (PUC)

  https://www.linkedin.com/in/d-r-dison

Definition: the first-order sensory relay nuclei of the thalamus (LGN, MGN, VPL, VPM) considered as a single unified interface between the sensory periphery and the cortex.

Each nucleus is named.

The set of them, considered as the unified gateway through which all sensation but smell must pass to reach consciousness, has never been named as a Layer 9 anatomical ensemble.

3.5.10 The Transthalamic Corticocortical Relay (PUC)

Definition: the route by which one cortical area communicates with another not directly, but by descending to a higher-order thalamic nucleus (chiefly the pulvinar or the mediodorsal nucleus) by way of lamina 5 driving projections, and ascending again to the target cortex.

Proof of existence: higher-order thalamic nuclei receive driving input from cortical lamina 5 and project to a different cortical area.

This is now well established in the driver-modulator literature.

Proof of namelessness: the phenomenon is described.

The nuclei are named.

But the route itself, considered as a named alternative to direct corticocortical communication, has no accepted proper name, despite being one of the most theoretically consequential findings in modern systems neuroscience: it implies that a substantial fraction of cortex-to-cortex traffic is routed through the thalamus, and that the thalamus is therefore a participant in cortical computation rather than a gateway to it.

Consequence of failure: the loss of coordinated timing between cortical areas.

3.5.11 The Attentional Gating Interface (PUC)

Origin (SEND): the prefrontal cortex.

Destination (RECEIVE): the thalamic reticular nucleus.

Cargo: the instruction to suppress the thalamic relay of a specified, currently irrelevant input.

Format: excitatory onto inhibitory cells, therefore net inhibitory upon the relay nuclei.

Consequence of failure: distractibility, and the failure of selective attention.

The TRN is named.

Prefrontal projections to it are documented.

But this specific interface, which is the physical mechanism by which top-down attention suppresses distraction, has never been named, despite being one of the most sought-after mechanisms in cognitive neuroscience.

  https://www.linkedin.com/in/d-r-dison

3.5.12 The Interoceptive Ascending Interface (PUC)

Origin (SEND): the nucleus tractus solitarius.

Destination (RECEIVE): the insula, by way of the parabrachial nucleus and the VMpo thalamic nucleus.

Cargo: the physiological condition of the body, delivered to the cortex that will render it as feeling.

Consequence of failure: the loss of the felt sense of bodily state, and, on the strong reading, of the bodily basis of emotion itself.

The nuclei are named.

The ascending route is described.

The interface, considered as the unified channel by which the body reports to consciousness, is not named.

3.5.13 The Neuroendocrine Command Interface (PUC)

Origin (SEND): the hypothalamic parvocellular and magnocellular neurons.

Destination (RECEIVE): the pituitary, by way of the portal system and the hypothalamo-hypophyseal tract, and thence the entire body.

Definition: the unified interface at which a neural signal, encoded as action potentials, is transduced into an endocrine signal, encoded as hormone concentration in the blood.

Proof of existence: this transduction demonstrably occurs, and it is the mechanism by which every neuroendocrine axis is driven.

Proof of namelessness: the portal system is named.

The hypothalamo-hypophyseal tract is named.

The individual hormones are named.

But the transduction event itself, meaning the conversion of neural code into endocrine code, has never been named as a unified interface, despite being the single point at which the nervous system and the endocrine system become one system.

It is, in the vocabulary of the companion volume, the entry port to the endocrine broadcast medium.

Consequence of failure: the complete dissociation of the brain from the body's chemical state.

3.6 THE CENSUS OF COMPUTATIONAL INTERFACES

Applying the Perfect Wire Test across the catalogue above yields the following census.

These are the interfaces in which a faculty resides, and which are therefore not pipes.

  https://www.linkedin.com/in/d-r-dison

• The thalamic reticular nucleus.

FACULTY: selective attention.

It is the cleanest case in the brain.

Structurally it is pure interface: it sits astride every fibre between thalamus and cortex and sends no output to the cortex at all.

It exists only to sit in the path of other traffic, and to decide what passes.

• The optic chiasm.

FACULTY: construction of the visual hemifield, and therefore stereopsis.

• The intercalated cell masses of the amygdala.

FACULTY: fear extinction.

The prefrontal cortex cannot inhibit the central amygdala directly; it excites these GABAergic cells, which inhibit it.

The sign inversion happens in the interface.

Replace them with a perfect wire and the attempt to calm yourself would frighten you instead.

The capacity to extinguish fear exists in neither the prefrontal cortex nor the amygdala.

It exists between them.

• The dentate gyrus.

FACULTY: pattern separation.

• The inferior olive and climbing fibre.

FACULTY: generation of the error signal.

• The superior olivary complex.

FACULTY: sound localization.

Neither cochlea can localize a sound; each knows only what reached it.

The location is computed by comparison, in a structure interposed between the two ears' pathways, and the medial superior olive resolves interaural time differences of about ten microseconds, an interval shorter than an action potential.

• The hypophyseal portal system and the hypothalamo-hypophyseal tract.

FACULTY: transduction between signalling media.

• The pyramidal decussation.

FACULTY: contralateral control.

  https://www.linkedin.com/in/d-r-dison

• The corticothalamic feedback projection.

FACULTY: predictive gain control.

Outnumbering the forward projection roughly ten to one, it carries the cortex's expectations back to the thalamic gate, so that what is relayed onward is weighted by what was unexpected.

• The Schaffer collaterals.

FACULTY: storage.

An interface whose weights change with use is an interface that remembers, and an interface that remembers is not a pipe.

By contrast, the following are genuinely transmissive, and the perfect wire would substitute for each: the corticospinal tract, the optic radiation, the medial lemniscus, the fornix, the arcuate fasciculus, and the acoustic radiation.

Note the pattern, because it explains a historical blindness.

The transmissive interfaces are the long, myelinated, macroscopic bundles.

The computational interfaces are the nuclei and cell groups interposed within a pathway.

An anatomist dissecting a brain sees the bundles and names them as tracts.

The computational interfaces do not look like tracts; they look like small grey structures, and so they were classified as regions rather than as interfaces, and their interfacial nature was lost.

This forces the completion of the claim made in Section 3.1.

There it was said that a region is defined by its interfaces.

That is true, and incomplete.

The completion is this: some interfaces are themselves regions.

Region and interface are not two kinds of object at Layer 9.

They are two roles, and a single structure may occupy both.

This is not an embarrassment to the framework but the expected consequence of taking encapsulation seriously, and the engineering analogy is exact.

In a layered software architecture the hardest work is frequently done by the components that sit between the layers: the compiler, the device driver, the protocol stack.

These are neither the application nor the hardware.

They are the interface, and they are where the difficulty lives.

The brain appears to be built the same way.

  https://www.linkedin.com/in/d-r-dison

The faculty is often in the gate, not in the room.

3.7bis THE SPECIFICATION CONVENTION FOR INTERFACES

Every interface catalogued in this chapter and in the master table of Chapter 4 carries a specification block of the same four-line form used for regions and molecules.

Because interfaces fall cleanly into two dynamical kinds, the convention is stated once here and applied to all, rather than repeated forty times.

FOR EVERY TRANSMISSIVE INTERFACE (a pipe, transform absent):

SPECIFICATION BLOCK

DYNAMICS: archetype RELAY. State: signals in transit. Rule: apply the interface's DELAY (from its length and myelination, Chapter 6) and its current WEIGHT, then deliver the signal unchanged.

DRIVES: Not Applicable, for the reason that this is a transmissive interface: it conveys a signal and has no goal of its own, so it defends no set-point.

PLASTICITY: its WEIGHT (effective strength) changes slowly with use, by the three-factor rule of Chapter 10 where the pathway participates in learned behaviour; its ROUTING is fixed.

MODULATION: minimal; where a neuromodulator sets its gain, that molecule is named at its entry.

FOR EVERY COMPUTATIONAL INTERFACE (transform present, listed in the census of Section 3.6):

SPECIFICATION BLOCK

DYNAMICS: the archetype of its faculty, from Chapter 8: a GATE (the thalamic reticular nucleus, the intercalated cell masses, the optic chiasm as a routing gate), a COMPARATOR (the inferior olive, the superior olive), or an ATTRACTOR stage (the dentate gyrus as pattern separator). State and rule as given in that interface's own entry. It applies its transform, the Causal Relation Object (CRO, the reified causal claim of the companion data standard) it performs, to the signal in transit.

DRIVES: Not Applicable, for the reason that an interface serves no drive unless its transform happens to serve a named one; the intercalated cells are the exception, serving the down-regulation of the safety drive.

PLASTICITY: where the interface itself stores information (the Schaffer collaterals, the intercalated cells, the parallel-fibre-to-Purkinje synapse), its weights change by the three-factor rule with the third factor named at its entry; where it is a fixed computer (the inferior olive), Not Applicable, for the reason that a fixed computer has no learnable weights and its transform does not change with experience.

MODULATION: as named at its entry.

Thus every one of the roughly fifty interfaces has a complete specification block by this convention: a pipe is a RELAY with delay and weight, and a computer is its named archetype applying its transform.

  https://www.linkedin.com/in/d-r-dison

The individual computational interfaces whose learning matters, chiefly the memory synapses and the fear-extinction cells, carry their own explicit block in Chapter 11.

3.7 Summary: The Interface Census

The catalogue above names, in total, thirty-eight established interfaces and fourteen previously unnamed constructs, the fourteenth being the Computational Interface itself, for fifty-two in all.

The proportion is itself the finding.

Roughly one in four of the functionally necessary interfaces at Layer 9 lacks an accepted proper name, and the pattern of what is named and what is not is not random.

Neuroscience has named what a nineteenth-century anatomist could dissect: the compact, myelinated, macroscopically visible bundles.

It has left unnamed the diffuse, the functionally defined, the reciprocal return arms, and everything characterized after the classical naming period closed.

The fornix has a name because you can hold it in forceps.

The pathway by which the prefrontal cortex suppresses a distraction does not, because you cannot.

This is a methodological artefact masquerading as a fact about the brain, and the purpose of naming these constructs is to remove it.

CHAPTER 4. THE MASTER INTERFACE TABLE: FROM WHERE, TO WHERE

This chapter restates every interface in the volume in a single uniform format, so that the routing of the entire brain can be read at a glance.

The format is: NAME.

FROM: origin.

TO: destination.

CARGO: what is carried.

Constructs marked (PUC) are Previously Unnamed Constructs, named in this volume.

4.1 SENSORY INPUT INTERFACES

• Optic nerve and optic tract.

FROM: retinal ganglion cells.

TO: LGN, superior colliculus, suprachiasmatic nucleus, pretectum.

CARGO: the visual image.

  https://www.linkedin.com/in/d-r-dison

• Optic radiation.

FROM: LGN.

TO: V1, lamina 4C.

CARGO: the retinotopic image.

• Lateral olfactory tract.

FROM: olfactory bulb.

TO: piriform cortex, amygdala, entorhinal cortex.

CARGO: odour.

Uniquely, no thalamic relay.

• Acoustic radiation.

FROM: medial geniculate nucleus.

TO: A1.

CARGO: tonotopic sound.

• Brachium of the inferior colliculus.

FROM: inferior colliculus.

TO: medial geniculate nucleus.

CARGO: processed auditory signal.

• Medial lemniscus.

FROM: gracile and cuneate nuclei of the medulla.

TO: VPL thalamus.

CARGO: fine touch, vibration, proprioception.

• Spinothalamic tract.

FROM: spinal dorsal horn.

TO: VPL, intralaminar nuclei, VMpo.

CARGO: pain, temperature, crude touch.

• Trigeminal lemniscus.

FROM: trigeminal nuclei.

TO: VPM thalamus.

CARGO: facial sensation.

• Spinocerebellar tracts.

  https://www.linkedin.com/in/d-r-dison

FROM: spinal cord.

TO: cerebellum, by the inferior and superior peduncles.

CARGO: unconscious proprioception.

• Solitary tract.

FROM: vagus, glossopharyngeal, and facial nerves.

TO: nucleus tractus solitarius.

CARGO: all visceral afference, and taste.

• Retinohypothalamic tract.

FROM: intrinsically photosensitive retinal ganglion cells.

TO: suprachiasmatic nucleus.

CARGO: ambient light level only.

4.2 MOTOR OUTPUT INTERFACES

• Corticospinal (pyramidal) tract.

FROM: M1, premotor, SMA, S1, lamina 5.

TO: spinal motor neurons.

CARGO: voluntary movement commands.

• Corticobulbar tract.

FROM: motor cortex.

TO: cranial nerve motor nuclei.

CARGO: control of face, jaw, tongue, larynx.

• Rubrospinal tract.

FROM: red nucleus.

TO: spinal cord.

CARGO: upper limb flexor facilitation.

• Reticulospinal tracts.

FROM: reticular formation.

TO: spinal cord.

CARGO: postural tone, locomotion.

• Vestibulospinal tracts.

  https://www.linkedin.com/in/d-r-dison

FROM: vestibular nuclei.

TO: spinal cord.

CARGO: balance and antigravity tone.

• Tectospinal tract.

FROM: superior colliculus.

TO: cervical spinal cord.

CARGO: reflexive head orienting.

• Ventral and dorsal roots.

FROM: dorsal root ganglia (in) and ventral horn (out).

TO: spinal cord, and the muscles.

CARGO: all somatic sensation in, all somatic motor command out.

4.3 BASAL GANGLIA INTERFACES

• Corticostriatal projection.

FROM: the entire cortex, lamina 5.

TO: caudate, putamen, nucleus accumbens.

CARGO: candidate actions awaiting selection.

• Hyperdirect pathway.

FROM: motor and premotor cortex.

TO: subthalamic nucleus.

CARGO: a global stop signal.

Bypasses the striatum.

• Direct pathway.

FROM: striatal D1 neurons.

TO: GPi and SNr.

CARGO: the "go" signal, by inhibition of the inhibitor.

• Indirect pathway.

FROM: striatal D2 neurons.

TO: GPe, then STN, then GPi and SNr.

CARGO: the "no-go" signal.

• Nigrostriatal pathway.

  https://www.linkedin.com/in/d-r-dison

FROM: substantia nigra pars compacta.

TO: caudate and putamen.

CARGO: dopamine, encoding reward prediction error.

• Mesolimbic pathway.

FROM: ventral tegmental area.

TO: nucleus accumbens.

CARGO: incentive salience, "wanting."

• Mesocortical pathway.

FROM: ventral tegmental area.

TO: prefrontal cortex.

CARGO: working memory gain control.

• Pallidothalamic tract (ansa lenticularis and lenticular fasciculus).

FROM: GPi.

TO: VA and VL thalamus.

CARGO: tonic inhibition, selectively released.

• Thalamostriatal projection.

FROM: intralaminar thalamic nuclei.

TO: striatum.

CARGO: attentional and arousal signals.

• Stria medullaris.

FROM: GPi and limbic forebrain.

TO: habenula.

CARGO: the disappointment signal.

• Fasciculus retroflexus.

FROM: lateral habenula.

TO: VTA (via RMTg) and raphe nuclei.

CARGO: inhibition of both monoamine systems.

4.4 MEMORY AND LIMBIC INTERFACES

• Perforant path.

  https://www.linkedin.com/in/d-r-dison

FROM: entorhinal cortex, layer II.

TO: dentate gyrus and CA3.

CARGO: the convergent output of the entire association cortex.

• Temporoammonic path.

FROM: entorhinal cortex, layer III.

TO: CA1 directly.

CARGO: current sensory reality.

A within-structure shortcut.

• Mossy fibre pathway.

FROM: dentate granule cells.

TO: CA3.

CARGO: sparse, pattern-separated representations.

• Schaffer collaterals.

FROM: CA3.

TO: CA1.

CARGO: the pattern-completed reconstruction.

• Fornix.

FROM: subiculum and CA1.

TO: mammillary bodies, anterior thalamus, septal nuclei, nucleus accumbens.

CARGO: hippocampal output.

• Mammillothalamic tract.

FROM: mammillary bodies.

TO: anterior thalamic nuclei.

CARGO: the Papez circuit signal.

• Cingulum bundle.

FROM: cingulate and medial prefrontal cortex.

TO: parahippocampal cortex and hippocampus, reciprocally.

CARGO: default mode and Papez traffic.

• Stria terminalis.

FROM: amygdala.

  https://www.linkedin.com/in/d-r-dison

TO: BNST, hypothalamus, septal nuclei.

CARGO: emotional signals to subcortical targets.

• Ventral amygdalofugal pathway.

FROM: central and basolateral amygdala.

TO: mediodorsal thalamus, hypothalamus, brainstem, nucleus accumbens, prefrontal and cingulate cortex.

CARGO: emotional signals to cortical targets.

• Uncinate fasciculus.

FROM: orbitofrontal cortex.

TO: anterior temporal lobe and amygdala, reciprocally.

CARGO: emotional regulation and valenced memory.

• Thalamo-amygdalar projection (the "low road").

FROM: medial geniculate and posterior thalamus.

TO: amygdala.

CARGO: a fast, crude threat signal, arriving before cortical identification.

4.5 CEREBELLAR INTERFACES

• Corticopontine projection.

FROM: the entire cerebral cortex.

TO: pontine nuclei.

CARGO: the intended action.

Roughly twenty million fibres.

• Pontocerebellar projection (middle cerebellar peduncle).

FROM: pontine nuclei.

TO: contralateral cerebellar cortex, as mossy fibres.

CARGO: intention and context.

• Olivocerebellar projection (inferior cerebellar peduncle).

FROM: inferior olivary nucleus.

TO: cerebellar cortex, as climbing fibres.

CARGO: the error signal.

  https://www.linkedin.com/in/d-r-dison

One fibre per Purkinje cell.

• Dentato-thalamo-cortical pathway (superior cerebellar peduncle).

FROM: dentate nucleus.

TO: VL thalamus, then motor and prefrontal cortex.

CARGO: the correction.

• Central tegmental tract.

FROM: red nucleus.

TO: inferior olive.

CARGO: the efference copy against which error is computed.

4.6 CORTICOCORTICAL INTERFACES

• Corpus callosum.

FROM: each cortical hemisphere.

TO: the homologous region of the other.

CARGO: interhemispheric integration.

Roughly two hundred million fibres.

• Arcuate fasciculus.

FROM: Wernicke's area.

TO: Broca's area, reciprocally.

CARGO: phonological representations for repetition.

• Superior longitudinal fasciculus.

FROM: parietal and temporal cortex.

TO: frontal cortex.

CARGO: spatial attention and working memory traffic.

• Inferior fronto-occipital fasciculus.

FROM: occipital and temporal cortex.

TO: frontal cortex.

CARGO: the ventral stream carried forward.

• Inferior longitudinal fasciculus.

FROM: occipital cortex.

TO: anterior temporal lobe.

  https://www.linkedin.com/in/d-r-dison

CARGO: visual object information toward memory.

• Ventral stream (occipitotemporal).

FROM: V1 and V4.

TO: inferotemporal cortex.

CARGO: object identity, the "what."

• Dorsal stream (occipitoparietal).

FROM: V1 and V5/MT.

TO: posterior parietal cortex.

CARGO: location and action guidance, the "where and how."

• Corticothalamic feedback (lamina 6).

FROM: every cortical area.

TO: its source thalamic nucleus.

CARGO: expectation and gain control.

Outnumbers the forward projection roughly ten to one.

4.7 NEUROMODULATORY BROADCAST INTERFACES

• Ascending noradrenergic bundle.

FROM: locus coeruleus.

TO: essentially the entire brain.

CARGO: arousal, vigilance, and surprise.

• Ascending serotonergic projections.

FROM: raphe nuclei.

TO: essentially the entire brain, and the spinal dorsal horn.

CARGO: mood, impulse control, pain gating.

• Basal forebrain cholinergic projection.

FROM: nucleus basalis of Meynert.

TO: the entire cerebral cortex.

CARGO: attention, and the signal that something is worth learning.

• Septohippocampal projection.

FROM: medial septal nucleus and diagonal band, by the fornix.

  https://www.linkedin.com/in/d-r-dison

TO: hippocampus.

CARGO: acetylcholine, and the theta rhythm pacemaker.

• Tuberomammillary histaminergic projection.

FROM: tuberomammillary nucleus.

TO: the entire cortex.

CARGO: wakefulness.

• Orexinergic projections.

FROM: lateral hypothalamic orexin neurons.

TO: locus coeruleus, VTA, raphe, tuberomammillary nucleus, basal forebrain.

CARGO: the stabilization of the wake state.

4.8 THALAMIC AND ATTENTIONAL INTERFACES

• Thalamocortical radiations.

FROM: each thalamic relay nucleus.

TO: its cortical target, lamina 4.

CARGO: the driving sensory or motor message.

• Reticulothalamic inhibition.

FROM: thalamic reticular nucleus.

TO: the thalamic relay nuclei, exclusively.

CARGO: inhibitory gating.

The TRN projects to the cortex not at all.

• Ascending reticular activating system.

FROM: the reticular formation and the arousal nuclei.

TO: the intralaminar thalamus and, diffusely, the cortex.

CARGO: the level of consciousness itself.

4.9 NEUROENDOCRINE INTERFACES

• Hypophyseal portal system.

FROM: hypothalamic parvocellular neurons, at the median eminence.

TO: the anterior pituitary.

CARGO: the releasing and inhibiting hormones.

Format: a private capillary bed, not an axon.

  https://www.linkedin.com/in/d-r-dison

• Hypothalamo-hypophyseal tract.

FROM: magnocellular neurons of the PVN and SON.

TO: the posterior pituitary, and thence the systemic bloodstream.

CARGO: oxytocin and vasopressin.

• Sympathetic pineal pathway.

FROM: SCN, via PVN, the intermediolateral cell column, and the superior cervical ganglion.

TO: the pineal gland.

CARGO: the command to secrete melatonin.

• Sympathetic preganglionic outflow.

FROM: the intermediolateral cell column.

TO: the sympathetic chain and prevertebral ganglia, and the adrenal medulla directly.

CARGO: sympathetic command.

Note that the adrenal medulla is innervated directly, with no ganglionic relay, which is why adrenaline release is so fast.

• Vagal efferent outflow.

FROM: the dorsal motor nucleus of the vagus and nucleus ambiguus.

TO: the heart, lungs, and gut.

CARGO: parasympathetic command.

4.10 THE PREVIOUSLY UNNAMED INTERFACES, IN FROM-TO FORM

• Temporo-Frontal Semantic Integration Pathway (PUC).

FROM: temporal and inferotemporal association cortex.

TO: prefrontal cortex.

CARGO: object identity and lexical meaning, for reasoning.

• Temporo-Hippocampal Encoding Channel (PUC).

FROM: temporal association cortex, via perirhinal and parahippocampal cortex.

TO: hippocampus, via entorhinal cortex.

CARGO: recognized experience, submitted for storage.

• Parieto-Frontal Spatial Planning Pathway (PUC).

  https://www.linkedin.com/in/d-r-dison

FROM: posterior parietal cortex.

TO: premotor cortex, frontal eye fields, DLPFC.

CARGO: spatial coordinates for action.

• Cortico-Cerebellar Refinement Loop (PUC).

FROM: cerebral cortex.

TO: cerebellum, and back to cortex.

CARGO: intention out, correction back.

• Cerebello-Cortical Error Correction Pathway (PUC).

FROM: deep cerebellar nuclei.

TO: motor and parietal cortex, via VL thalamus.

CARGO: the discrepancy between intended and actual movement.

• Hippocampo-Cortical Retrieval Network (PUC).

FROM: CA1 and subiculum.

TO: the neocortex, via entorhinal, retrosplenial, and posterior cingulate cortex.

CARGO: the reinstatement of a stored pattern.

This is the physical route of remembering.

• Amygdalo-Frontal Threat Assessment Channel (PUC).

FROM: basolateral and central amygdala.

TO: OFC, VMPFC, ACC.

CARGO: the flag of emotional significance.

• Fronto-Amygdalar Regulatory Channel (PUC).

FROM: VMPFC and OFC.

TO: the intercalated cell masses, and thence the central amygdala.

CARGO: the instruction to suppress fear.

This is the pathway strengthened by exposure therapy.

• Thalamic Sensory Relay Ensemble (PUC).

FROM: the sensory periphery.

TO: the cortex.

CARGO: all sensation but smell.

  https://www.linkedin.com/in/d-r-dison

The set of LGN, MGN, VPL, and VPM considered as one gateway.

• Transthalamic Corticocortical Relay (PUC).

FROM: cortical area A, lamina 5.

TO: cortical area B, via a higher-order thalamic nucleus (pulvinar or MD).

CARGO: cortex-to-cortex traffic routed through the thalamus rather than directly.

• Attentional Gating Interface (PUC).

FROM: prefrontal cortex.

TO: thalamic reticular nucleus.

CARGO: the instruction to suppress an irrelevant thalamic relay.

The mechanism of top-down attention.

• Interoceptive Ascending Interface (PUC).

FROM: nucleus tractus solitarius.

TO: the insula, via the parabrachial nucleus and VMpo thalamus.

CARGO: the physiological condition of the body, delivered to be felt.

• Neuroendocrine Command Interface (PUC).

FROM: hypothalamic neurons.

TO: the pituitary, and thence the entire body.

CARGO: a neural code, transduced into an endocrine code.

The point at which the two systems become one.

CHAPTER 5. THE RHYTHMS: HOW REGIONS TALK IN TIME

5.1 Why a Chapter on Rhythm

The preceding chapters treated the brain as an architecture of places and wires.

That is true, and it is incomplete, because it leaves out the dimension in which the architecture actually operates, which is time.

Regions do not merely connect.

They connect on a schedule.

The brain is a rhythmic organ, and its rhythms are not a byproduct of its activity but a mechanism of its coordination.

Two regions that must exchange information do so by oscillating together, and two regions that must be kept apart are held at incompatible rhythms.

Rhythm is how the connectome, a fixed anatomical object, becomes a dynamic one: the wires are constant, but which wires are open for traffic at any instant is set by the phase of an oscillation.

This chapter adds the temporal dimension to the spatial catalogue, and it is essential to the framework, because a Conduit (in the vocabulary of the companion data standard) that is anatomically present but rhythmically closed is not conducting, and a region's faculty may depend as much on when it fires as on what it connects to.

5.2 The Frequency Bands, and What Each Is For

Neural oscillations are conventionally divided into bands, and each band has a characteristic spatial reach and a characteristic function.

The division is not arbitrary: slow oscillations travel far and coordinate large territories, while fast oscillations are local and perform computation.

This is a general principle worth stating plainly, because it recurs at every band: LOW FREQUENCY IS FOR LONG-RANGE COMMUNICATION, HIGH FREQUENCY IS FOR LOCAL COMPUTATION.

• Delta, roughly 0.5 to 4 hertz.

The slowest and largest.

Dominant in deep sleep, where its slow oscillation, below 1 hertz, orchestrates the Up and Down states of entire cortical territories.

It is the master clock of memory consolidation, and it coordinates the dialogue between hippocampus and cortex during the night.

• Theta, roughly 4 to 8 hertz.

The rhythm of the hippocampus during active exploration and of the cortex during the long-range coordination of cognition.

Theta is the carrier wave of memory: items are encoded and retrieved at specific phases of the theta cycle, and the hippocampal theta rhythm is paced from the medial septum, the structure described in Section 2.2.7.

When theta stops, so does the orderly writing of experience into memory.

• Alpha, roughly 8 to 13 hertz.

The rhythm of the resting, relaxed cortex, and, crucially, the rhythm of INHIBITION.

Alpha is not idling; it is active suppression.

When you close your eyes, occipital alpha surges, and it represents the visual cortex being actively silenced.

Alpha is how the brain turns a region down, and its power rises over cortex that is being told to be quiet so that resources go elsewhere.

This is attention implemented as rhythm.

  https://www.linkedin.com/in/d-r-dison

• Beta, roughly 13 to 30 hertz.

The rhythm of the engaged, attentive, and specifically the STATUS-QUO-MAINTAINING cortex.

Beta dominates the motor system when a posture is being held, and its abnormal excess in the basal ganglia is the signature of Parkinsonian rigidity: the beta rhythm is the brain's way of saying "keep things as they are," and in Parkinson's disease it says so pathologically, and movement freezes.

• Gamma, roughly 30 to 100-plus hertz.

The fast, local rhythm of active computation.

Gamma is generated by the interplay of excitatory pyramidal cells and fast-spiking parvalbumin interneurons (Section 2.2.1 and 2.2.4 of the cellular volume), and it is present wherever cortex is actively processing.

Gamma is the leading candidate solution to the binding problem: the puzzle of how the redness, the roundness, and the motion of a ball, processed in separate cortical areas, become one experienced ball.

The proposed answer is that the neurons representing features of the same object fire together in gamma synchrony, and that this temporal togetherness is what binds them into a unified percept.

5.3 The Master Mechanism: Cross-Frequency Coupling

The single most important idea in this chapter, and the one least known outside specialist circles, is cross-frequency coupling.

The bands do not operate in isolation.

They nest.

The phase of a slow oscillation controls the amplitude of a fast one, so that fast local computation happens only during permitted windows of the slow long-range rhythm.

The canonical example is theta-gamma coupling in the hippocampus.

Gamma bursts, each carrying one item of information, are slotted into successive phases of the slower theta cycle.

This produces a physical mechanism for working memory capacity: roughly seven gamma cycles fit into one theta cycle, and this has been proposed as the deep reason that working memory holds about seven items.

The famous "magical number seven" may be a consequence of the ratio of two brain rhythms.

The general principle is this.

Cross-frequency coupling is how the brain reconciles its two competing needs: fast local computation and slow global coordination.

The slow rhythm is the postal schedule; the fast rhythm is the letters.

  https://www.linkedin.com/in/d-r-dison

A letter written outside the scheduled window does not get delivered.

For the framework of this volume, cross-frequency coupling is the temporal implementation of the interface.

A Conduit conducts when the sending and receiving regions are coupled at the right phase, and does not when they are not.

The anatomy says two regions CAN talk; the rhythm says whether they ARE talking, right now.

5.4 Communication Through Coherence

The theory that unifies this chapter is called communication through coherence, and it states the mechanism precisely.

Two regions exchange information effectively when their oscillations are coherent, meaning their excitable phases line up, so that spikes sent from one arrive at the other during its window of maximum receptivity.

When the phases are misaligned, the same anatomical connection carries little effective information, because the arriving spikes land during the receiver's refractory trough and are ignored.

This means the effective connectivity of the brain, the pattern of who is actually influencing whom, is a dynamic subset of its structural connectivity, the pattern of who is anatomically wired to whom, and the selection is made by rhythm.

This is a profound addition to the connectional-fingerprint principle of Chapter 1.

A region is defined by its interfaces, yes, but which of its interfaces are live at any moment is set by the coherence structure of the ongoing oscillations.

The map is fixed.

The traffic is rhythmic.

5.5 Sleep: The Architecture Rebuilt Nightly

Nowhere is the rhythmic reorganization of the brain more dramatic than in sleep, which is not an off state but a sequence of distinct, actively generated global rhythms, each doing different work.

THE STAGES

• Stage N1.

The drowsy transition.

Alpha gives way to theta, and vertex sharp waves appear.

This is the border of consciousness, and it is where hypnic jerks and hallucinatory images arise.

• Stage N2.

  https://www.linkedin.com/in/d-r-dison

Light sleep, and the hallmark of it is two events: the sleep spindle, a brief 11 to 16 hertz burst generated by the thalamic reticular nucleus (Section 2.3.1.9), and the K-complex, a large solitary slow wave.

Spindles are now understood to gate memory: they are the thalamus briefly rehearsing and protecting a memory trace, and their density predicts overnight memory improvement.

• Stage N3.

Slow-wave sleep, deep sleep, dominated by the delta slow oscillation below 1 hertz.

This is the stage of greatest restoration and of the most powerful memory consolidation.

• REM sleep.

Rapid eye movement sleep, with a paradoxical waking-like EEG of mixed frequency, vivid dreaming, theta oscillation, and near-total muscle atonia, the paralysis that keeps you from acting out your dreams.

Its failure is REM sleep behaviour disorder, and that disorder is one of the strongest early predictors of future Parkinson's disease.

THE CONSOLIDATION TRIAD

The central discovery of modern sleep science is that memory is consolidated during N3 by the precise nesting of three rhythms, one from each of three structures, in a strict temporal hierarchy:

The cortical slow oscillation (the delta Up state) sets the master timing.

Nested within its Up states come thalamic sleep spindles.

And nested within the spindles come hippocampal sharp-wave ripples, brief 150 to 250 hertz bursts during which the day's experiences are replayed at high speed, compressed roughly twenty-fold, and handed from hippocampus to cortex for permanent storage.

Slow oscillation, then spindle, then ripple, in that order, in a window of a few hundred milliseconds, hundreds of times a night.

This triple nesting is the physical event of a memory moving from its temporary hippocampal store to its permanent cortical home.

It is the Hippocampo-Cortical Retrieval Network (PUC) of Section 3.5.6 running in reverse and at night, and it is arguably the most important thing the sleeping brain does.

THE SWITCH

Sleep and waking are not a dimmer but a switch, and the switch is the flip-flop of Section 2.3.2: the sleep-promoting ventrolateral preoptic nucleus and the wake-promoting arousal nuclei mutually inhibit one another, producing two stable states with a fast transition between them, stabilized by orexin.

  https://www.linkedin.com/in/d-r-dison

A flip-flop has no comfortable middle, which is why you are usually either clearly asleep or clearly awake, and why the loss of orexin (narcolepsy) makes the switch flip uncontrollably, dropping a person into sleep or into REM atonia (cataplexy) without warning.

THE NIGHTLY CLEANING

Sleep is also when the brain is physically washed.

The glymphatic system (Section 2.11 of the whole-system volume) increases its flow dramatically during slow-wave sleep, as the interstitial space between neurons expands by roughly sixty percent, and cerebrospinal fluid pulses through and flushes out the metabolic waste of the day, including the amyloid beta whose accumulation drives Alzheimer's disease.

This is a leading candidate answer to why sleep is not optional, and why chronic sleep loss is a risk factor for dementia: the brain that does not sleep does not clean, and the debris that would have been washed away accumulates.

5.6 The Rhythmic Signatures of Disease

The rhythms are diagnostic, and their disruption is not incidental to brain disease but often central to it.

• Parkinson's disease is a disease of pathological beta synchrony in the basal ganglia.

The subthalamic nucleus and globus pallidus lock into excessive beta oscillation, and this is why deep brain stimulation works: high-frequency stimulation disrupts the pathological rhythm.

The therapy is, quite literally, jamming a broken clock.

• Epilepsy is the pathology of runaway synchrony: the failure of the inhibitory restraint that normally keeps excitation local, so that a hypersynchronous discharge spreads across tissue that should have stayed independent.

A seizure is the encapsulation of Chapter 1 catastrophically failing in the time domain.

• Schizophrenia is increasingly understood as a disorder of gamma and of the parvalbumin interneurons that generate it.

Reduced gamma synchrony may underlie the failure to bind experience coherently, and the thalamic reticular nucleus, the attentional gate, is implicated in the failure to filter irrelevant input.

• Alzheimer's disease disrupts the sleep rhythms that clear it, producing a vicious loop: less slow-wave sleep means less amyloid clearance means more amyloid means less slow-wave sleep.

Experimental work driving 40-hertz gamma with light and sound to clear amyloid is a direct attempt to intervene in the rhythmic dimension.

The lesson for the framework is that a complete Layer 9 account cannot be purely anatomical.

  https://www.linkedin.com/in/d-r-dison

The regions and interfaces of Chapters 2 through 4 are the instrument; the rhythms of this chapter are the music; and much of brain disease is the music going wrong while the instrument stays intact.

CHAPTER 6. THE NUMBERS: A QUANTITATIVE GROUNDING

6.1 Why Numbers Matter

A framework that is to become a working model, and eventually a running system, cannot rest on qualitative description alone.

It needs magnitudes.

This chapter collects the quantitative facts that ground the anatomy of the preceding chapters, because the difference between a caricature and a model is that a model has numbers, and because many of these numbers are themselves surprising and load-bearing.

The numbers below are drawn from the current literature and are given as the consensus figures a working scientist would use.

Where a figure is contested or approximate, it is marked as such.

6.2 The Whole Brain

• Neurons in the human brain: approximately 86 billion.

The older figure of 100 billion was an estimate; the 86 billion figure comes from actually counting, by dissolving brains into a uniform suspension and counting nuclei, a method called the isotropic fractionator.

• Glial cells: roughly 85 billion, giving a glia-to-neuron ratio near 1 to 1, not the 10 to 1 long asserted in textbooks.

This correction matters, and it overturned a century of received wisdom.

• Of the 86 billion neurons, roughly 69 billion, about 80 percent, are in the CEREBELLUM, not the cerebrum.

The cerebellum contains more neurons than all the rest of the brain combined, a fact that sits uneasily with its traditional demotion to a mere motor structure, and which is part of the modern case for its role in cognition.

• The cerebral cortex contains roughly 16 billion neurons.

• Synapses: on the order of 100 trillion to 1 quadrillion.

A single cortical pyramidal neuron may bear tens of thousands of synapses.

• Mass: roughly 1.4 kilograms, about 2 percent of body weight.

• Energy: the brain consumes roughly 20 percent of the body's oxygen and glucose at rest.

This is a tenfold overrepresentation relative to its mass, and most of it is spent not on task-evoked activity but on the ongoing intrinsic activity

  https://www.linkedin.com/in/d-r-dison

of the resting brain, which is why the default mode network was such a surprise.

• Total axonal length in the human brain has been estimated in the range of a hundred thousand to hundreds of thousands of kilometres.

The wiring, laid end to end, would circle the Earth several times.

6.3 The Cellular and Synaptic Scale

• Resting membrane potential: approximately negative 70 millivolts.

• Action potential amplitude: approximately 100 millivolts, from negative 70 to about positive 30 and back, in roughly 1 to 2 milliseconds.

• Conduction velocity: from about 0.5 metres per second in unmyelinated fibres to about 120 metres per second in the thickest myelinated ones, which is roughly 430 kilometres per hour.

Myelination buys a hundredfold speedup.

• Synaptic delay: roughly 0.5 to 1 millisecond, the time for a chemical synapse to release, diffuse, and bind.

• Synaptic cleft width: roughly 20 nanometres.

• Number of neurotransmitter molecules in one vesicle: a few thousand.

• Firing rates: cortical neurons idle at fractions of a hertz to a few hertz, and fire in bursts up to hundreds of hertz.

Purkinje cells of the cerebellum fire tonically at a remarkable 50 to 100 hertz continuously.

6.4 The Neuromodulatory Nuclei: Small Sources, Vast Reach

The disproportion between the size of these nuclei and their influence is one of the most striking quantitative facts in the brain, and it is the anatomical basis of the broadcast principle.

• Locus coeruleus: roughly 15,000 noradrenergic neurons per side in humans, perhaps 50,000 total, and their axons reach essentially the entire brain.

A rounding error's worth of cells sets the arousal of the whole organ.

• Dopaminergic neurons: only about 400,000 to 600,000 in the entire human midbrain.

Parkinson's disease becomes symptomatic when roughly 60 to 80 percent of the substantia nigra pars compacta neurons are already dead, meaning the system has enormous reserve and the disease is far advanced before it is visible.

This reserve is why early diagnosis is so difficult.

• Serotonergic neurons of the raphe: on the order of a few hundred thousand, projecting brain-wide and to the spinal cord.

  https://www.linkedin.com/in/d-r-dison

• Histaminergic neurons of the tuberomammillary nucleus: only a few thousand, yet they help set the entire brain's wakefulness.

This is why antihistamines that cross into the brain cause such profound drowsiness: a few thousand cells, blocked, and consciousness dims.

The general fact: the systems that set the global state of the brain are among the smallest populations in it.

Power in this architecture is not proportional to size.

It is proportional to reach.

6.5 The Dopamine Signal, Quantified

Because it is the most computationally characterized neuromodulator, dopamine deserves a quantitative note of its own, and recent work has refined the classic picture.

The classic finding is that midbrain dopamine neurons encode a reward prediction error: they burst above baseline when reward exceeds expectation, stay silent when reward exactly matches expectation, and pause below baseline when reward falls short.

This is the biological implementation of the temporal-difference learning signal from reinforcement learning, and it is one of the clearest bridges between a brain measurement and a computational theory.

The refinement is twofold.

First, dopamine operates in two modes: a fast PHASIC burst carrying the prediction error, and a slow TONIC level now thought to track the ongoing value of the current state, rising as an animal approaches a reward in space or time.

Second, the population is not uniform.

Medial ventral tegmental area neurons carry the classic prediction-error signal, while lateral neurons and the substantia nigra carry a SALIENCE signal that responds to important events regardless of whether they are good or bad, including aversive ones.

The single label "the dopamine reward signal" conceals at least two distinct computations in two distinct populations, and a working model must represent both.

6.6 The Cerebellar Transform, Quantified

The cerebellum's uniformity is quantitative and is the key to its computational interpretation.

Its circuit is repeated with extraordinary regularity: each Purkinje cell receives input from roughly 150,000 to 200,000 parallel fibres, and exactly one climbing fibre.

That ratio, hundreds of thousands to one, is the anatomical signature of its learning rule: a vast number of weak, adjustable inputs (the parallel fibres, carrying context and intention) and one overwhelming teaching input (the climbing fibre, carrying error).

  https://www.linkedin.com/in/d-r-dison

The modern interpretation, the universal cerebellar transform, is that the cerebellum applies ONE computation everywhere: it builds a forward model, a predictor of the sensory consequences of an action, and reports the discrepancy when the prediction fails.

What differs across the cerebellum is only what is plugged into the identical circuit.

Plug in motor commands and it predicts and corrects movement.

Plug in cognition, via the loops with prefrontal cortex, and it predicts and corrects thought.

This is why cerebellar damage produces not only ataxia of movement but, in the posterior cerebellum, the cerebellar cognitive affective syndrome: a dysmetria of thought and emotion, the mental equivalent of overshooting a target.

For the framework, the cerebellum is the clearest existing proof of the volume's deepest thesis: one canonical computation, replicated and rewired to many purposes, its function set entirely by its interfaces.

It is Chapter 1 written in a single structure.

6.7 The Networks, Quantified

The intrinsic networks of Section 2.10 are defined quantitatively, by correlation.

The brain at rest is not silent.

It is organized into a small number of large-scale networks whose activity waxes and wanes coherently, and the triple-network model captures the three most consequential:

• The default mode network is active at rest and during internally directed thought, and it DEACTIVATES during externally directed tasks.

It and the central executive network are ANTICORRELATED: when one rises, the other falls.

This anticorrelation is a fundamental organizing fact of brain dynamics, and its breakdown is a marker across psychiatric disorders.

• The central executive network, anchored in dorsolateral prefrontal and posterior parietal cortex, is active during externally directed, effortful cognition.

• The salience network, anchored in the anterior insula and dorsal anterior cingulate, is the SWITCH.

It detects behaviourally salient events and toggles the brain between the default and executive modes.

The anterior insula is thought to be the specific structure that initiates the switch, and its dysfunction is implicated in the failure to disengage from internal preoccupation that characterizes depression, or from irrelevant stimuli in attention disorders.

  https://www.linkedin.com/in/d-r-dison

The quantitative point is that these networks are defined not by anatomy alone but by the temporal correlation of their activity, which returns us to the theme of Chapter 5: the functional brain is a dynamic object, and its units are as much rhythmic and correlational as they are anatomical.

6.8 A Closing Number

One figure deserves to stand alone.

The human cerebral cortex contains roughly 16 billion neurons, and each is connected, on average, to thousands of others, yielding a connectivity on the order of a hundred trillion synapses in the cortex alone.

The number of possible states of such a system is not merely astronomical; it exceeds, by an incomprehensible margin, the number of atoms in the observable universe.

This is the physical object that this volume has attempted to map at the level of its regions and interfaces.

The map is coarse, deliberately: Layer 9 is the scale of the hundred regions and the forty tracts, not the hundred trillion synapses.

But it is worth ending the quantitative chapter by remembering what the map is a map OF.

The territory is the most complex object known to exist, and the fact that it can be described at all, at any level, in a finite number of named parts with named connections, is itself the central bet of the framework: that complexity this vast is nonetheless LAYERED, and that a layered thing can be understood one level at a time.

CHAPTER 7. THE MOLECULES: THE CHEMICAL MESSENGERS OF THE BRAIN

7.1 Why the Molecules Are a Layer of Their Own

The preceding chapters described the brain as regions, interfaces, and rhythms.

That description is incomplete in a way that this chapter repairs, because it leaves out the substance that actually travels along the interfaces.

When one region influences another, something physical crosses the gap, and that something is a molecule.

In the vocabulary of the companion framework, these molecules live at Layer 3, the ionic and simple molecular layer, and at Layer 4, the macromolecular layer, well below the Layer 9 regions that are the subject of this volume.

That is precisely why they demand a chapter here.

A chemical messenger is the clearest case in the whole framework of a lower-layer entity whose behaviour cannot be understood without reference to the regions it connects.

Dopamine is a small Layer 3 molecule, but its meaning is a Layer 9 and Layer 10 fact: what dopamine DOES depends entirely on which region released

  https://www.linkedin.com/in/d-r-dison

it, which region received it, and which of a dozen receptor subtypes was waiting.

The molecule is simple; its significance is distributed across the map.

This is the reason the earlier working title of this material, which referred only to Layer 9, was set aside.

Dopamine and its cousins are too important, and too cross-cutting, to be excluded on the grounds that they are technically molecules.

They are the medium in which the regions and interfaces of this volume conduct their business, and any model that represents the regions but not the messengers will be a model of a switchboard with the current turned off.

A NECESSARY DISTINCTION, STATED ONCE

Three overlapping terms are used loosely in the literature and must be separated, because the distinction is functional and it will matter for any model built from this material.

A NEUROTRANSMITTER is released at a synapse, crosses a cleft roughly twenty nanometres wide, and acts on the immediately opposing cell, fast and privately.

This is wiring transmission, and it is the point-to-point telephone call of the nervous system.

A NEUROMODULATOR is released more diffusely, often from varicosities along an axon rather than from a classical synapse, and it spreads through the extracellular space to act on many cells at once, more slowly and more persistently.

This is volume transmission, and it is the radio broadcast.

The same molecule can be a transmitter in one place and a modulator in another; dopamine, serotonin, and acetylcholine all do double duty.

A HORMONE is released into the bloodstream and acts on any cell in the body bearing the right receptor, slowest and most broadly of all.

Some molecules are simultaneously all three: norepinephrine is a neurotransmitter at a synapse, a neuromodulator in a diffuse cortical projection, and a hormone when the adrenal medulla dumps it into the blood.

These are not three kinds of molecule.

They are three MODES of use, and a molecule's mode is a fact about the interface it is using, not about the molecule itself.

This is the single most important idea in the chapter, and it recurs throughout: the molecule is the same; the mode is set by the anatomy.

7.2 HOW TO READ THIS CATALOGUE

Each messenger is given under a uniform set of headings, so the entries can be compared and, later, modelled.

  https://www.linkedin.com/in/d-r-dison

ORIGIN (PUBLISH): where it is synthesized, and by which cells or regions.

RECEPTORS: the receptor families it acts on, and whether they are ionotropic (fast ion channels) or metabotropic (slow G-protein-coupled receptors).

This distinction governs speed and is the first thing a model must represent.

TARGETS AND EFFECTS (SUBSCRIBE): which regions it reaches and what it does to them.

FUNCTION AND FACULTY: the behaviour or capacity it serves.

WHEN IT FAILS: the clinical consequence of too much or too little, which is often the clearest evidence of the normal function.

DIGITAL NOTE: a preliminary observation on how the messenger might be represented in a complex digital system.

These notes are deliberately tentative.

They are not a design; they are the beginning of one, and they are included because the eventual purpose of this catalogue is to be modelled.

WHY THE ORIGIN AND TARGET FIELDS ARE LABELLED PUBLISH AND SUBSCRIBE

Two of these headings carry a parenthetical borrowed from computer and information science, where the terms have been used for decades: ORIGIN is marked (PUBLISH), and TARGETS AND EFFECTS is marked (SUBSCRIBE).

The reason is that a chemical messenger behaves, at the level of the whole system, exactly like a publish-and-subscribe channel, the same pattern used in message-oriented software since at least the early information-highway systems of the turn of the century.

A source region PUBLISHES a molecule: it synthesizes and releases the substance without addressing any particular recipient, and that is what the ORIGIN field records.

Any region bearing the right receptor SUBSCRIBES to that molecule: it receives the broadcast and is changed by it, and that is what the TARGETS AND EFFECTS field records.

The publisher does not know its subscribers; the subscribers are defined by the receptors they carry, not by a private wire the publisher ran to them.

This is precisely the volume-transmission broadcast described earlier in this chapter, and it is the deep reason a neuromodulator's effect is set by the anatomy of its subscribers rather than by the molecule itself.

The labels are not decoration.

They are placed here so that a coding-assistant artificial-intelligence system building the eventual digital model will recognize the messaging semantics at once, and will map ORIGIN onto a publisher and TARGETS AND EFFECTS onto its subscribers, rather than having to rediscover from prose that the molecules of the brain are a publish-and-subscribe fabric.

  https://www.linkedin.com/in/d-r-dison

The same publish-and-subscribe pattern can be seen in the interfaces of Chapter 3; whether to label it there in the same way is a separate question, considered at the end of this volume.

The messengers are grouped into eight classes: the small-molecule amino-acid transmitters, the monoamines, acetylcholine, the neuropeptides, the gasotransmitters, the endocannabinoids, the purines, and the trophic and immune signals.

The endocrine hormones that act on the brain are treated in Section 7.11.

7.3 CLASS ONE: THE AMINO-ACID TRANSMITTERS (THE FAST WORKHORSES)

These four carry the overwhelming majority of fast synaptic traffic in the brain.

If the monoamines are the volume knobs, these are the signal itself.

A NOTE ON THE PUBLISH-AND-SUBSCRIBE LABELS FOR THESE FOUR: strictly, these fast amino-acid transmitters are wiring transmission, the point-to-point telephone call of the nervous system, not the diffuse broadcast of the monoamines, so at the level of a single synapse their truest label is send-and-receive rather than publish-and-subscribe.

They keep the publish-and-subscribe headings here for one deliberate reason: this volume models the major cognitive regions, not the eighty-six billion individual neurons, and at the level of a whole region a source population publishes glutamate or gamma-aminobutyric acid to a target population that subscribes to it, which is the resolution the digital system is built at.

Where single-synapse fidelity matters, read these four as send-and-receive; at the region scale this volume works in, the publish-and-subscribe label holds.

GLUTAMATE

ORIGIN (PUBLISH): synthesized from glutamine in virtually all excitatory neurons throughout the brain.

It is the most abundant neurotransmitter in the central nervous system; the great majority of excitatory synapses are glutamatergic.

RECEPTORS: both ionotropic and metabotropic.

The ionotropic receptors are the AMPA receptor (fast excitation, admits sodium), the NMDA receptor (the coincidence detector, admits calcium, but only when glutamate is bound AND the membrane is already depolarized enough to expel its magnesium block), and the kainate receptor.

The metabotropic receptors are the eight mGluR subtypes, slower and modulatory.

The NMDA receptor additionally requires a co-agonist, either glycine or D-serine, before it will open, which means it is a triple-coincidence detector: glutamate, depolarization, and co-agonist all at once.

TARGETS AND EFFECTS (SUBSCRIBE): everywhere.

  https://www.linkedin.com/in/d-r-dison

Glutamate is the principal driver of the cortex, the hippocampal trisynaptic circuit, the thalamocortical radiations, and essentially every long-range excitatory pathway in this volume.

FUNCTION AND FACULTY: fast excitatory signalling, and, through the NMDA receptor's calcium influx, the induction of long-term potentiation, the leading molecular mechanism of learning and memory.

Glutamate is both the brain's signal and the trigger for changing the wiring that carries it.

WHEN IT FAILS: too much glutamate is excitotoxic.

The massive glutamate release during a stroke or seizure floods neurons with calcium and kills them, and this excitotoxicity is a final common pathway of neural death across many diseases.

Too little NMDA function, by contrast, produces the dissociative and psychotic states caused by ketamine and PCP, which block the NMDA receptor, and NMDA hypofunction is a leading hypothesis for schizophrenia.

DIGITAL NOTE: glutamate maps naturally onto the excitatory signal-passing edge of a network, with the NMDA receptor's coincidence requirement corresponding to a conjunctive gate that both passes signal and, when its conditions are met, writes a weight change.

The dual role of carrying the signal and triggering the learning is unusual and worth preserving in any model rather than splitting into two separate mechanisms.

SPECIFICATION BLOCK

DYNAMICS: Fast DIRECTED SIGNAL on ionotropic receptors; carries excitation. Modeled as the excitatory edge; NMDA as a coincidence gate that also triggers weight change.

DRIVES: Not Applicable, for the reason that this molecule is the signal itself and serves no drive of its own; it carries signal on behalf of all systems.

PLASTICITY: IS the plasticity trigger (NMDA calcium influx = the coincidence factor).

MODULATION: Not Applicable, for the reason that this molecule IS the fast signal, not a target of modulation; nothing sets its gain because it carries content rather than adjusting it.

GABA (GAMMA-AMINOBUTYRIC ACID)

ORIGIN (PUBLISH): synthesized from glutamate by the enzyme GAD (glutamic acid decarboxylase) in inhibitory interneurons throughout the brain.

It is the principal inhibitory transmitter of the adult brain, the exact functional mirror of glutamate.

RECEPTORS: the ionotropic GABA-A receptor (admits chloride, fast inhibition) and the metabotropic GABA-B receptor (slow inhibition via potassium channels).

  https://www.linkedin.com/in/d-r-dison

The GABA-A receptor is the target of an enormous class of drugs: benzodiazepines, barbiturates, alcohol, and many general anaesthetics all bind it and enhance inhibition.

TARGETS AND EFFECTS (SUBSCRIBE): local inhibition everywhere, sculpting the activity of every excitatory circuit.

The fast-spiking parvalbumin interneurons that generate gamma rhythms are GABAergic, so GABA is also the pacemaker of the brain's fastest coordinating rhythm.

FUNCTION AND FACULTY: inhibition, and therefore the shaping, timing, and containment of excitation.

Without GABA, excitation has no brake.

It is also, through the interneurons, the source of the oscillations that bind and segregate.

WHEN IT FAILS: too little GABA function produces epilepsy (uncontrolled excitation) and anxiety; drugs that enhance it are anticonvulsant and anxiolytic.

A notable developmental fact: in the immature brain, GABA is briefly EXCITATORY, because the chloride gradient is reversed early in development, and this switch is essential for normal brain wiring.

The same molecule, the same receptor, opposite sign, depending on developmental state.

DIGITAL NOTE: GABA maps onto the inhibitory edge, but its role in generating rhythm means it cannot be modelled as mere subtraction.

The parvalbumin-GABA circuit is an oscillator, and a model that treats inhibition only as a negative weight will lose the timing functions that inhibition actually performs.

SPECIFICATION BLOCK

DYNAMICS: Fast DIRECTED SIGNAL, inhibitory; also the OSCILLATOR substrate (parvalbumin interneurons pace gamma). Modeled as inhibitory edge AND rhythm generator.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Inhibitory synapses are plastic; sign can invert in development.

MODULATION: Not Applicable, for the reason that this molecule IS the fast signal, not a target of modulation; nothing sets its gain because it carries content rather than adjusting it.

GLYCINE

ORIGIN (PUBLISH): released by inhibitory interneurons concentrated in the spinal cord and brainstem.

RECEPTORS: the ionotropic glycine receptor (admits chloride, fast inhibition).

Separately, glycine is a required CO-AGONIST at the NMDA receptor, a completely distinct role from its inhibitory one.

TARGETS AND EFFECTS (SUBSCRIBE): fast inhibition in the spinal cord and brainstem, including the reciprocal inhibition that coordinates opposing muscles.

FUNCTION AND FACULTY: spinal and brainstem inhibition, motor coordination, and the reflex circuitry.

Its co-agonist role additionally gates excitatory plasticity brain-wide.

WHEN IT FAILS: blockade of the glycine receptor by strychnine removes inhibition and produces convulsions and extreme muscle rigidity; the poisoning is a picture of inhibition catastrophically lost.

The genetic disorder hyperekplexia, exaggerated startle, is a glycine receptor defect.

DIGITAL NOTE: glycine's two unrelated jobs, inhibitory transmitter and excitatory co-agonist, are a caution for modelling: the same token cannot be assumed to have one meaning.

Its effect is defined by which receptor it meets.

SPECIFICATION BLOCK

DYNAMICS: Fast DIRECTED SIGNAL, inhibitory (spine/brainstem); also NMDA co-agonist. Two roles, receptor-defined.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Co-agonist role gates plasticity elsewhere.

MODULATION: Not Applicable, for the reason that this molecule is itself a modulator or a signal, and so is not the recipient of modulation; nothing sets a gain upon it.

D-SERINE

ORIGIN (PUBLISH): synthesized by astrocytes and some neurons; a striking example of a signalling molecule made largely by glia rather than neurons.

RECEPTORS: the co-agonist site of the NMDA receptor, the same site glycine uses; in the forebrain, D-serine is now thought to be the dominant co-agonist.

FUNCTION AND FACULTY: gating NMDA-receptor-dependent plasticity, and therefore learning.

Because it comes from astrocytes, D-serine is a principal channel by which glia participate directly in synaptic plasticity, a cornerstone of the tripartite synapse.

DIGITAL NOTE: D-serine is the clearest small-molecule evidence that the glia are not passive support.

A model that represents only neurons omits an entire population that gates learning.

SPECIFICATION BLOCK

DYNAMICS: NMDA co-agonist, glia-sourced. Modeled as a plasticity-enabling gate, not an edge.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Gates NMDA plasticity; glial control of learning.

MODULATION: Not Applicable, for the reason that this molecule is itself a modulator or a signal, and so is not the recipient of modulation; nothing sets a gain upon it.

7.4 CLASS TWO: THE MONOAMINES (THE NEUROMODULATORS)

These are the great broadcast systems.

Each arises from a small nucleus, described in Chapter 2, and projects diffusely to vast territories, setting the global state within which the fast transmitters operate.

They do not carry the message; they set how the message is treated.

DOPAMINE

ORIGIN (PUBLISH): two principal sources, and the distinction is functionally decisive.

The substantia nigra pars compacta supplies the nigrostriatal pathway to the dorsal striatum (movement and habit).

The ventral tegmental area supplies the mesolimbic pathway to the nucleus accumbens (reward and motivation) and the mesocortical pathway to the prefrontal cortex (working memory).

A fourth, smaller source, the arcuate nucleus of the hypothalamus, supplies the tuberoinfundibular pathway that inhibits prolactin release, a pure endocrine role.

RECEPTORS: five metabotropic subtypes in two families.

The D1 family (D1, D5) is excitatory in effect; the D2 family (D2, D3, D4) is inhibitory.

All are G-protein-coupled and therefore slow and modulatory.

  https://www.linkedin.com/in/d-r-dison

There is no ionotropic dopamine receptor: dopamine never carries a fast signal, only ever adjusts one.

TARGETS AND EFFECTS (SUBSCRIBE): the striatum, where it biases action selection toward "go"; the nucleus accumbens, where it assigns incentive salience, the property of "wanting"; and the prefrontal cortex, where it tunes the signal-to-noise ratio of working memory along an inverted-U curve, so that both too little and too much impair cognition.

FUNCTION AND FACULTY: reward prediction error, the surprise signal of learning; motivation and incentive salience; motor initiation; and the gain control of working memory.

As Chapter 6 detailed, dopamine speaks in two tenses, phasic bursts carrying prediction error and tonic levels carrying ongoing value, and in at least two populations, a medial value population and a lateral salience population that responds to any important event including aversive ones.

WHEN IT FAILS: too little dopamine in the nigrostriatal pathway is Parkinson's disease, with its bradykinesia, rigidity, and tremor, symptomatic only after 60 to 80 percent of the source neurons have died.

Too much dopamine transmission, or too many D2 receptors, is implicated in the positive symptoms of schizophrenia, which is why every antipsychotic drug blocks D2 receptors.

Every drug of abuse, without exception, raises dopamine in the nucleus accumbens.

Dopamine dysregulation thus sits at the centre of movement disorders, psychosis, and addiction alike.

DIGITAL NOTE: dopamine is the best-characterized bridge between a neurotransmitter and a computational quantity: its phasic signal IS a temporal-difference reward prediction error, the same object used in reinforcement learning.

A digital system could represent phasic dopamine as a scalar error signal broadcast to update value estimates, and tonic dopamine as a slowly varying gain or motivation parameter.

The two-population structure, value versus salience, means a single scalar is insufficient; at least two channels are needed.

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD. State: phasic (reward prediction error) and tonic (state value) level per target. Modeled as a global scalar with per-target levels, NOT an edge.

DRIVES: Carries the reward signal that IS drive-reduction; the common currency of all drives.

PLASTICITY: THE principal third factor: gates three-factor plasticity across striatum, hippocampus, cortex.

MODULATION: It is modulation.

NOREPINEPHRINE (NORADRENALINE)

  https://www.linkedin.com/in/d-r-dison

ORIGIN (PUBLISH): almost entirely from the locus coeruleus, a nucleus of only about fifteen thousand neurons per side whose axons reach essentially the entire brain and spinal cord.

Also released peripherally by the sympathetic nervous system and, as a hormone, by the adrenal medulla.

RECEPTORS: the metabotropic adrenergic receptors, alpha-1, alpha-2, and beta, all G-protein-coupled.

TARGETS AND EFFECTS (SUBSCRIBE): the whole cortex, hippocampus, amygdala, thalamus, and cerebellum.

It raises the gain of target neurons, sharpening their response to strong inputs and suppressing weak ones, which is arousal implemented at the level of single-neuron responsiveness.

FUNCTION AND FACULTY: arousal, vigilance, and the reorienting of attention to salient and unexpected events.

Its phasic burst marks a surprise; its tonic level sets the overall level of alertness, along another inverted-U, so that both drowsiness and extreme stress impair performance.

It is silent during REM sleep.

It also tags emotionally arousing events for stronger memory, via the amygdala.

WHEN IT FAILS: too little produces inattention and the cognitive fog of underarousal; too much, as in chronic stress or panic, produces hypervigilance and anxiety.

Its degeneration in the locus coeruleus is one of the earliest lesions in both Alzheimer's and Parkinson's disease, possibly the earliest detectable of all.

DIGITAL NOTE: norepinephrine maps naturally onto a global gain or temperature parameter that sharpens or softens the whole system's responsiveness, and onto an interrupt signal that reorients processing when something unexpected occurs.

It is the clearest biological analogue of a system-wide "pay attention now" broadcast.

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD. State: tonic (arousal) and phasic (surprise) level. Modeled as a global GAIN/temperature parameter plus a reorienting interrupt.

DRIVES: Serves arousal, the enabling condition of every drive.

PLASTICITY: A third factor: gates and strengthens plasticity, especially emotional memory.

MODULATION: It is modulation.

  https://www.linkedin.com/in/d-r-dison

SEROTONIN (5-HYDROXYTRYPTAMINE, 5-HT)

ORIGIN (PUBLISH): the raphe nuclei of the brainstem midline, projecting brain-wide and to the spinal cord.

Strikingly, roughly ninety percent of the body's serotonin is not in the brain at all but in the gut, where it regulates motility.

RECEPTORS: at least fourteen subtypes across seven families, the most diverse receptor family of any transmitter.

Thirteen are metabotropic; one, the 5-HT3 receptor, is ionotropic, a fast excitatory channel and the sole exception.

This receptor diversity is why serotonin's effects are so various and so hard to summarize.

TARGETS AND EFFECTS (SUBSCRIBE): nearly the entire brain.

Effects depend heavily on which receptor subtype dominates in the target, which is why serotonin resists a single functional label.

FUNCTION AND FACULTY: modulation of mood, but also of aggression, impulse control, appetite, sleep, and the sense of satiety and social rank.

A useful, if simplified, framing is that serotonin promotes patience, behavioural restraint, and tolerance of delay.

It is not simply the "happiness molecule" of popular writing; that framing is misleading.

WHEN IT FAILS: dysregulation is associated with depression and anxiety, which is why the selective serotonin reuptake inhibitors, which raise synaptic serotonin, are first-line antidepressants, though their mechanism is clearly more complex than mere serotonin elevation given their delayed action.

Serotonin is also the system engaged by classic psychedelics, which act at the 5-HT2A receptor.

The 5-HT3 receptor mediates nausea, and blocking it is the basis of the most effective anti-emetic drugs.

DIGITAL NOTE: serotonin's fourteen receptors are a warning against representing a neuromodulator as a single scalar.

Its effect is not one quantity but a family of quantities selected by receptor context.

A model might represent serotonin as a modulator whose sign and magnitude at each target are set by a per-target receptor profile, rather than as one global level.

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD with 14 receptor subtypes. State: level per target, with per-target sign set by receptor profile. NOT a single scalar.

  https://www.linkedin.com/in/d-r-dison

DRIVES: Modulates mood, patience, impulse control; biases delay tolerance across drives.

PLASTICITY: Modulates plasticity and mood set-points.

MODULATION: It is modulation.

HISTAMINE

ORIGIN (PUBLISH): a small population of neurons in the tuberomammillary nucleus of the hypothalamus, only a few thousand cells, projecting widely.

RECEPTORS: four metabotropic subtypes, H1 through H4.

TARGETS AND EFFECTS (SUBSCRIBE): broad cortical and subcortical projections that promote wakefulness and arousal.

FUNCTION AND FACULTY: the maintenance of wakefulness.

Histamine is a core component of the ascending arousal system.

WHEN IT FAILS: antihistamines that cross the blood-brain barrier block H1 receptors and cause profound drowsiness; this is the whole reason first-generation allergy drugs make people sleepy.

A few thousand cells, blocked, and consciousness dims, a vivid demonstration of the reach-not-size principle.

DIGITAL NOTE: histamine is another arousal or wakefulness gate, overlapping with norepinephrine and orexin.

Its inclusion emphasizes that arousal is not a single system but a coalition of redundant broadcast systems, which is itself a design worth noting: the brain does not trust wakefulness to one nucleus.

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD. State: wakefulness-promoting level. A redundant arousal gate.

DRIVES: Serves wakefulness (sleep/wake drive).

PLASTICITY: Not Applicable, for the reason that this construct is a cause of plasticity elsewhere or a fixed signalling element, not itself a site whose weights change with experience.

MODULATION: It is modulation.

7.5 CLASS THREE: ACETYLCHOLINE (IN A CLASS OF ITS OWN)

Acetylcholine deserves separate treatment because it is both a fast, precise transmitter and a diffuse modulator, and it straddles the central and peripheral nervous systems as no other molecule does.

ORIGIN (PUBLISH): two very different sources.

In the periphery, it is the transmitter of the neuromuscular junction (every voluntary movement) and of the autonomic ganglia and all parasympathetic targets.

In the brain, it arises from two nuclear groups: the basal forebrain complex (the nucleus basalis of Meynert and the medial septum), which supplies the cortex and hippocampus, and the brainstem pedunculopontine and laterodorsal tegmental nuclei, which supply the thalamus.

RECEPTORS: two entirely different families.

The nicotinic receptors are ionotropic, fast, and are the receptors at the neuromuscular junction and the target of nicotine.

The muscarinic receptors (M1 through M5) are metabotropic, slow, and mediate most of acetylcholine's modulatory effects in the brain and its parasympathetic effects in the body.

TARGETS AND EFFECTS (SUBSCRIBE): at the neuromuscular junction, it commands muscle contraction, reliably and without fail.

In the cortex, released from the basal forebrain, it acts as a diffuse modulator that raises attention, enhances the encoding of new information, and increases the signal-to-noise ratio of sensory processing.

FUNCTION AND FACULTY: peripherally, movement and parasympathetic "rest and digest" control.

Centrally, attention, arousal, and the gating of learning: acetylcholine essentially signals to the cortex "this is worth encoding." It also drives REM sleep and the hippocampal theta rhythm.

WHEN IT FAILS: degeneration of the basal forebrain cholinergic neurons is an early and severe feature of Alzheimer's disease, and the loss of cortical acetylcholine tracks the memory decline; the cholinesterase-inhibitor drugs, which raise acetylcholine, are a first-line, if modest, treatment.

At the neuromuscular junction, autoimmune attack on nicotinic receptors is myasthenia gravis, and blockade of acetylcholine release by botulinum toxin is botulism, and, in dilute form, Botox.

Nerve agents and organophosphate insecticides kill by blocking the enzyme that breaks acetylcholine down, causing catastrophic, unending stimulation.

DIGITAL NOTE: acetylcholine's split personality, precise commander in the periphery and diffuse learning-gate in the cortex, is a caution against assuming a molecule has one computational role.

In a model, the cortical cholinergic signal maps onto a learning-rate or attention gate that says "encode this now," which is distinct from and complementary to dopamine's "this was better or worse than expected."

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD. State: wakefulness-promoting level. A redundant arousal gate.

DRIVES: Serves wakefulness (sleep/wake drive).

  https://www.linkedin.com/in/d-r-dison

PLASTICITY: Not Applicable, for the reason that this construct is a cause of plasticity elsewhere or a fixed signalling element, not itself a site whose weights change with experience.

MODULATION: It is modulation.

7.6 CLASS FOUR: THE NEUROPEPTIDES (THE SLOW, SPECIFIC MODULATORS)

There are dozens of neuropeptides, short chains of amino acids, and they form the largest and most diverse class of signalling molecules in the brain.

They are co-released with the classical transmitters, but from separate dense-core vesicles that require sustained high-frequency firing to release, so peptides are the signal the brain sends only when it means it.

They act almost exclusively through metabotropic G-protein-coupled receptors, so they are slow and long-lasting, and they are typically not cleared by reuptake but by degradation, so their effects persist.

If the monoamines are the volume knobs, the neuropeptides are the slow mood and drive settings.

Only the principal ones can be catalogued here; the class is very large.

THE OPIOID PEPTIDES (ENDORPHINS, ENKEPHALINS, DYNORPHINS)

ORIGIN (PUBLISH): widely distributed, including the periaqueductal gray, hypothalamus, and limbic regions.

RECEPTORS: the mu, delta, and kappa opioid receptors, all metabotropic.

FUNCTION AND FACULTY: the suppression of pain (the mu receptor in the periaqueductal gray is the target of morphine and of the descending analgesia system), and the mediation of reward, pleasure, and social bonding.

The endogenous opioids are why a wound may not be felt in the heat of a crisis, and part of why social connection feels good.

WHEN IT FAILS: the exogenous opioids that mimic them, from morphine to fentanyl, are supremely effective analgesics and supremely addictive, and their action on brainstem respiratory centres is what makes overdose lethal.

DIGITAL NOTE: an internal reward-and-relief signal distinct from dopamine's wanting; the opioid system corresponds more to "liking" and to relief from aversive states than to the pursuit that dopamine drives.

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD (peptide). State: level in pain and reward territories. Slow, long-lasting.

DRIVES: Serves pain relief and social bonding; the 'liking'/relief signal (vs dopamine's wanting).

PLASTICITY: Modulates reward learning.

  https://www.linkedin.com/in/d-r-dison

MODULATION: It is modulation.

OXYTOCIN AND VASOPRESSIN

ORIGIN (PUBLISH): the paraventricular and supraoptic nuclei of the hypothalamus, released both into the bloodstream (from the posterior pituitary) and within the brain (from dendrites), and the two releases are independently regulated.

RECEPTORS: metabotropic oxytocin and vasopressin receptors.

FUNCTION AND FACULTY: oxytocin in social bonding, maternal behaviour, trust, and the salience of social cues; vasopressin in social behaviour, affiliation, aggression, and, as a hormone, water retention.

The popular framing of oxytocin as a simple "love hormone" is not supported; it appears rather to amplify the salience of social information, whether positive or negative.

DIGITAL NOTE: a social-salience modulator that up-weights social inputs; not a valence signal but a relevance signal specific to the social domain.

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD (peptide), also blood-borne hormone. State: level in social circuits.

DRIVES: Serves the SOCIAL CONNECTION drive; up-weights social salience.

PLASTICITY: Modulates social-memory plasticity.

MODULATION: It is modulation.

CRH (CORTICOTROPIN-RELEASING HORMONE)

ORIGIN (PUBLISH): the paraventricular nucleus of the hypothalamus, and the central amygdala and BNST.

FUNCTION AND FACULTY: the apex trigger of the HPA stress axis (Chapter 2), and, acting within the brain, a direct driver of anxiety and stress behaviour.

CRH is, functionally, the brain's "there is a threat, mobilize" signal.

WHEN IT FAILS: chronic CRH overactivity is implicated in depression and anxiety disorders.

DIGITAL NOTE: a global stress or threat-mobilization parameter that reconfigures priorities toward defense, distinct from moment-to-moment fear.

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD (peptide) and endocrine trigger. State: stress-mobilization level.

DRIVES: The 'threat, mobilize' signal; serves the safety/stress drive and triggers the HPA axis.

  https://www.linkedin.com/in/d-r-dison

PLASTICITY: Chronic elevation shifts stress set-points (allostasis).

MODULATION: It is modulation.

OREXIN (HYPOCRETIN)

ORIGIN (PUBLISH): a small population in the lateral hypothalamus.

FUNCTION AND FACULTY: the stabilizer of wakefulness.

Orexin holds the sleep-wake flip-flop switch (Chapter 5) in the waking state and prevents unwanted transitions.

WHEN IT FAILS: loss of orexin neurons causes narcolepsy with cataplexy: the switch flips uncontrollably, dropping the person into sleep or into REM muscle atonia without warning.

DIGITAL NOTE: a state-stabilization signal that adds hysteresis to a bistable state switch, preventing rapid oscillation between modes.

SPECIFICATION BLOCK

DYNAMICS: BROADCAST FIELD (peptide). State: wake-stabilization level. Adds hysteresis to the sleep-wake GATE.

DRIVES: Stabilizes the sleep/wake drive switch.

PLASTICITY: Not Applicable, for the reason that this construct is a cause of plasticity elsewhere or a fixed signalling element, not itself a site whose weights change with experience.

MODULATION: It is modulation.

THE OTHER PRINCIPAL NEUROPEPTIDES, IN BRIEF

Substance P: transmits pain signals and mediates stress responses; acts at the NK1 receptor.

Neuropeptide Y (NPY): a powerful driver of feeding and an endogenous anxiolytic; it opposes CRH, promoting calm and energy conservation.

Cholecystokinin (CCK): signals satiety after a meal, and, in the brain, can be anxiogenic.

Galanin: promotes sleep and modulates mood; implicated in stress resilience.

Somatostatin: broadly inhibitory; also a hypothalamic hormone that restrains growth hormone.

VIP (vasoactive intestinal peptide): a transmitter of the suprachiasmatic clock and of disinhibitory cortical interneurons.

The general lesson of the class: neuropeptides provide a large vocabulary of slow, specific, long-lasting drive and mood signals, each with its own receptor and its own behavioural domain.

They are how the brain sets not what it is doing this millisecond but what it is disposed to do over the next minutes and hours.

  https://www.linkedin.com/in/d-r-dison

7.7 CLASS FIVE: THE GASOTRANSMITTERS (THE RULE-BREAKERS)

Three gases, nitric oxide, carbon monoxide, and hydrogen sulfide, act as messengers while violating nearly every rule of classical neurotransmission.

They are not stored in vesicles; they are synthesized on demand.

They are not released by exocytosis; they diffuse freely through membranes in all directions.

They bind no surface receptor; they act on intracellular targets, typically by binding metal centres in enzymes such as guanylyl cyclase.

And they cannot be cleared by reuptake; they simply decay.

NITRIC OXIDE (NO) is the best characterized.

Synthesized on demand by nitric oxide synthase when calcium enters a neuron, it diffuses backward across the synapse from the postsynaptic to the presynaptic cell, making it a RETROGRADE messenger: the receiving cell talks back to the sending cell.

This retrograde signalling is a mechanism of long-term potentiation, allowing a strengthened synapse to signal its own strengthening back to its input.

Nitric oxide also controls local blood flow, coupling neural activity to the blood supply that functional imaging measures.

CARBON MONOXIDE and HYDROGEN SULFIDE act on similar principles and are implicated in plasticity, blood-flow control, and, increasingly, pain signalling, though their roles are less fully mapped.

FUNCTION AND FACULTY: synaptic plasticity through retrograde signalling, and neurovascular coupling.

DIGITAL NOTE: the gasotransmitters are the clearest biological case of a signal that is not addressed to anyone.

They diffuse to whatever is nearby, in all directions, including backward.

In a model, they correspond not to an edge but to a local field: a quantity that spreads from its source to affect a neighbourhood, including the very units that produced it.

This is a fundamentally different primitive from a directed connection, and a model with only directed edges cannot represent it.

SPECIFICATION BLOCK

DYNAMICS: LOCAL RETROGRADE FIELD. State: local concentration. Diffuses in all directions including backward to its source; NOT an edge. A distinct primitive.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or

  https://www.linkedin.com/in/d-r-dison

motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Retrograde signal that helps consolidate LTP (a plasticity messenger).

MODULATION: Not Applicable, for the reason that this molecule is itself a modulator or a signal, and so is not the recipient of modulation; nothing sets a gain upon it.

7.8 CLASS SIX: THE ENDOCANNABINOIDS (THE DIMMER SWITCH)

The endocannabinoids, principally anandamide and 2-AG, are lipid messengers synthesized on demand from the cell membrane.

Like the gases, they are retrograde: they are made by the postsynaptic cell and travel backward to act on CB1 receptors on the presynaptic terminal, where they SUPPRESS further neurotransmitter release.

FUNCTION AND FACULTY: this makes the endocannabinoid system a demand-driven brake.

When a postsynaptic neuron is over-excited, it releases endocannabinoids that turn down its own inputs, a negative feedback that protects against runaway excitation and shapes plasticity.

The system regulates appetite, pain, mood, and memory.

It is the system engaged by cannabis, whose THC mimics the endocannabinoids at the CB1 receptor.

DIGITAL NOTE: a retrograde, demand-driven gain-reduction signal, mathematically a local negative-feedback controller that a downstream unit uses to throttle its own upstream inputs.

Again a retrograde primitive, and again not representable by a forward directed edge alone.

SPECIFICATION BLOCK

DYNAMICS: LOCAL RETROGRADE FIELD. State: on-demand level. Made post-synaptically, acts back on the pre-synaptic terminal to SUPPRESS release. A demand-driven brake.

DRIVES: Not Applicable, for the reason that this construct serves no homeostatic drive directly; it contributes to perception, cognition, or motor output on which the drives act, rather than monitoring a bodily variable of its own.

PLASTICITY: Retrograde gain-reduction; shapes plasticity and prevents runaway excitation.

MODULATION: Not Applicable, for the reason that this molecule is itself a modulator or a signal, and so is not the recipient of modulation; nothing sets a gain upon it.

7.9 CLASS SEVEN: THE PURINES (THE ENERGY MESSENGERS)

  https://www.linkedin.com/in/d-r-dison

ATP, the universal energy currency, doubles as an extracellular signalling molecule, as does its breakdown product adenosine.

ATP is co-released with classical transmitters and acts on purinergic (P2) receptors, both ionotropic and metabotropic.

It is a fast co-transmitter and a major signal between neurons and glia.

ADENOSINE, produced as ATP is consumed, acts on its own receptors and accumulates in the brain during waking, producing the mounting SLEEP PRESSURE that eventually forces sleep.

Caffeine works precisely by blocking adenosine receptors: it does not add energy; it masks the accumulating signal of fatigue.

FUNCTION AND FACULTY: fast co-transmission and neuron-glia signalling (ATP); the homeostatic tracking of time-awake and the generation of sleep pressure (adenosine).

DIGITAL NOTE: adenosine is a beautifully clean example of a homeostatic integrator: a quantity that accumulates monotonically with activity or time-awake and, on crossing a threshold, forces a state transition to sleep, then is cleared during that state.

It maps onto an accumulator or leaky integrator gating a mandatory maintenance cycle.

SPECIFICATION BLOCK

DYNAMICS: ATP = fast co-signal; ADENOSINE = a homeostatic INTEGRATOR (accumulates with time-awake). State: adenosine level. Gates the sleep switch.

DRIVES: Adenosine IS the sleep-pressure drive variable.

PLASTICITY: Not Applicable, for the reason that this construct is a cause of plasticity elsewhere or a fixed signalling element, not itself a site whose weights change with experience.

MODULATION: Caffeine blocks adenosine (masks the signal).

7.10 CLASS EIGHT: THE TROPHIC AND IMMUNE SIGNALS (THE SLOW BUILDERS)

These do not carry moment-to-moment signals; they build, maintain, and remodel the machinery over hours to days.

BDNF (BRAIN-DERIVED NEUROTROPHIC FACTOR) is the principal member.

Released by active neurons and by glia, acting on the TrkB receptor, it promotes neuronal survival, the growth of new synapses, and the strengthening of existing ones.

BDNF is the molecular link between activity and structural plasticity: it is largely how "neurons that fire together wire together" becomes physical.

It is elevated by exercise, learning, and antidepressant treatment, and suppressed by chronic stress, which is a direct molecular pathway from a Layer 12 experience to a Layer 4 change in the machinery of the brain.

THE CYTOKINES, the signalling proteins of the immune system, also act on the brain.

Circulating cytokines from an infection produce sickness behaviour, the lethargy, withdrawal, and low mood that accompany illness, and chronic inflammation is increasingly implicated in depression.

The microglia, the brain's immune cells, both respond to and release these signals, and they physically prune synapses.

FUNCTION AND FACULTY: BDNF, structural plasticity, survival, and the physical substrate of long-term learning.

Cytokines, the neuroimmune interface, sickness behaviour, and inflammation-linked mood change.

DIGITAL NOTE: BDNF corresponds to a slow structural-plasticity signal that gates whether an active connection is consolidated into permanent structure, distinct from the fast weight changes of ordinary learning.

It is the difference between writing to cache and committing to disk.

The cytokines correspond to a whole-system state signal, originating outside the cognitive core, that can reconfigure global mood and energy, a reminder that in a biological system the "hardware" can message the "software."

SPECIFICATION BLOCK

DYNAMICS: ATP = fast co-signal; ADENOSINE = a homeostatic INTEGRATOR (accumulates with time-awake). State: adenosine level. Gates the sleep switch.

DRIVES: Adenosine IS the sleep-pressure drive variable.

PLASTICITY: Not Applicable, for the reason that this construct is a cause of plasticity elsewhere or a fixed signalling element, not itself a site whose weights change with experience.

MODULATION: Caffeine blocks adenosine (masks the signal).

7.11 THE ENDOCRINE HORMONES THAT ACT ON THE BRAIN

Finally, the bloodstream-borne hormones reach the brain and reconfigure it over the longest timescales of all, from minutes to a lifetime.

They are the slowest and broadest messengers, and several are the very molecules that make the layer-skipping causation of the companion framework possible.

CORTISOL, the glucocorticoid stress hormone from the adrenal cortex, is small and fat-soluble, so it crosses every membrane, enters neurons, and binds intracellular receptors that act directly on DNA, altering gene expression.

It sharpens memory acutely but, in chronic excess, damages the hippocampus.

Cortisol is the exemplar of a signal that reaches from a social event all the way down to the gene.

  https://www.linkedin.com/in/d-r-dison

THE SEX HORMONES, testosterone, estrogen, and progesterone, exert two kinds of effect: organizational effects that permanently shape the developing brain, and activational effects that modulate mood, cognition, and behaviour in the adult.

Estrogen is notably neuroprotective and interacts with the serotonin and BDNF systems.

THYROID HORMONE sets the metabolic rate of every neuron and is absolutely required for normal brain development; its deficiency in early life causes irreversible intellectual disability.

INSULIN, LEPTIN, and GHRELIN report the body's energy state to the hypothalamus, and increasingly appear to affect cognition and mood, linking metabolism to mind.

MELATONIN, from the pineal gland, is the hormonal signal of darkness that entrains the circadian system.

DIGITAL NOTE: the hormones correspond to the slowest and most global parameters of the system, the ones that reconfigure the operating regime over long epochs and, in the case of the developmental hormones, set structural parameters once and for life.

In a model they are less like messages and more like slowly varying configuration state, some of it set during a developmental phase and thereafter largely fixed.

SPECIFICATION BLOCK

DYNAMICS: SLOWEST BROADCAST / CONFIGURATION STATE. State: blood level, reconfiguring the operating regime over minutes to a lifetime. Cortisol crosses to DNA (layer-skipping).

DRIVES: Report and defend the slowest bodily set-points (stress, energy, reproduction, circadian).

PLASTICITY: Reconfigure plasticity globally; developmental hormones set structure once.

MODULATION: They are the slowest modulation.

7.12 SYNTHESIS: THE MESSENGERS AS A SYSTEM

Stepping back from the catalogue, a small number of organizing principles govern the whole molecular layer, and these principles, more than any individual molecule, are what a model must capture.

FIRST, SPEED IS STRATIFIED BY RECEPTOR TYPE.

Ionotropic receptors are fast, in milliseconds, and carry signal.

Metabotropic receptors are slow, in hundreds of milliseconds to seconds, and carry modulation.

Hormonal and genomic actions are slower still, minutes to a lifetime.

  https://www.linkedin.com/in/d-r-dison

A messenger's speed, and therefore its role, is set by the receptor it meets, not by the molecule alone.

Glutamate on an AMPA receptor is a signal; glutamate on an mGluR is a modulation.

SECOND, THE SAME MOLECULE MEANS DIFFERENT THINGS IN DIFFERENT PLACES.

Dopamine in the striatum selects actions; in the prefrontal cortex it tunes working memory.

Acetylcholine at the muscle commands contraction; in the cortex it gates learning.

GABA in the adult inhibits; in the infant it excites.

The meaning is not in the molecule; it is in the interface, the pairing of a molecule with a receptor in a region.

This is the molecular restatement of the volume's central thesis: significance is a fact about connections, not about parts.

THIRD, THERE ARE FEW SIGNALS AND MANY MODULATORS.

The fast point-to-point traffic is carried by essentially two molecules, glutamate and GABA, excitation and inhibition.

Everything else, the dozens of monoamines, peptides, gases, and hormones, is modulation: systems that set the gain, the mood, the drive, the arousal, and the plasticity within which those two signals operate.

The brain is a simple fast signalling system wrapped in an extraordinarily rich slow modulatory one.

FOURTH, SOME MESSENGERS ARE NOT ADDRESSED AT ALL.

The gases and endocannabinoids diffuse to whoever is near, including backward to their own source.

The monoamines and hormones broadcast to vast territories.

These are not directed connections; they are fields and broadcasts, and they require a modelling primitive fundamentally different from the directed edge.

A model built only of point-to-point connections can represent glutamate and GABA and nothing else.

FIFTH, THE SLOW MESSENGERS REACH ACROSS THE LAYERS.

Cortisol carries a social event to the gene.

BDNF carries an experience to a structural change.

Cytokines carry a bodily infection to a mood.

These are the molecular vehicles of the layer-skipping causation that the companion framework was built to express, and it is no accident that most

of them are hormones and trophic factors, the messengers that travel by blood and act on the machinery rather than on the signal.

The molecules that skip layers are, almost without exception, the slow, broadcast, membrane-crossing ones.

For the eventual digital system, the catalogue suggests that a faithful model needs at least four distinct messenger primitives, not one: a fast directed signal (glutamate and GABA on ionotropic receptors), a slow directed modulation (metabotropic transmission), a diffuse broadcast (the monoamines and peptides), and a local retrograde field (the gases and endocannabinoids), with the hormones as a fifth, slowest tier of global configuration state.

Any model with a single kind of connection is modelling only the first, and the first is only two molecules out of dozens.

The richness of cognition, this chapter suggests, lives substantially in the modulatory classes that a naive network omits.

7.13 THE REMAINING MESSENGERS, WITH PUBLISH AND SUBSCRIBE

The classes below were named in the class introductions above; here each principal remaining messenger is given with its publisher, its subscribers, and its faculty, so the catalogue is complete.

THE FURTHER NEUROPEPTIDES

NEUROPEPTIDE Y (NPY)

ORIGIN (PUBLISH): hypothalamic arcuate neurons, and interneurons of the cortex, hippocampus, and amygdala.

TARGETS AND EFFECTS (SUBSCRIBE): the Y-family receptors (Y1, Y2, Y5) in the hypothalamus, amygdala, and hippocampus.

FUNCTION AND FACULTY: the most abundant peptide in the brain; drives feeding, and publishes a powerful anxiolytic and stress-resilience signal.

SUBSTANCE P

ORIGIN (PUBLISH): dorsal-horn, limbic, and hypothalamic neurons.

TARGETS AND EFFECTS (SUBSCRIBE): the tachykinin NK1 receptor in the spinal cord, amygdala, and hypothalamus.

FUNCTION AND FACULTY: pain transmission, stress, and emotional response.

CHOLECYSTOKININ (CCK)

ORIGIN (PUBLISH): cortical, hippocampal, and amygdalar interneurons, and the gut.

TARGETS AND EFFECTS (SUBSCRIBE): the CCK2 receptor throughout the brain.

FUNCTION AND FACULTY: satiety, anxiety and panic, memory, and the pacing of network oscillations.

SOMATOSTATIN (SST)

  https://www.linkedin.com/in/d-r-dison

ORIGIN (PUBLISH): cortical and hippocampal inhibitory interneurons, and the hypothalamus.

TARGETS AND EFFECTS (SUBSCRIBE): the somatostatin receptors (sst1 to sst5) in cortex and hippocampus.

FUNCTION AND FACULTY: inhibitory dendritic control, memory modulation, and the restraint of growth-hormone release.

VASOACTIVE INTESTINAL PEPTIDE (VIP)

ORIGIN (PUBLISH): suprachiasmatic-nucleus neurons and cortical VIP interneurons.

TARGETS AND EFFECTS (SUBSCRIBE): the VPAC receptors in the suprachiasmatic nucleus and cortex.

FUNCTION AND FACULTY: circadian synchronization of the master clock, and cortical disinhibition that opens windows for plasticity.

GALANIN

ORIGIN (PUBLISH): the locus coeruleus, hypothalamus, sleep-promoting preoptic neurons, and basal forebrain.

TARGETS AND EFFECTS (SUBSCRIBE): the galanin receptors (GALR1 to GALR3).

FUNCTION AND FACULTY: the co-transmitter of restraint: it dampens arousal, promotes sleep, and modulates mood and memory.

NEUROTENSIN

ORIGIN (PUBLISH): the hypothalamus, central amygdala, and neurons associated with the ventral tegmental area.

TARGETS AND EFFECTS (SUBSCRIBE): the neurotensin receptors (NTS1, NTS2).

FUNCTION AND FACULTY: modulation of dopamine and reward, analgesia, and thermoregulation.

THE GASOTRANSMITTERS (MADE ON DEMAND, NOT STORED)

NITRIC OXIDE (NO)

ORIGIN (PUBLISH): postsynaptic and interneuronal neurons bearing nitric-oxide synthase (made on demand, not from a vesicle).

TARGETS AND EFFECTS (SUBSCRIBE): soluble guanylate cyclase in nearby cells, including the presynaptic terminal it diffuses back to.

FUNCTION AND FACULTY: a retrograde messenger of long-term potentiation and synaptic plasticity, and a vasodilator coupling activity to blood flow.

CARBON MONOXIDE (CO)

ORIGIN (PUBLISH): neurons expressing heme oxygenase.

  https://www.linkedin.com/in/d-r-dison

TARGETS AND EFFECTS (SUBSCRIBE): soluble guanylate cyclase (the same second-messenger pathway as nitric oxide).

FUNCTION AND FACULTY: neuromodulation and the maintenance of plasticity.

HYDROGEN SULFIDE (H_2S)

ORIGIN (PUBLISH): neurons and glia expressing the synthesizing enzymes.

TARGETS AND EFFECTS (SUBSCRIBE): the N-methyl-D-aspartate receptor and transmitter-release machinery.

FUNCTION AND FACULTY: facilitation of long-term potentiation and neuroprotection.

THE ENDOCANNABINOIDS (RETROGRADE LIPID SIGNALS)

ANANDAMIDE AND 2-ARACHIDONOYLGLYCEROL (2-AG)

ORIGIN (PUBLISH): postsynaptic neurons, made on demand from membrane lipids and released backward across the synapse.

TARGETS AND EFFECTS (SUBSCRIBE): the presynaptic cannabinoid receptor (CB1), the most abundant metabotropic receptor in the brain.

FUNCTION AND FACULTY: retrograde suppression of transmitter release, the basis of a synapse quieting its own inputs; mood, memory, appetite, and pain.

THE PURINES

ADENOSINE TRIPHOSPHATE (ATP) AND ADENOSINE

ORIGIN (PUBLISH): co-released with the classical transmitters and from astrocytes (adenosine triphosphate); adenosine forms as it breaks down.

TARGETS AND EFFECTS (SUBSCRIBE): the purinergic P2X and P2Y receptors (adenosine triphosphate) and the P1 receptors (adenosine).

FUNCTION AND FACULTY: fast purinergic co-transmission; adenosine is the accumulating signal of sleep pressure and a brake on arousal (the target of caffeine).

THE ENDOCRINE HORMONES THAT ACT ON COGNITION

CORTISOL (THE PRINCIPAL GLUCOCORTICOID)

ORIGIN (PUBLISH): the adrenal cortex, driven by the hypothalamic-pituitary-adrenal axis.

TARGETS AND EFFECTS (SUBSCRIBE): the mineralocorticoid and glucocorticoid receptors in the hippocampus, amygdala, and prefrontal cortex.

FUNCTION AND FACULTY: the slow arm of the stress response; it modulates memory along an inverted-U and provides negative feedback to its own axis.

LEPTIN AND GHRELIN

ORIGIN (PUBLISH): leptin from fat, ghrelin from the stomach (published into the bloodstream).

TARGETS AND EFFECTS (SUBSCRIBE): their receptors in the hypothalamic arcuate nucleus and the hippocampus.

FUNCTION AND FACULTY: the body's energy report to the brain; both also modulate hippocampal plasticity and memory (leptin enhancing, ghrelin enhancing).

INSULIN

ORIGIN (PUBLISH): the pancreas (with some central synthesis).

TARGETS AND EFFECTS (SUBSCRIBE): the insulin receptor in the hippocampus, cortex, and hypothalamus.

FUNCTION AND FACULTY: central metabolism, synaptic plasticity, and memory; its central resistance is implicated in cognitive decline.

THYROID HORMONE (T3 AND T4)

ORIGIN (PUBLISH): the thyroid gland, driven by the hypothalamic-pituitary-thyroid axis.

TARGETS AND EFFECTS (SUBSCRIBE): nuclear thyroid receptors throughout the brain.

FUNCTION AND FACULTY: brain development and the maintenance of metabolism and cognition.

MELATONIN

ORIGIN (PUBLISH): the pineal gland, driven by the suprachiasmatic clock.

TARGETS AND EFFECTS (SUBSCRIBE): the melatonin receptors (MT1, MT2) in the clock, thalamus, and hippocampus.

FUNCTION AND FACULTY: the hormonal signal of darkness, timing sleep and coupling the circadian clock to cognition.

CHAPTER 8. THE DYNAMICS: HOW A REGION CHANGES IN TIME

8.1 The Gap This Chapter Closes

Everything in this volume up to now has been a map.

It says what connects to what, what each region does, which molecules modulate it, and on what rhythm they talk.

A map is necessary, and it is not sufficient, because a map is frozen.

It shows the roads but not the traffic, the parts but not their motion.

To turn this volume from a description into something an engine can run, one thing must be added to every construct: a rule for how its state changes from one moment to the next.

  https://www.linkedin.com/in/d-r-dison

In the language of the coming specification, every region, interface, and molecule needs a STATE and an UPDATE RULE, a function of the form

```
next_state = update(current_state, inputs, neuromodulators)
```

This is the single most important addition in the final three chapters, and it is what makes the sideways turn into a runnable specification possible at all.

Without it, an implementer can read what the hippocampus IS and still have no idea how to step it forward one tick.

This chapter establishes the vocabulary of dynamics.

The two chapters after it add the drives that give the system goals and the plasticity that lets it learn.

Then Chapter 11 attaches all three, as a compact per-construct block, to the constructs themselves, because a specification wants the rule written ON the component, not gathered in an essay elsewhere.

A NOTE ON HONESTY, STATED ONCE AND MEANT THROUGHOUT

The anatomy in this book is established neuroscience.

The dynamics, drives, and plasticity rules that follow are, in part, MODELING COMMITMENTS: choices this framework makes in order to be buildable, drawn from computational neuroscience but selected and simplified by us.

Where a statement is settled science, it is stated plainly.

Where it is a design decision, it is marked as one, with the phrase THIS FRAMEWORK'S CHOICE or equivalent.

A specification is allowed to commit; it is not allowed to pretend its commitments are facts.

This is the same discipline the companion data standard applied to itself, and it is not optional here.

8.2 The Grain of the Model

A first decision governs everything: at what grain is a region's state represented?

This framework does NOT model individual neurons, individual spikes, or individual synapses.

That was decided earlier and it holds.

Instead it adopts the grain of the NEURAL MASS or FIRING-RATE model: each region is represented by a small number of aggregate state variables describing the coarse activity of the whole population, not the firing of its cells.

This is the standard grain of large-scale brain modeling, it is the grain at which the region-level functions in this book were actually learned, and it is the grain at which a running simulation is tractable.

  https://www.linkedin.com/in/d-r-dison

Concretely, a region's state is a handful of numbers: its overall ACTIVATION (how active the population is now), and where relevant a small set of additional variables such as its current TUNING (what it is currently representing) or its ADAPTATION (how fatigued or sensitized it is).

The update rule says how those numbers change given the region's inputs and the ambient neuromodulatory levels.

8.3 The Six Dynamical Archetypes

The central proposal of this chapter, and THIS FRAMEWORK'S CHOICE, is that the roughly one hundred regions of this volume do not each need a unique, hand-built dynamical rule.

They fall into a small number of DYNAMICAL ARCHETYPES, each a recognized pattern in computational neuroscience, and each region is assigned to one archetype whose parameters it then specializes.

Six archetypes cover the great majority of the brain.

This is the dynamical restatement of the volume's oldest theme: one canonical computation, replicated and rewired.

The cerebellar transform argument, generalized to the whole brain.

ARCHETYPE 1: THE RELAY. State: an activation that tracks its input with a gain and a delay. Update rule, in words: the region's output is its input, scaled by a gain that neuromodulators set, after a fixed delay. A relay does not transform content; it passes it, faithfully or with adjusted strength. This is the dynamics of the transmissive interfaces (Chapter 3), of the first-order thalamic relay nuclei, and of the long white-matter tracts. It is the pipe, made temporal.

ARCHETYPE 2: THE INTEGRATOR. State: an activation that ACCUMULATES its input over time and decays slowly when input stops. Update rule, in words: the new activation is the old activation, slightly decayed, plus the current input. An integrator remembers its recent history and sums evidence. This is the dynamics of evidence accumulation in decision-making, of the buildup of sleep pressure by adenosine, and of any region that must hold a running total. It is a leaky bucket.

ARCHETYPE 3: THE OSCILLATOR. State: an activation plus a phase, which cycles. Update rule, in words: the phase advances at a characteristic frequency, and the region's receptivity and output vary with the phase. An oscillator generates rhythm and gates communication by phase, the mechanism of Chapter 5. This is the dynamics of the thalamic reticular nucleus generating spindles, of the medial septum pacing hippocampal theta, and of the parvalbumin-interneuron networks generating gamma. It is a metronome that also gates.

ARCHETYPE 4: THE ATTRACTOR, or MEMORY. State: a pattern of activation that, once set, holds itself stable against perturbation, and that can be completed from a fragment. Update rule, in words: recurrent self-excitation pulls the current activation toward the nearest of a set of stored stable patterns. This is the dynamics of working memory (a point attractor holding a value), of the head-direction and place systems (a ring or grid attractor

  https://www.linkedin.com/in/d-r-dison

holding a position), and of CA3 completing a memory from a cue. It is a landscape of basins, and activity rolls into the nearest one.

ARCHETYPE 5: THE GATE, or SWITCH. State: a discrete or near-discrete mode, held stable, that flips when driven hard enough. Update rule, in words: the region stays in its current mode until a threshold is crossed, then switches, and resists switching back (hysteresis). This is the dynamics of the sleep-wake flip-flop, of the basal-ganglia action-selection gate that is inhibitory-by-default and releases one channel, and of attentional gating. It is a switch with a deliberate stiffness.

ARCHETYPE 6: THE COMPARATOR, or ERROR UNIT. State: an activation representing the DIFFERENCE between two of its inputs, typically a prediction and an outcome. Update rule, in words: the output is a function of the mismatch between an expected input and an actual one. This is the dynamics of the inferior olive computing motor error, of the dopamine system computing reward prediction error, and of the habenula computing disappointment. It is a subtractor, and its output is surprise.

Most regions are one archetype.

Some are a small combination, a relay whose gain is set by an oscillator, or an integrator feeding a comparator.

The archetype is not the whole truth about a region; it is the truth a runnable model needs first.

8.4 The Universal State Variables

Beneath the archetypes, every region shares a common minimal state, so that the simulation can step uniformly.

THIS FRAMEWORK'S CHOICE is that every region carries at least:

ACTIVATION: a number, roughly the population firing rate, normalized.

How active is this region right now.

GAIN: a multiplier on its responsiveness, set largely by neuromodulators.

How strongly is it reacting to input right now.

This is where norepinephrine and acetylcholine act.

ADAPTATION: a slow variable that rises with sustained activity and reduces responsiveness, capturing fatigue, satiety, and habituation.

Why a constant stimulus stops being noticed.

And where the archetype requires it, a TUNING or CONTENT variable (what the region currently represents) and a PHASE variable (for oscillators).

The update rule for the whole brain, at this grain, is then the parallel update of every region's activation given its inputs through the Conduits of Chapter 3, its gain set by the neuromodulatory broadcast of Chapter 7, and its archetype-specific rule.

One tick of the simulation is one round of this update across all regions at once.

  https://www.linkedin.com/in/d-r-dison

This is the temporal engine the specification will formalize.

8.5 The Interfaces Are Dynamical Too

The Conduits of Chapter 3 are not instantaneous wires.

Each carries its signal with a DELAY (set by its length and myelination, from Chapter 6's conduction velocities) and a WEIGHT (its current strength, which plasticity will change).

A transmissive Conduit applies delay and weight and nothing else.

A computational Conduit additionally applies its transform, the CRO it performs, which for the framework means it runs its own small update rule on the signal in transit.

The thalamic reticular nucleus, as an interface, has the GATE archetype; the inferior olive, as an interface, has the COMPARATOR archetype.

This is what it means, in dynamical terms, for an interface to compute rather than merely convey.

CHAPTER 9. THE DRIVES: WHAT THE SYSTEM WANTS

9.1 Why a Mind Needs Wants

A system that only maps inputs to outputs is a calculator.

A mind is not a calculator, because a mind has GOALS it is trying to maintain, and it evaluates every input and every action by whether it serves those goals.

The difference between a reflex and a motivation is that a motivation has a target state it is defending.

This chapter gives the system its wants.

It is short, because the idea is simple and powerful, but it is not optional: a specification of a mind that omits what the mind wants has specified the wrong kind of object.

Chapter 12's argument that there is no orchestrator already depends on this chapter, because the override hierarchy it describes (breathing beats deliberation) is precisely a ranking of DRIVES.

9.2 The Thermostat, and Then the Generalization

The simple form, the on-ramp, is the thermostat.

A drive has four parts:

A MONITORED VARIABLE: the quantity it watches (body temperature, blood glucose, blood oxygen, hydration, safety, social connection).

A SET-POINT: the value it defends (37 degrees, and so on).

An ERROR: the current distance from the set-point.

  https://www.linkedin.com/in/d-r-dison

A CORRECTIVE RESPONSE: the behavior a deviation triggers, which grows with the error.

When the monitored variable drifts from its set-point, the error rises, and the drive commands corrective behavior until the error falls.

Cold triggers shivering and seeking warmth; low glucose triggers hunger and seeking food; rising carbon dioxide triggers the overwhelming urge to breathe.

This is negative-feedback control, the oldest idea in the study of living systems, and it is enough to state most of the brain's drives.

THE GENERALIZATION, and it is beautiful, is this.

A recent and influential body of work (homeostatic reinforcement learning (reinforcement learning, or RL, is the branch of machine learning in which an agent learns by trial and error from reward; the framework is due to Keramati and Gutkin)) proved that two things long thought separate are in fact identical: maximizing reward and minimizing drive.

If REWARD is defined as the reduction of the distance from a set-point, then a system that maximizes reward is exactly a system that defends its set-points, and the reward-learning machinery and the homeostatic machinery are one machine.

This is a load-bearing convergence for the framework, and it is THIS FRAMEWORK'S CHOICE to adopt it.

It means the dopamine reward-prediction-error signal of Chapter 7 and the drives of this chapter are not two systems to be bolted together.

They are the same system.

Dopamine reports progress toward a set-point.

A drive defines which set-point.

The reward signal is the derivative of drive satisfaction.

One loop, not two.

9.3 The Principal Drives, and Their Override Order

The drives, roughly from most to least overriding, with their monitored variable and their seat in the anatomy of this book:

RESPIRATION.

Variable: blood carbon dioxide and oxygen.

Seat: the medullary respiratory centres.

Override rank: absolute.

It cannot be suppressed voluntarily to the point of death; the drive seizes control first.

  https://www.linkedin.com/in/d-r-dison

This is the top of the override hierarchy.

THERMOREGULATION.

Variable: core temperature.

Seat: the preoptic hypothalamus.

FLUID AND OSMOTIC BALANCE.

Variable: blood osmolality.

Seat: the subfornical organ and hypothalamus, reading the blood directly through the circumventricular ports.

ENERGY.

Variable: glucose and fat stores.

Seat: the arcuate nucleus, reading leptin and ghrelin.

The hunger drive.

SLEEP.

Variable: accumulated adenosine (time awake).

Seat: the ventrolateral preoptic nucleus and the sleep switch.

The integrator of Chapter 8 feeding the gate of Chapter 8.

SAFETY, or THREAT AVOIDANCE.

Variable: presence of danger.

Seat: the amygdala and periaqueductal gray.

This drive can override most others acutely, which is why fear beats hunger.

SOCIAL CONNECTION.

Variable: belonging and attachment.

Seat: the oxytocin and opioid systems, and the social-pain circuitry of the anterior cingulate.

That social rejection recruits physical-pain circuitry is the evidence that connection is a defended set-point, not a luxury.

REPRODUCTION.

Variable: reproductive opportunity.

Seat: the hypothalamic-pituitary-gonadal axis.

Notably suppressed by the stress axis, which is why chronic stress reduces fertility, a direct competition between two drives.

  https://www.linkedin.com/in/d-r-dison

CURIOSITY, or INFORMATION.

Variable: prediction error and uncertainty.

Seat: the dopamine system itself.

This framework treats the reduction of uncertainty as an intrinsic drive, which is what makes the system explore rather than merely satisfy its bodily needs.

THIS FRAMEWORK'S CHOICE, and a consequential one: it is what makes the mind want to learn.

The override order is not fixed and absolute except at the extremes.

It is itself modulated: acute threat reorders everything, satiety lowers the hunger drive's priority, and the arbitration among simultaneously active drives is exactly the action-selection problem the basal ganglia solve.

The drives are the competitors; the basal ganglia are the referee; the winner gets the muscles.

This is the Arbitrated Federation of Chapter 12, now given its contestants.

9.4 How a Drive Connects to the Dynamics

A drive, in the running model, is a special INTEGRATOR-plus-COMPARATOR construct.

It integrates its monitored variable over time, compares it to its set-point, and emits an error signal.

That error does two things: it biases action selection toward behaviors that have reduced this error before (learned via the plasticity of Chapter 10), and, when it crosses a threshold, it can invoke the override hierarchy and seize control from slower deliberation.

The drive is thus the bridge between the body's state and the system's behavior, and it is the reason the system does anything at all rather than sitting inert.

CHAPTER 10. THE PLASTICITY: HOW THE SYSTEM CHANGES ITSELF

10.1 The Gap, Stated as a Patient

A brain with dynamics and drives but no plasticity is a brain frozen at one instant.

It can perceive, it can want, it can act, and it cannot LEARN.

It will make the same mistakes forever, because nothing it experiences changes what it is.

This is not a hypothetical.

It is the exact condition of patient H.M., whose story opened this volume: after the loss of his hippocampi he retained perception, intelligence,

  https://www.linkedin.com/in/d-r-dison

personality, and every skill and fact learned before the surgery, and he could form no new memories at all.

He lived permanently in the present.

A specification without a learning rule specifies H.M., which is a poignant irony given that H.M. is the case this whole framework centres on.

So this chapter supplies the one thing that separates a mind from a recording of a mind: a rule by which experience rewrites the machinery.

10.2 The Simplest Rule, and Why It Is Not Enough

The on-ramp is Hebb's rule, stated in 1949 and usually paraphrased as: neurons that fire together wire together.

In the vocabulary of this book: when two connected regions are active at the same time, the Conduit between them strengthens; when they are not, it weakens.

Formally, the change in a connection's weight is proportional to the product of the activity at its two ends.

This is enough to state as a first approximation, and it captures the core of long-term potentiation, the synaptic strengthening that Chapter 7 named as the leading mechanism of memory.

But by itself it has two fatal problems, and a specification must face them.

FIRST, it is unstable.

If firing together always strengthens, and strengthening makes firing-together more likely, then connections grow without bound until everything is maximally connected to everything, which is a seizure, not a memory.

Pure Hebbian learning explodes.

SECOND, it learns indiscriminately.

It strengthens every coincidence, whether the coincidence mattered or not, whether the outcome was good or bad.

It has no way to learn only what is worth learning.

Both problems are solved by the same two additions, and they turn the coarse rule into the one this framework adopts.

10.3 The Rule This Framework Adopts: Three-Factor, Homeostatically Bounded

The plasticity rule of this framework, and THIS FRAMEWORK'S CHOICE, has two components, each repairing one of the failures above.

COMPONENT ONE: THE THIRD FACTOR. Modern neuroscience has established that biological synapses do not change on coincidence alone. They change on coincidence GATED BY A NEUROMODULATOR. This is called three-factor, or neo-Hebbian, learning, and the rule is:

  https://www.linkedin.com/in/d-r-dison

weight change = (activity at one end) times (activity at the other end) times (a neuromodulatory signal)

The first two factors are Hebb's coincidence.

The third factor is the broadcast chemistry of Chapter 7, chiefly dopamine (carrying reward prediction error) and acetylcholine (carrying attention and the signal that something is worth encoding).

The third factor solves the indiscriminate-learning problem directly: a coincidence is written into the weights only when the neuromodulator says it mattered.

The brain does not learn everything that co-occurs; it learns what co-occurs AND is followed by reward, surprise, or attention.

This is the convergence the last three chapters have been building toward, and it is worth stating flatly.

The molecules of Chapter 7 are not a separate topic from learning.

Dopamine, the reward signal, IS the third factor.

The drive-reduction of Chapter 9 defines the reward.

The dynamics of Chapter 8 produce the coincidences.

Plasticity binds all three into one loop: the drives say what matters, the dopamine signal broadcasts when it happens, and the weights change accordingly, which alters the dynamics, which changes the behavior that serves the drives.

This is the whole engine of a learning mind, and every part of it is already in this book.

A mechanism called the ELIGIBILITY TRACE makes the timing work: a recently active connection leaves a fading tag, so that a reward arriving a second or two after the action can still find and strengthen the connection that earned it.

This bridges the gap between a fast coincidence and a slower reward, and it is well supported experimentally.

COMPONENT TWO: HOMEOSTATIC BOUNDING. To solve the instability, the framework adds a slow, stabilizing process called homeostatic plasticity or synaptic scaling: each region continuously adjusts its overall connection strengths to keep its average activity near a target level. If a region becomes too active, all its inputs are scaled down; if too quiet, scaled up. This is a negative feedback on the learning itself, and it keeps the fast three-factor learning from exploding. It is the plasticity of the drives applied to plasticity: the system defends a set-point for its own activity.

Together: fast, selective, neuromodulator-gated strengthening for learning what matters, bounded by slow homeostatic scaling for stability.

That pair is the complete learning rule the specification needs, and it is stated in one paragraph.

  https://www.linkedin.com/in/d-r-dison

10.4 The Timescales of Change

Plasticity operates on a hierarchy of timescales, and a faithful model represents more than one.

FAST, seconds to minutes: the eligibility trace and short-term facilitation, the temporary marking of what just happened.

MEDIUM, minutes to hours: long-term potentiation and depression, the weight changes that are the substance of new memory, requiring the neuromodulatory third factor to consolidate.

SLOW, hours to a lifetime: structural plasticity, the actual growth and pruning of connections, driven by brain-derived neurotrophic factor (Chapter 7), and consolidated during the sleep replay of Chapter 5.

This is the difference between writing to cache and committing to disk, and it is why sleep is a part of learning and not a pause in it.

VERY SLOW, developmental: the critical periods early in life when the system is maximally plastic and its large-scale wiring is set, after which the architecture largely fixes and only the weights within it continue to change.

THIS FRAMEWORK'S CHOICE is to model the adult regime, weights changing within a fixed architecture, and to treat development as the setting of initial conditions rather than as an ongoing process, for the same reason the framework models the adult brain throughout: it is the tractable starting point.

10.5 What Plasticity Does to the Map

The deepest consequence, and the one that closes this volume's argument, is that plasticity makes the map ALIVE.

Every prior chapter described fixed structure: these regions, these interfaces, these molecules, these rhythms.

Plasticity is the rule by which that structure edits itself in response to experience.

The connectome is not a diagram the brain is born with and keeps; it is a diagram the brain continuously rewrites, within limits, according to what happens to it.

A brain region is defined by its interfaces, as Chapter 1 insisted; plasticity is the process by which experience changes those interfaces, and therefore changes what the region is.

The map is not just alive.

It draws itself.

CHAPTER 11. THE SPECIFICATION BLOCK: DYNAMICS, DRIVES, AND PLASTICITY PER CONSTRUCT

11.1 Why On the Construct, Not in an Essay

  https://www.linkedin.com/in/d-r-dison

A specification wants the rule written on the component.

An implementer reading the entry for the hippocampus should find, in that entry, everything needed to code it: its function (already given in Chapter 2), and now its archetype, its state variables, its update rule, the drives it serves, and how its connections learn.

Gathering these in a separate essay would force the implementer to cross-reference; putting them on the construct makes each entry self-contained.

This is why the qualities are added as a per-construct BLOCK.

The block has a fixed, compact shape, so that it can be read uniformly and, later, parsed.

For each construct:

DYNAMICS: its archetype (from Chapter 8's six), its principal state variables, and its update rule in one line.

DRIVES: which drive or drives it serves or reports to, if any (many regions serve none directly).

PLASTICITY: whether and how its connections change, including which neuromodulatory third factor gates its learning, if any.

MODULATION: which neuromodulators set its gain, from Chapter 7.

Chapter 11 gives worked, expanded blocks for a set of constructs chosen to cover every archetype and every drive, so that each archetype is fully illustrated once.

A compact block of the same four-line shape is then attached DIRECTLY to every named region, interface, and molecule at its own entry in Chapters 2, 3, and 7, so that no construct lacks its specification.

Where a field genuinely does not apply to a construct, it is marked (NA), meaning Not Applicable, together with the reason.

11.1b A PLAIN-LANGUAGE KEY TO THE BLOCK

Before the blocks themselves, here is what each line means, in ordinary words, so that any reader can read any block in the book without a neuroscience or engineering background.

The block is deliberately terse because a specification must be; this key is its translation.

DYNAMICS answers: what kind of thing is this, and how does it change moment to moment?

It names one of six simple behaviours, and you can picture each with an everyday object.

A RELAY is a wire or a pipe. It passes its input along, maybe louder or quieter, without changing the message. A microphone cable.

  https://www.linkedin.com/in/d-r-dison

An INTEGRATOR is a bucket filling with water. It adds up what comes in and slowly leaks. It remembers recent history as a running total. A kitchen tally.

An OSCILLATOR is a metronome. It keeps a beat, and it opens and closes a gate in time with that beat, so things can only pass on the beat.

An ATTRACTOR is a marble rolling into the nearest of several dips in a landscape. Once it settles in a dip it stays there, and if you nudge it with half a memory it rolls to the whole one. This is how memory holds a pattern and completes it from a hint.

A GATE is a light switch with a stiff spring. It sits in one position, on or off, and won't flip until pushed hard, then resists flipping back. This is how the brain decides one action and sticks with it.

A COMPARATOR is a pair of scales. It weighs what it expected against what actually happened and reports only the difference, the surprise. This is how the brain notices error and reward.

STATE answers: what few numbers describe this part right now?

Usually just how active it is, how strongly it is reacting (its GAIN), and how tired or used-up it is (its ADAPTATION).

DRIVES answers: what, if anything, does this part want?

Most parts want nothing; they just do their job, and the block says (NA), meaning not applicable.

A few parts are the body's needs made into goals, hunger, safety, sleep, and those are named.

PLASTICITY answers: does this part learn, and how?

Learning here means connections getting stronger or weaker with experience.

The key phrase is the THIRD FACTOR: a connection changes not just when two parts are active together, but only when a chemical messenger, usually dopamine, signals that the moment MATTERED, because it brought reward or surprise.

So the block often reads "dopamine-gated": it learns what pays off, not everything that happens.

Where a part does not learn, it reads (NA).

MODULATION answers: which chemical messengers turn this part up or down?

These are the mood-and-arousal chemicals of Chapter 7, dopamine, norepinephrine, serotonin, acetylcholine, setting how strongly the part reacts, the way a dimmer sets a light.

With this key, every block in the book can be read as a plain sentence.

For example, the hippocampus block reads, in translation: it behaves like a marble-in-a-landscape memory driven by a metronome rhythm; it wants nothing directly but serves curiosity; it learns heavily, and only when dopamine

  https://www.linkedin.com/in/d-r-dison

says the moment mattered; and acetylcholine sets whether it is in writing or reading mode.

That is the whole of what an implementer, or a curious reader, needs to know to picture it working.

11.2 THE SPECIFICATION BLOCKS

PRIMARY VISUAL CORTEX (V1)

DYNAMICS: archetype RELAY, with local ATTRACTOR dynamics for contour completion. State: activation per orientation-and-location channel; gain. Update rule: output tracks lateral geniculate input, scaled by attentional gain and sharpened by local lateral inhibition, in near-real-time.

DRIVES: none directly; supplies the perceptual input on which all drives act.

PLASTICITY: strong in the developmental critical period, limited in the adult; connection weights refine with visual experience, gated by acetylcholine (attention) and behavioral relevance.

MODULATION: acetylcholine and norepinephrine set gain; attention raises the gain of attended channels.

THE HIPPOCAMPUS

DYNAMICS: archetype ATTRACTOR (CA3 is an autoassociative point-and-pattern attractor) driven by an OSCILLATOR (theta). State: current activation pattern (the represented place or memory), theta phase. Update rule: recurrent CA3 connections pull partial cues toward the nearest stored pattern (pattern completion); encoding and retrieval alternate by theta phase; sharp-wave ripples replay patterns during rest.

DRIVES: serves the information/curiosity drive (novelty) and, through spatial memory, the foraging aspect of the energy drive.

PLASTICITY: the highest plasticity in the brain. Schaffer-collateral weights change by three-factor rule: coincidence gated by dopamine (novelty signal from the ventral tegmental area) and consolidated during sleep replay. This is the site where the learning rule of Chapter 10 is most directly instantiated.

MODULATION: acetylcholine (from the medial septum) sets the encoding-versus-retrieval balance and paces theta; dopamine gates what is stored.

THE AMYGDALA (basolateral and central)

DYNAMICS: basolateral is a COMPARATOR-and-ATTRACTOR (threat detection, association); central is a GATE (threat output). State: threat activation. Update rule: basolateral integrates sensory and contextual input and fires when a learned threat pattern is matched, including via the fast subcortical low road; central gates the autonomic and behavioral threat response.

DRIVES: the seat of the SAFETY drive. Reports threat error to the whole system and can invoke the override hierarchy acutely.

  https://www.linkedin.com/in/d-r-dison

PLASTICITY: fear conditioning is rapid three-factor plasticity at basolateral synapses; fear extinction is new learning at the intercalated cells gated by prefrontal input. This is why fear is learned in one trial and unlearned slowly.

MODULATION: norepinephrine (arousal) strengthens threat learning, which is why arousing events are vividly remembered.

THE BASAL GANGLIA (striatum and output nuclei)

DYNAMICS: archetype GATE, inhibitory-by-default. State: activation of each competing action channel. Update rule: the striatum receives candidate actions, and the direct and indirect pathways select one by releasing its inhibition while suppressing competitors; the output is action selection by disinhibition.

DRIVES: the ARBITER of all drives. It does not have preferences; it selects among the actions the drives propose. This is the referee of the Arbitrated Federation.

PLASTICITY: corticostriatal weights change by three-factor rule with DOPAMINE as the third factor, carrying reward prediction error (RPE) from the substantia nigra and ventral tegmental area. This is reinforcement learning (RL) made physical: actions followed by reward are strengthened. The single clearest bridge from biology to the reinforcement-learning algorithm.

MODULATION: dopamine biases selection toward go (via D1) and away from no-go (via D2); its loss is Parkinsonian freezing.

THE THALAMIC RETICULAR NUCLEUS (TRN)

DYNAMICS: archetype OSCILLATOR-and-GATE. State: inhibitory activation per thalamic channel, phase (for spindles). Update rule: generates rhythmic inhibition that gates which thalamic relays reach cortex; prefrontal input biases the gate toward task-relevant channels.

DRIVES: serves attention, which is in service of whichever drive is currently dominant.

PLASTICITY: limited; its gating is set largely by top-down prefrontal signals rather than by local learning.

MODULATION: acetylcholine shifts it between the spindle-generating (sleep) and relay-gating (waking) modes.

THE INFERIOR OLIVE (as region and as interface)

DYNAMICS: archetype COMPARATOR. State: error activation. Update rule: computes the mismatch between an intended movement (efference copy from the red nucleus) and its sensory outcome, and emits the difference as the climbing-fibre error signal.

DRIVES: none directly; serves motor accuracy.

PLASTICITY: it is the TEACHER of cerebellar plasticity, not the site of it. Its error signal is the third factor that gates weight change at the parallel-fibre-to-Purkinje synapse. This is supervised learning made physical.

  https://www.linkedin.com/in/d-r-dison

MODULATION: minimal; it is a dedicated error computer.

THE LOCUS COERULEUS

DYNAMICS: archetype RELAY-and-COMPARATOR broadcasting globally. State: tonic firing level (arousal) and phasic bursts (surprise). Update rule: tonic output tracks overall arousal; phasic output fires on unexpected salient events; both broadcast norepinephrine brain-wide.

DRIVES: serves arousal, the enabling condition of all drives.

PLASTICITY: itself relatively fixed; it is a source of the third factor (norepinephrine) that gates plasticity elsewhere.

MODULATION: it is a modulator; orexin and CRH (corticotropin-releasing hormone) drive it.

THE VENTRAL TEGMENTAL AREA (VTA) AND SUBSTANTIA NIGRA

DYNAMICS: archetype COMPARATOR broadcasting globally. State: tonic dopamine level (ongoing value) and phasic bursts and pauses (reward prediction error). Update rule: phasic output equals actual minus expected reward; tonic output tracks state value; both broadcast dopamine.

DRIVES: computes the reward signal that IS drive-reduction (Chapter 9), and so it is the common currency in which all drives are compared.

PLASTICITY: the source of the dopamine third factor that gates plasticity across the striatum, hippocampus, and cortex. The single most important plasticity-gating signal in the brain.

MODULATION: the lateral habenula inhibits it on disappointment; orexin and many inputs drive it.

THE HYPOTHALAMUS (as the drive hub)

DYNAMICS: archetype INTEGRATOR-and-COMPARATOR, one per drive. State: the monitored variable and its error, per drive (temperature, osmolality, energy, and the rest). Update rule: each nucleus integrates its bodily variable, compares to a set-point, and emits a drive error that biases behavior and commands the endocrine and autonomic response.

DRIVES: the SEAT of most homeostatic drives (Chapter 9). This is where the body's needs become the system's goals.

PLASTICITY: set-points can shift slowly (allostasis), for instance the raising of the stress set-point under chronic stress; otherwise the drive machinery is largely fixed.

MODULATION: reads the blood directly at the circumventricular organs, bypassing the usual barriers; a layer-skipping input.

A TRANSMISSIVE INTERFACE (for example, the corticospinal tract)

DYNAMICS: archetype RELAY. State: signals in transit. Update rule: apply delay (from length and myelination) and weight; deliver unchanged. A perfect wire with a lag.

  https://www.linkedin.com/in/d-r-dison

DRIVES: none; serves motor output.

PLASTICITY: the weight (effective strength) changes slowly with use and practice, which is part of motor skill learning; the routing is fixed.

MODULATION: minimal.

A COMPUTATIONAL INTERFACE (for example, the intercalated cells of the amygdala)

DYNAMICS: archetype GATE performing sign inversion. State: inhibitory activation. Update rule: convert excitatory prefrontal input into inhibition of the central amygdala; the transform is the function.

DRIVES: serves the safety drive, specifically its down-regulation (extinction).

PLASTICITY: THIS is where fear extinction is stored; its weights change by three-factor rule gated by prefrontal input and success signals. The clinical target of exposure therapy.

MODULATION: set by prefrontal top-down control.

A NEUROMODULATOR (for example, dopamine, as a construct in its own right)

DYNAMICS: it is not a region but a BROADCAST FIELD. State: its concentration in each target territory. Update rule: released by its source nuclei (phasic and tonic), it diffuses to its targets and sets their gain and their plasticity third factor, then is cleared by reuptake. Modeled not as an edge but as a global parameter with per-target levels.

DRIVES: dopamine specifically carries the reward-prediction-error that unifies the drives.

PLASTICITY: it is a CAUSE of plasticity elsewhere, being the third factor, not a site of it.

MODULATION: it is modulation, by definition.

11.3 The Template, For the Remaining Constructs

The eleven blocks above cover all six dynamical archetypes, the arbiter, the drive hub, both interface kinds, and the broadcast fields, each illustrated in full.

The same four-line block, DYNAMICS, DRIVES, PLASTICITY, MODULATION, is now attached to every named construct at its own entry, throughout Chapters 2, 3, and 7.

Where a field does not apply, it reads (NA).

The reader will find, at the foot of each region, interface, and molecule, a SPECIFICATION BLOCK giving its archetype and state, the drive it serves, how it learns, and what modulates it.

This is the per-construct grounding the SPARC Specification phase consumes directly.

  https://www.linkedin.com/in/d-r-dison

CHAPTER 12. IS THERE AN ORCHESTRATOR?

12.1 The Question

Chapters 2 through 4 have catalogued roughly a hundred regions and forty interfaces, each with its own function, each talking to its neighbours.

The question a systems engineer asks upon seeing such a diagram is unavoidable, and it is the right question.

Who is in charge?

Is there a conductor of this orchestra: a central controller that reads the sensors, decides, and issues commands?

Or is this a federation of autonomous parts, with the appearance of unified behaviour arising from their interaction?

Or something else again?

The answer matters, because it determines what kind of thing a person is.

The short answer is that there is no orchestrator, that "federation" is closer but still wrong, and that the true architecture has no name in neuroscience.

We shall give it one, and we shall situate it carefully against the existing literature, because this question has a long history and a great many people have got it wrong in instructive ways.

12.2 The Four Candidate Architectures

Before adjudicating, the options must be stated clearly.

Borrowing the vocabulary of distributed systems engineering, there are four.

• A CENTRALIZED architecture has one controller.

All information flows to it; all commands flow from it.

The controller has global state and final authority.

This is the client-server model, the conductor, the chief executive.

• A FEDERATED architecture has autonomous peers.

Each acts on local information; global behaviour emerges from local interaction.

No peer has authority over another.

This is the ant colony, the market, the peer-to-peer network.

• A HYBRID architecture has autonomous peers plus one or more coordinating services that the peers voluntarily consult, but which cannot compel them.

  https://www.linkedin.com/in/d-r-dison

• A HETERARCHICAL architecture has multiple overlapping authority structures, each supreme in its own domain, with no single ordering that ranks them all.

We shall argue that the brain is none of these four exactly, but is closest to the fourth, and that its specific variant deserves its own name.

12.3 The Case Against a Central Orchestrator

The intuition that there must be a central controller is very strong, very old, and wrong.

It is wrong for five reasons, each independently sufficient.

FIRST: EVERY CANDIDATE HAS BEEN PROPOSED AND EVERY CANDIDATE HAS FAILED

The prefrontal cortex is the usual nominee.

But prefrontal damage does not abolish the self.

It produces a person with impaired planning who remains unmistakably a person, who continues to perceive, feel, remember, and act.

Phineas Gage survived and remained Phineas Gage, altered but present.

The thalamus is a better nominee, since everything passes through it and thalamic damage does abolish consciousness.

But the thalamus decides nothing.

It gates and relays.

A switchboard is not a chief executive, and the fact that cutting the phone lines silences the company does not make the switchboard operator the president.

The claustrum was Crick's candidate and remains the most interesting.

But its function is genuinely unknown, and one case of stimulation-induced unconsciousness is a clue, not a demonstration of command.

The reticular formation controls the level of consciousness but not its content.

It can turn the lights on and off.

It cannot decide what to do.

There is no region whose destruction removes the commander while leaving the machinery intact.

There is no throne room.

This is not for want of looking.

SECOND: THE INFINITE REGRESS

Suppose there were an orchestrator: a region receiving all information and making all decisions.

  https://www.linkedin.com/in/d-r-dison

We should then have to ask how that region decides.

Does it contain a smaller orchestrator?

The question recurs, and recurs forever.

This is the homunculus fallacy, and Dennett named its architectural form the Cartesian Theater: the tacit assumption that there is a place in the brain where "it all comes together" and is presented to an inner audience.

Positing a little person inside the head to explain the person does not explain the person.

It relocates the mystery and multiplies it.

The only escape is to accept that decisions are made by the interaction of parts, none of which is itself a decider.

This is not a concession.

It is the whole content of the naturalistic explanation of mind, and any architecture that quietly smuggles a controller back in has failed at the first step.

THIRD: THE ANATOMY DOES NOT PROVIDE A CONVERGENCE ZONE

A central controller requires a central place where all information converges.

Chapter 2 catalogued the convergence zones of the brain, and there is no such place.

Information does converge, repeatedly, but each convergence is partial and purpose-built.

The superior colliculus converges vision, hearing, and touch, but only into a spatial map for orienting.

The hippocampus converges the "what" and "where" streams, but only for memory.

The orbitofrontal cortex converges all five senses, but only to compute value.

The hypothalamus converges every interoceptive signal, but only to regulate the body.

Each convergence serves a purpose.

None serves a controller.

FOURTH: BEHAVIOUR SURVIVES REMOVAL OF THE SUPPOSED CONTROLLER

The spinal cord walks without the brain: its central pattern generators produce coordinated locomotion in a decerebrate preparation.

The pre-Botzinger complex breathes without instruction, and continues to do so in a dish.

  https://www.linkedin.com/in/d-r-dison

The enteric nervous system runs the gut with the vagus severed.

The cerebellum corrects movements the cortex never learns about.

If there were a controller, these behaviours should cease when it is removed.

They do not.

FIFTH: THE SPLIT BRAIN

When the corpus callosum is cut, the patient does not lose half a mind.

Two centres of experience appear to result, each capable of independent perception and choice, each unaware of the other's knowledge, and, most tellingly, the left hemisphere confabulates explanations for actions the right hemisphere initiated, with complete sincerity.

This is close to decisive.

A unified self that could be divided by a knife was never a unified thing sitting somewhere.

It was a process, and the process could be run twice.

12.4 The Case Against a Pure Federation

Having rejected the orchestrator, the tempting move is to swing to the opposite pole: a flat federation of autonomous modules, global behaviour emerging from local interaction, as in an ant colony.

This is closer to the truth, and it remains wrong, for three reasons.

FIRST: THERE ARE PURPOSE-BUILT ARBITRATION STRUCTURES

A pure federation has no arbiter.

Conflicts are settled by whoever pushes hardest.

The brain does not work this way.

It contains structures whose entire function is arbitration, and the argument for this is not speculative: it is the central thesis of one of the most influential papers in systems neuroscience.

Redgrave, Prescott, and Gurney posed the problem in its general form.

A selection problem arises whenever two or more competing systems seek simultaneous access to a restricted resource, and the body is exactly such a resource: there is one set of muscles, and only one action can be performed with them at a time.

They then made a decisive architectural observation.

A distributed selection mechanism, in which every module must negotiate with every other, requires connections growing as the square of the number of modules.

A central switch, reciprocally connected to all of them, requires only 2n connections and grows by two for each module added.

Centralized selection is vastly cheaper to wire, and evolution, they argued, found it.

The basal ganglia are that switch: a vertebrate solution to the selection problem.

The architecture described in Section 2.2.4 is exactly what such a switch would look like.

It is inhibitory by default: every candidate action is held suppressed, and selection consists of releasing exactly one.

This is not emergence.

This is a referee, and its failure modes prove it: too little output and unwanted actions escape as chorea, too much and no action can be released at all, which is Parkinsonian akinesia.

The thalamic reticular nucleus arbitrates attention on precisely the same inhibitory-by-default principle.

The ventrolateral preoptic nucleus and the ascending arousal system arbitrate sleep and waking through a bistable flip-flop.

These are dedicated referees, and a federation with referees is not a pure federation.

SECOND: THERE ARE BROADCAST AUTHORITIES

A pure federation has no global signals.

The brain has several, and they are not peers.

The locus coeruleus, from fifteen thousand cells, can change the operating state of the entire cortex at once.

The nucleus basalis can raise the plasticity of the whole cortical sheet.

The hypothalamus, by releasing a hormone, can change the state of every cell in the body simultaneously.

These are not negotiations among equals.

They are broadcasts with authority, and the recipients do not vote.

THIRD: THERE IS A DEMONSTRABLE OVERRIDE HIERARCHY

In a flat federation, no module outranks another.

In the brain, ranking is demonstrable and its order is precise.

The hyperdirect pathway vetoes an action already in progress, overriding the deliberative machinery entirely.

The amygdala commandeers attention, autonomic state, and behaviour before the cortex has finished identifying what frightened it.

  https://www.linkedin.com/in/d-r-dison

The brainstem's respiratory drive will override any voluntary attempt to hold the breath, and it will win, every time, without exception.

You cannot decide not to breathe.

The federation has a constitution, and some of its clauses are not amendable by the cortex.

12.5 THE SEVEN RIVAL ANSWERS: A SURVEY AND ADJUDICATION

The question of this chapter is old, and it has received at least seven serious answers.

Intellectual honesty requires that each be stated in its strongest form and adjudicated on the evidence, rather than dismissed.

Several are partly right, and the architecture proposed in Section 12.6 borrows from four of them.

12.5.1 ANSWER ONE: THE CARTESIAN THEATER (the orchestrator exists)

THE CLAIM: There is a place where it all comes together and is presented to an observer, who decides.

WHO HOLDS IT: Descartes explicitly, locating the meeting point at the pineal gland.

Nobody defends it explicitly today, and almost everybody defends it implicitly, every time they write that "the brain decides" or "the prefrontal cortex evaluates the options."

THE STRONGEST CASE FOR IT: Experience really does feel unified and singular.

There does seem to be one of you.

Any theory that denies this owes an explanation of why the illusion is so total.

THE VERDICT: Refuted, on the five grounds of Section 12.3, of which the infinite regress is fatal by itself.

Dennett's diagnosis is correct and worth stating precisely: the Cartesian Theater is not a theory that happens to be false.

It is a theory that cannot be true, because it does not explain what it purports to explain.

It merely postpones the explanation by installing an audience, and then owes us an account of the audience.

WHAT SURVIVES: The datum that motivated it.

Experience is unified, and that unification requires a mechanism.

The mechanism is not a place.

It is the broadcast and loop architecture of Section 12.6.

  https://www.linkedin.com/in/d-r-dison

12.5.2 ANSWER TWO: PURE MODULARITY, OR THE SOCIETY OF MIND

THE CLAIM: The mind is a collection of specialized agents, each individually mindless, with no central control.

Intelligence is what their interaction looks like from outside.

WHO HOLDS IT: Minsky, in The Society of Mind, where the agents are mindless and intelligence emerges from their organization.

Selfridge's Pandemonium, where "demons" shout with intensity proportional to their confidence and the loudest is heard.

Fodor's modularity of mind, in a restricted form.

The massive modularity of evolutionary psychology, in an unrestricted one.

THE STRONGEST CASE FOR IT: The brain is manifestly modular.

Chapter 2 is a catalogue of specialized modules.

Lesions produce selective deficits, which is precisely what modularity predicts and what a general-purpose central processor does not.

THE VERDICT: Right about the modules, wrong about the absence of arbitration.

This is the pure federation refuted in Section 12.4.

Minsky's agents have no basal ganglia: they have no dedicated structure whose job is to decide which of them gets the muscles, and so the theory has no account of why exactly one action occurs when many are proposed.

Pandemonium's answer, that the loudest demon wins, is a selection mechanism, but an unimplemented one.

Redgrave's contribution was to show that the brain implements it, and where.

The theory is also, as its critics note, largely unfalsifiable as stated: it offers descriptive narratives rather than testable mechanisms.

WHAT SURVIVES: The modules, and the insight that they are individually mindless.

Both are retained in Section 12.6.

12.5.3 ANSWER THREE: THE GLOBAL WORKSPACE

THE CLAIM: Many specialized processors operate in parallel and unconsciously, competing for access to a limited-capacity workspace.

The winner is broadcast globally to all the others, and that broadcast is consciousness.

WHO HOLDS IT: Baars, who proposed it; Dehaene and Changeux, who gave it a neural implementation as the Global Neuronal Workspace, identifying the

  https://www.linkedin.com/in/d-r-dison

substrate as long-range frontoparietal pyramidal neurons and the mechanism as "ignition," a sudden nonlinear synchronization.

THE STRONGEST CASE FOR IT: It explains the most robust fact about consciousness, namely its severe capacity limit.

You can attend to one thing.

It explains why unconscious processing is parallel and massive while conscious processing is serial and narrow.

It makes testable predictions about ignition dynamics, and those predictions have substantially held.

THE COMMON MISREADING: That the workspace is an orchestrator, and that GWT therefore smuggles the homunculus back in.

Baars's own theatre metaphor invites this reading, and it is wrong.

The critical point is that nobody sits in the audience.

The workspace is not a controller.

It is a communication bus.

It has no preferences, forms no intentions, and issues no commands.

It broadcasts whatever wins a competition that it does not adjudicate.

Remove every specialized processor and the workspace has nothing to broadcast and nothing to say.

THE VERDICT: Substantially correct, and adopted into the present framework, but incomplete as an answer to this chapter's question.

GWT describes a mechanism of coordination.

It does not describe the arbitration structures that determine what wins, nor the override hierarchy, nor the neuromodulatory state-setting.

It is one of the four mechanisms of Section 12.6, not the whole architecture.

WHERE IT LIVES IN THIS VOLUME: The workspace is implemented by the long-range corticocortical interfaces of Section 3.4.6, gated by the Attentional Gating Interface (PUC) of Section 3.5.11, and supplied with content by whichever coalition wins the competition arbitrated by the basal ganglia and the thalamic reticular nucleus.

12.5.4 ANSWER FOUR: INTEGRATED INFORMATION THEORY

THE CLAIM: There is no controller and no broadcast.

Consciousness simply is the integrated information of a system: the degree to which its causal structure is irreducible to the sum of its parts.

The relevant quantity is called phi.

WHO HOLDS IT: Tononi, and Koch.

  https://www.linkedin.com/in/d-r-dison

THE STRONGEST CASE FOR IT: It explains why the cerebellum, with four times as many neurons as the cortex, contributes essentially nothing to consciousness.

Its circuitry is feedforward and modular, so its parts are nearly independent, so its integrated information is low.

This is a genuinely impressive prediction that no other theory makes as cleanly.

THE VERDICT: Live, unresolved, and orthogonal to the present question.

IIT locates the substrate in the posterior cortical hot zone rather than in frontal broadcast networks, and an adversarial collaboration is currently adjudicating between IIT and GWT on precisely this point.

Nothing in this chapter depends on the outcome, because the two theories agree entirely on the matter at issue here: neither posits a central controller.

IIT is, if anything, more radical in its denial, since it holds that consciousness is not produced by any mechanism of coordination at all, but is simply what a certain kind of causal structure is.

12.5.5 ANSWER FIVE: PREDICTIVE PROCESSING AND ACTIVE INFERENCE

THE CLAIM: The brain has no controller because it has no need of one.

It is a hierarchical generative model that continuously predicts its own sensory input and acts to minimize the discrepancy.

There is no executive issuing commands; there are only predictions cascading downward and prediction errors cascading upward, and behaviour is what happens when the system changes the world to match its predictions rather than changing its predictions to match the world.

WHO HOLDS IT: Friston, under the free energy principle; Clark and Hohwy, philosophically.

THE STRONGEST CASE FOR IT: It unifies perception and action under a single principle, which is a considerable theoretical achievement.

And it is anatomically well-supported by the very fact this volume has emphasized repeatedly: the corticothalamic feedback projection outnumbers the forward projection roughly ten to one.

That ratio is bizarre on a bottom-up view of perception and exactly what predictive processing requires.

The brain really is mostly telling itself what to expect.

THE VERDICT: Powerful, and complementary rather than rival.

Predictive processing explains what the regions of Chapter 2 are computing.

It does not explain how conflicts among them are resolved when two incompatible actions would each minimize prediction error, and the honest

  https://www.linkedin.com/in/d-r-dison

active inference answer, that precision-weighting settles it, is a description of arbitration rather than an alternative to it.

Precision weighting is largely implemented by neuromodulation, which is Mechanism Two of Section 12.6.

The theory has also been criticized as unfalsifiable in its strongest formulations, since a sufficiently flexible generative model can accommodate any observation.

WHAT SURVIVES: A great deal, and it is incorporated.

The loops of Mechanism Three are predictive loops, and the state-setting of Mechanism Two is precision-weighting by another name.

12.5.6 ANSWER SIX: THE SUBSUMPTION ARCHITECTURE

THE CLAIM: There is no central model and no central planner.

Control is layered: simple behavioural modules operate in parallel, each competent to act alone, each connected directly to sensors and actuators, and higher layers may suppress or subsume the outputs of lower ones without replacing them.

WHO HOLDS IT: Brooks, in robotics, from 1986.

THE STRONGEST CASE FOR IT: It works.

Brooks's robots walked, avoided obstacles, and explored, while classical robots with central world-models and planners were still computing their first move.

And the architecture was arrived at independently of neuroscience, by an engineer facing the same constraint evolution faced: build something that must keep working while you add to it.

THE VERDICT: Correct in outline, and the closest engineering analogue to the true architecture.

The override hierarchy of Mechanism Four is subsumption: the brainstem is the lowest layer, competent alone, and the cortex is a later layer that can suppress it, but only up to a point, and never on the matter of breathing.

Its limitation as an account of the brain is that it has no equivalent of the basal ganglia.

Subsumption resolves conflict by fixed priority wiring, decided by the designer in advance.

The brain resolves conflict by a learned, dopamine-tuned, dynamically reweighted selection process.

Subsumption is arbitration with the arbiter compiled in; the brain's arbiter is trainable.

WHAT SURVIVES: The layered override principle, adopted as Mechanism Four.

12.5.7 ANSWER SEVEN: THE HETERARCHY

  https://www.linkedin.com/in/d-r-dison

THE CLAIM: There is no single ordering of authority.

There are multiple overlapping authority structures, each supreme within its own domain, none supreme overall, with circular causality among them.

WHO HOLDS IT: A tradition running from McCulloch's 1945 observation that the nervous system exhibits circularities of preference that cannot be arranged in any consistent ranking, through cybernetics, to contemporary work on circular causality in biology.

THE STRONGEST CASE FOR IT: It is the only one of the seven that takes the loops seriously as a structural fact rather than an inconvenience.

If A controls B, B controls C, and C controls A, then asking which is in charge is not a hard question.

It is a malformed one.

McCulloch's original point deserves restating because it is sharper than the modern literature usually renders it: he showed that the nervous system's preferences can be intransitive, meaning it can prefer A to B, B to C, and C to A.

No hierarchy can produce that.

Only a heterarchy can.

THE VERDICT: The most nearly correct of the seven, and the base on which Section 12.6 is built.

Its deficiency is that it is a negative thesis.

It tells you what the architecture is not, namely a hierarchy, and it does not tell you what the coordinating mechanisms are.

The Arbitrated Federation is a heterarchy with the mechanisms filled in.

12.5.8 SUMMARY OF THE ADJUDICATION

Of the seven, one is refuted outright (the Cartesian Theater).

One is orthogonal and unresolved (Integrated Information Theory).

Four are substantially correct but partial, and each contributes exactly one mechanism to the synthesis:

• The Society of Mind contributes the specialized, individually mindless modules.

• The Global Workspace contributes the broadcast bus.

• Predictive Processing contributes the loops, and precision-weighting as state-setting.

• Subsumption contributes the layered override hierarchy.

• Heterarchy contributes the formal denial that any consistent ranking exists.

  https://www.linkedin.com/in/d-r-dison

What none of them supplies, and what the anatomy of Chapter 2 plainly demands, is dedicated arbitration: the basal ganglia, the thalamic reticular nucleus, the sleep switch.

Redgrave's insight, that selection is a distinct problem requiring a distinct organ, is the missing piece, and it is the piece that turns a federation into an Arbitrated Federation.

12.6 THE ARBITRATED FEDERATION (PUC)

DEFINITION

The Arbitrated Federation is the organizational architecture of the human nervous system: a heterarchical federation of semi-autonomous specialized regions, containing no central controller, coordinated by four mechanisms which are themselves specialized and distributed rather than centralized.

We propose the name because neuroscience possesses no term for this architecture, and consequently the field oscillates between two positions that are both wrong: a covert homuncularism that speaks casually of "the brain deciding," and a flat emergentism that denies there is any organizing structure at all.

The truth is a third thing, and a third thing needs a word.

THE FOUR COORDINATING MECHANISMS

The first is DEDICATED ARBITRATION.

Purpose-built selector structures resolve competition among incompatible demands.

The basal ganglia select actions.

The thalamic reticular nucleus selects attentional targets.

The sleep switch selects global state.

The crucial property of these structures is that they are referees, not rulers.

They determine which competitor wins; they have no preferences of their own; they generate no candidates.

The basal ganglia do not decide what you should do.

They decide which of the things you were already considering gets the muscles.

Strip out every cortical proposal and the basal ganglia have nothing to select among, and nothing to say.

The second is BROADCAST STATE-SETTING.

The neuromodulatory nuclei and the endocrine system do not issue commands; they change the parameters within which every region computes.

Norepinephrine does not tell the cortex what to do.

  https://www.linkedin.com/in/d-r-dison

It tells the cortex how urgently to do whatever it is already doing.

This is governance by adjusting the constants, not by issuing instructions, and it is the Layer 9 instantiation of the Broadcast Principle set out in the companion volume: it exists precisely where the message is for everyone and an address would be waste.

The third is LOOP CLOSURE.

Nearly every major pathway catalogued in this volume is a loop, not an arc.

Cortex to basal ganglia to thalamus to cortex.

Cortex to pons to cerebellum to thalamus to cortex.

Hippocampus to mammillary bodies to anterior thalamus to cingulate to hippocampus.

In an architecture of loops, no region is upstream in any absolute sense, because every region is downstream of itself.

Control is circular, and circular control has no top.

This is the formal reason the search for a controller in Section 12.3 was doomed: you cannot find the top of a circle.

The fourth is OVERRIDE HIERARCHY.

A partial ordering exists, and it is not arbitrary.

It is ordered by evolutionary age and by lethality: the oldest and most vital systems can override the newest.

Breathing beats deliberation.

Threat detection beats reasoning.

A sufficiently strong interoceptive signal beats any plan.

The cortex is not the master of this house.

It is the most recently arrived tenant, with the largest room and the least authority in an emergency.

THE PRECEDENT IN ENGINEERING

This architecture is not unprecedented.

It was independently discovered by robotics.

Brooks's subsumption architecture abandoned central planning in favour of layered behavioural modules operating in parallel, each competent to act on its own, with higher layers able to suppress or subsume the outputs of lower ones but never to replace them.

There is no central model of the world and no central planner.

  https://www.linkedin.com/in/d-r-dison

Intelligence, Brooks argued, is determined by the dynamics of interaction with the world, and his robots walked while classical planning robots were still computing.

The convergence is striking, and it is not a coincidence.

Both systems were built under the same constraint: incremental construction, with no permission to stop working during the upgrade.

THE CONSEQUENCE

Under this architecture, a decision is not made anywhere.

It is a settlement, reached among competing coalitions of regions, adjudicated by structures that do not care about the outcome, under global parameters set by broadcast systems that do not know what is being decided, within loops that have no beginning, subject to override by older systems that cannot be reasoned with.

The feeling that there is someone inside making the decisions is itself produced by the system, and it is produced late.

It is generated substantially by the Default Mode Network Core Triad (Section 2.1.11), whose function is precisely the construction of a coherent self-narrative, and the split-brain confabulation results show what that machinery does when it is denied the facts: it does not report uncertainty.

It invents a reason, and believes it.

The narrator is real.

The narrator is a region.

The narrator is not in charge.

12.7 Why Evolution Built It This Way

The architecture is not an accident, and it is not a failure to achieve centralization.

It is what you get when a control system must be assembled incrementally, over hundreds of millions of years, and is never once permitted to stop working.

A central orchestrator cannot be built incrementally.

It must be designed in advance, with knowledge of every subsystem it will eventually command.

Evolution has no such foresight, and it cannot take the organism offline for a rewrite.

A creature whose brain is under reconstruction is a creature that is eaten.

What evolution can do is add a module alongside the existing ones, wire it in, and let it compete for control of the muscles.

  https://www.linkedin.com/in/d-r-dison

That is the observed history, and it is legible in the anatomy of Chapter 2.

The brainstem runs the body.

The midbrain adds orienting.

The basal ganglia add selection.

The limbic system adds valuation.

The cortex adds modelling.

Nothing was ever removed.

Each new system was added on top and granted partial, negotiated, revocable authority over what was already there.

The old systems kept their veto, which is why you cannot hold your breath to death, and why fear is faster than thought.

The Arbitrated Federation is therefore not merely a description of how the human brain happens to be organized.

It is a prediction about how any control system must be organized if it is grown rather than designed.

It is the reason the brain is a city rather than a machine: not planned, never demolished, everything built on the foundations of something older, and governed by a negotiation among districts rather than by a mayor.

12.8 A Note on Metaphors: Why It Is Not an Orchestra

The orchestra is the metaphor that everyone reaches for, and it is worth being explicit about why it must be rejected, because it fails in exactly the place where the question is decided.

An orchestra has a conductor.

The conductor has global information, sets the tempo, cues the entrances, and possesses unambiguous authority.

That is precisely the centralized architecture refuted in Section 12.3, and to invoke the orchestra as an image of the brain is to smuggle the homunculus back in through the metaphor after having thrown it out through the argument.

The metaphor works against the claim it is usually offered to support.

The better image is a jazz ensemble with no conductor at all.

Each player is autonomous and competent alone.

Each is listening to the others and responding in real time.

Leadership exists, but it is transient and negotiated: one instrument takes the solo, and the others accompany, and then it passes.

  https://www.linkedin.com/in/d-r-dison

Coherence is real and audible, and it is produced by mutual constraint rather than by command.

Nobody is in charge, and the result is not chaos.

The analogy is imperfect in one respect that is worth naming, and the imperfection is instructive.

A jazz ensemble has no basal ganglia.

Its players can, in principle, all solo at once, and the ensemble degrades gracefully into noise if they do.

The brain cannot afford that failure mode, because there is only one body and only one action can be performed with it.

This is exactly the selection problem of Redgrave, Prescott, and Gurney, and it is why the brain contains a dedicated arbiter and a jazz quartet does not.

So the corrected metaphor is a jazz ensemble with a referee who has no musical opinions: someone who does not choose the notes, and who cannot play, but who ensures that exactly one solo happens at a time.

That is an odd image, and its oddness is the point.

The architecture has no everyday analogue, which is precisely why it needed a name.

12.9 The Design Implication, for Those Building Systems

This chapter has an engineering corollary, and it is worth stating because the question of this chapter is being asked with increasing urgency outside neuroscience, by people building multi-agent software systems.

The lesson of the Arbitrated Federation is that the choice between central orchestration and pure emergence is a false one, and that both failure modes are real.

Pure emergence, with no arbitration, yields subsystems that compete for the same resource and deadlock, or that all act at once and produce incoherence.

This is a nervous system without basal ganglia, and its clinical name is chorea.

Rigid central orchestration yields brittleness: a single point of failure, an inability to degrade gracefully, and an architecture that cannot be extended without redesigning the controller.

It is also, as Section 12.7 argued, unbuildable incrementally, which is why evolution never built it.

The biological solution is neither, and it decomposes into four design rules.

- Grant local autonomy.

  https://www.linkedin.com/in/d-r-dison

Each module should be competent alone, and should act on local information without waiting for permission.

• Define the interfaces rigorously.

This is the burden of Chapters 3 and 4 of this volume, and it is the load-bearing wall of the whole architecture.

Modules influence each other only through defined interfaces, and the interfaces, as Section 3.6 showed, are frequently where the real computation lives.

• Add dedicated arbiters, not commanders.

The basal ganglia are the model: a structure that resolves conflict without having preferences of its own, that generates no candidates, and that selects among proposals it did not author.

An arbiter that starts having opinions has become a dictator, and the architecture has failed.

• Coordinate by broadcast state, not by instruction.

Global signals should set the parameters within which modules compute, rather than telling them what to compute.

This is what neuromodulation does, and it is cheap, fast, and robust precisely because it carries no content that could be wrong.

The resulting system is what might be called organized emergence: neither chaos nor command.

It is what evolution arrived at over six hundred million years, under the hardest possible constraint, which is that the system was never once allowed to stop working while it was being modified.

Anyone building a system under that same constraint should expect to arrive somewhere similar.

12.10 What Would Falsify the Arbitrated Federation

A theory that cannot be refuted is not a theory, and the four mechanisms make specific claims that could fail.

The framework would be falsified, or seriously damaged, by any of the following.

• The discovery of a region whose destruction abolishes decision-making while leaving perception, memory, emotion, and motor capacity intact.

This would be an orchestrator, and it would refute the central claim.

No such region has been found, and the framework predicts none will be.

• A demonstration that the basal ganglia generate action candidates rather than merely selecting among them.

This would convert the referee into a ruler, and would collapse Mechanism One.

  https://www.linkedin.com/in/d-r-dison

• A demonstration that the neuromodulatory systems carry specific content rather than setting global parameters.

If it were shown that a particular locus coeruleus firing pattern instructs the cortex what to do, rather than how urgently to do what it is already doing, Mechanism Two would fail.

• A demonstration that the override hierarchy is not ordered by evolutionary age and lethality.

The framework predicts, concretely, that no amount of cortical training will permit a person to voluntarily suppress breathing to the point of death, and that no psychotherapy will abolish the amygdala's capacity to preempt cortical processing.

If either were achieved, Mechanism Four would fail.

• The demonstration of a consistent, transitive ranking of authority across all brain systems.

This would show the architecture is a hierarchy after all, and would refute the heterarchical claim inherited from McCulloch.

The framework accepts these risks.

12.11 The Answer, Stated Plainly

The question was: is there an orchestrator, or do the parts act independently, and is the architecture federated, hybrid, or other?

The answer is: other, and here it is in one paragraph.

There is no orchestrator, and there could not be one, because a controller of the parts would itself require a controller, and the regress does not terminate.

But neither are the parts independent, because they compete for a single body and something must decide which of them gets it.

The brain resolves this by an architecture that has no name in the literature and which we have called the Arbitrated Federation: a heterarchy of specialized, semi-autonomous, individually mindless regions, with no consistent ranking of authority, coordinated by four mechanisms that are themselves specialized organs rather than central executives.

Dedicated referees settle competition without having preferences of their own.

Broadcast systems set the parameters of computation without knowing what is being computed.

Loops close, so that no region is upstream in any absolute sense and control is circular.

And an override hierarchy, ordered by evolutionary age and by lethality, ensures that the oldest systems can always veto the newest, which is why fear is faster than thought and why you cannot decide not to breathe.

  https://www.linkedin.com/in/d-r-dison

Decisions are not made anywhere.

They are settlements.

And the part of you that feels certain it was in charge the whole time is a region on a map, in the medial prefrontal cortex and the posterior cingulate, whose job is to tell a coherent story afterwards, and which, when the facts are withheld from it, does not fall silent.

It makes something up, and believes it.

CHAPTER 13. THE THREE-STAGE PROVING SYSTEM OF PREVIOUSLY UNNAMED CONSTRUCTS (PUC)

13.1 Why the Names Must Be Proved

Twenty times in this volume, a construct has been marked with the tag (PUC), Previously Unnamed Construct, and given a name.

That is a strong act, and it carries a burden.

To claim that a thing is real, is used, and yet has never been named is to make three assertions at once, and each can be wrong.

The thing might not be real.

It might already have a name the author overlooked.

Or the proposed name might be a poor one that the field is right to reject.

A responsible naming, therefore, is not a christening but a proof.

This chapter supplies that proof for every one of the twenty constructs, using the same three-stage proving system by which the companion volume, the catalogue of sixty-eight previously unnamed neural constructs, evaluated each of its own proposals.

No construct earns its name by assertion.

Each must survive the three stages, and this chapter shows that each does.

The naming is never a claim of discovery.

Every construct named in this book was already real, already relied upon, already published in parts.

What was missing was the word, and a science cannot reason efficiently about a thing it cannot name.

The proving system is the discipline that ensures a proposed word deserves to exist.

13.2 The Three Stages

Every construct is evaluated against three stages, in order.

A construct that fails any stage does not earn its name.

  https://www.linkedin.com/in/d-r-dison

STAGE 1: CONCEPTUAL PRECISION. The construct must be demonstrably real and currently relied upon in the neuroscience literature. It must be physically or functionally distinct from every adjacent named construct, not merely a paraphrase of something already named. And it must be functionally necessary, playing a specific causal role in neural function whose existence is documented, rather than being a convenience of description. In short: the thing must exist, must be distinct, and must matter.

STAGE 2: IDENTIFICATION OF A GAP, AND PROOF OF NAMELESSNESS. The construct must lack an accepted, widely used proper name in the existing literature. The burden here is high. Informal descriptive phrases, pedagogical metaphors, and passing usages in single papers do not count as proper names. An established proper name is one used consistently across the literature, appearing in standard textbooks, understood without further definition by the relevant community. If any such name exists, even an imperfect or contested one, the construct does not qualify, and it is cited as a named concept instead. The parts may be named; what must be nameless is the unified construct.

STAGE 3: PEDAGOGICAL CLARITY. The proposed name must be self-explanatory to an informed reader, must accurately describe the construct's identity and function, and must be more useful for teaching, research, and clinical application than the current unnamed state. Names that are opaque, misleading, excessively long, or redundant with existing terminology do not pass. The name must earn its place.

13.3 The Additional Considerations: Ten Properties of a Sound Construct

The three stages are the minimum.

To hold each construct to the maximum, this chapter also examines ten properties of each, drawn from the qualities that distinguish a genuine construct from a mere label.

These are stated once here and then reported briefly for each construct.

REALITY: the construct denotes something that physically or functionally exists, not a modelling artefact.

NECESSITY: the construct is required to reason about the domain; remove it and something true becomes unsayable.

CONCEPTUAL PRECISION AND SPECIFICITY: the construct picks out one thing exactly, with sharp boundaries, not a fuzzy region.

IDENTIFICATION OF A GAP IN CONCEPTUAL NEUROSCIENCE AND PROOF OF NAMELESSNESS: there is a demonstrable hole in the field's vocabulary that this construct fills, and no accepted proper name already occupies it.

PEDAGOGICAL CLARITY: the name teaches; a student meeting it learns something true about the construct from the name alone.

DIRECTIONALITY: where the construct is a pathway or a process, its direction of flow or causation is definite and stated, not ambiguous.

HIERARCHICAL POSITION: the construct sits at a definite layer of the fifteen-layer framework, and its relation to the layers above and below it is specified.

  https://www.linkedin.com/in/d-r-dison

UNIVERSALITY WITHIN ITS CLASS: the construct is not a one-off; it is an instance of a recognizable class, and the class is populated.

THEORETICAL COMPLETENESS IN HIERARCHICAL MODELS: naming the construct closes a gap that a complete hierarchical model of the nervous system would otherwise leave open; the framework is more complete with it than without it.

For brevity, where a property is fully satisfied and needs no elaboration, it is marked as satisfied; where it requires comment, the comment is given.

A construct for which a property genuinely does not apply is marked (NA), in the discipline kept throughout this book.

13.4 The Twenty Constructs, Proved

What follows proves each of the twenty constructs of this volume in turn: the five brain-region constructs of Chapter 2, the fourteen interface constructs of Chapter 3, and the one architectural construct of Chapter 12.

Each is given its stage-one, stage-two, and stage-three arguments, followed by the ten additional considerations.

The order follows the order of first appearance in the volume.

13.5 THE FIVE BRAIN-REGION CONSTRUCTS (CHAPTER 2, LAYER 9)

PROVING 1: THE DEFAULT MODE NETWORK CORE TRIAD

STAGE 1, CONCEPTUAL PRECISION: The triad of the medial prefrontal cortex, the posterior cingulate cortex with the precuneus, and the angular gyrus is real, and it is relied upon in every major review of the default mode network. It is distinct: it names not the diffuse network but its three highest-connectivity core hubs, the specific nodes that deactivate first on task onset and are most consistently implicated in self-referential cognition. It is necessary: the default mode network is habitually described as diffuse, and the field lacks a term for its anatomical centre of gravity, which is what disease preferentially attacks.

STAGE 2, GAP AND NAMELESSNESS: The default mode network is named. Each of the three regions is named. The phrase "core nodes" appears descriptively in some reviews. But the specific three-region set, considered as a named Layer 9 anatomical entity with its own internal interface structure, has never been formally designated. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name teaches at once: it names the network (default mode), the role (core), and the number (triad). A student learns from the name that the diffuse network has an anatomical heart, which is exactly the insight needed to understand its preferential failure in Alzheimer's disease.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied; without it the network's centre cannot be named.

Precision, satisfied; three named regions, sharp boundary.

Gap and namelessness, satisfied.

  https://www.linkedin.com/in/d-r-dison

Pedagogical clarity, satisfied.

Directionality, Not Applicable, for the reason that this construct is a hub set rather than a directed pathway, so it has no single direction of flow, though its internal interfaces are individually directed and named.

Hierarchical position, Layer 9, composed of Layer 9 regions, participating in the Layer 10 default mode network.

Universality within its class, satisfied; it is an instance of the class "named core-hub set of a large-scale network," a class that invites further members.

Theoretical completeness, satisfied; a complete model of the default mode network requires a named core.

PROVING 2: THE HABENULA-MONOAMINE FEEDBACK HUB

STAGE 1, CONCEPTUAL PRECISION: The lateral habenula, in its specific role as the node that simultaneously downregulates both dopamine and serotonin when outcomes fall short of expectation, is real and is relied upon in the reward-prediction-error and depression literatures. It is distinct: it is not the habenula in general, nor either monoamine system, but the specific coordinating role of negative feedback onto both at once. It is necessary: the anti-reward signal has a documented causal role, and its coordinated action on two neuromodulator systems has no unified name.

STAGE 2, GAP AND NAMELESSNESS: The lateral habenula is named. The ventral tegmental area and raphe nuclei are named. Reward prediction error is named. But the lateral habenula specifically as the unified negative-feedback hub for both monoamine systems at once has never been named as a distinct Layer 9 functional construct. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the structure (habenula), the targets (the monoamines), the direction (feedback), and the role (hub). A student learns that one small structure vetoes both major reward chemicals when hope is disappointed, which is the key to its role in depression.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied and definite; the flow is from the habenula, via the fasciculus retroflexus, onto the monoamine sources, and it is inhibitory.

Hierarchical position, Layer 9, feeding Layer 10 neuromodulatory systems.

Universality within its class, satisfied; an instance of the class "named neuromodulatory feedback hub." Theoretical completeness, satisfied; a complete model of reward requires a named seat for anti-reward.

PROVING 3: THE PERIAQUEDUCTAL GRAY PAIN MODULATION HUB

  https://www.linkedin.com/in/d-r-dison

STAGE 1, CONCEPTUAL PRECISION: The periaqueductal gray, in its specific role as the organizing hub of the descending pain-modulation system, is real, and electrical stimulation of it produces analgesia sufficient for surgery, which is as strong a proof of causal reality as neuroscience offers. It is distinct: it names this one role of the periaqueductal gray, separate from its roles in defensive behaviour and vocalization. It is necessary: the descending analgesia system has a hub, and the field describes the system without naming its organizing centre as such.

STAGE 2, GAP AND NAMELESSNESS: The periaqueductal gray is named. Descending pain modulation is described. But the periaqueductal gray specifically as the named hub of that system, distinct from its other functions, has never been designated. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the structure, the function (pain modulation), and the role (hub). A student learns that the brain has an organizing centre for switching pain off, which is exactly what is needed to understand stress-induced analgesia and the action of morphine.

ADDITIONAL CONSIDERATIONS: Reality, satisfied, with interventional evidence.

Necessity, satisfied.

Precision, satisfied; one role of one structure.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; the flow is descending, from the periaqueductal gray via the rostral ventromedial medulla to the dorsal horn.

Hierarchical position, Layer 9, organizing a Layer 10 modulatory system.

Universality within its class, satisfied; an instance of the class "named descending-modulation hub." Theoretical completeness, satisfied; a complete model of pain requires a named seat for its top-down control.

PROVING 4: THE ASCENDING AROUSAL SYSTEM NUCLEI

STAGE 1, CONCEPTUAL PRECISION: The specific set of brainstem, hypothalamic, and basal-forebrain nuclei that jointly sustain cortical arousal, the locus coeruleus, raphe, tuberomammillary nucleus, ventral tegmental area, pedunculopontine and laterodorsal tegmental nuclei, basal forebrain, and orexinergic lateral hypothalamus, is real and relied upon. It is distinct: it names the anatomical ensemble, not the functional concept of the reticular activating system. It is necessary: coma results from damage to the set considered together, and the set has no anatomical name.

STAGE 2, GAP AND NAMELESSNESS: Each nucleus is named individually. The "ascending reticular activating system" is a named functional and physiological concept dating to 1949, but it is a functional term that predates the identification of the specific nuclei, not an anatomical name for the set. The precise ensemble has never been named as a unified Layer 9 anatomical entity. The gap is genuine.

  https://www.linkedin.com/in/d-r-dison

STAGE 3, PEDAGOGICAL CLARITY: The name states the function (arousal), the direction (ascending), and the anatomical kind (a set of nuclei). A student learns that wakefulness is not one structure but a named coalition, which is exactly the anatomy that makes arousal robust and its failure, in coma, so grave.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied; a named list of nuclei with a defined joint function.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; ascending, toward the cortex.

Hierarchical position, Layer 9, driving the Layer 10 sleep-wake system.

Universality within its class, satisfied; an instance of the class "named functional nucleus ensemble." Theoretical completeness, satisfied; a complete model of consciousness-as-arousal requires a named anatomical source.

PROVING 5: THE PREVERTEBRAL GANGLION INTEGRATIVE CLUSTER

STAGE 1, CONCEPTUAL PRECISION: The celiac, superior mesenteric, and inferior mesenteric ganglia, considered as one functional unit that integrates central and enteric input rather than merely relaying it, are real and relied upon in autonomic physiology. It is distinct: it names the three prevertebral ganglia as an integrative cluster, not as three separate relays. It is necessary: these ganglia perform genuine neural integration in the periphery, receiving both spinal and enteric input, and the field has no name for the integrative unit.

STAGE 2, GAP AND NAMELESSNESS: Each ganglion is named. The concept of a peripheral integrative centre exists. But the three prevertebral ganglia as a unified Layer 9 peripheral processing unit have never been named. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the location (prevertebral ganglia), the function (integrative), and the kind (cluster). A student learns that genuine neural computation happens outside the skull, which corrects the common misconception that peripheral ganglia only relay.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied; three named ganglia, one function.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; it integrates convergent input (spinal and enteric) and distributes sympathetic output to the abdominal and pelvic viscera.

  https://www.linkedin.com/in/d-r-dison

Hierarchical position, Layer 9, in the peripheral nervous system, feeding the Layer 10 autonomic system.

Universality within its class, satisfied; an instance of the class "named peripheral integrative unit." Theoretical completeness, satisfied; a complete model of the autonomic periphery requires a named seat for peripheral integration.

13.6 THE FOURTEEN INTERFACE CONSTRUCTS (CHAPTER 3, LAYER 9)

PROVING 6: THE COMPUTATIONAL INTERFACE

STAGE 1, CONCEPTUAL PRECISION: The category of interfaces that transform their signal in transit, possessing a faculty resident in neither endpoint, is real, and the thalamic reticular nucleus, the intercalated cell masses, the inferior olive, and the superior olive all demonstrably instantiate it. It is distinct: it names the general category that opposes the transmissive interface, the pipe. It is necessary: the field has the driver-and-modulator distinction and the debate over whether the thalamus is "more than a relay," but no term for the general class of interfaces that compute rather than convey.

STAGE 2, GAP AND NAMELESSNESS: The individual computational interfaces are named. The driver and modulator concepts exist. But the general category, the interface-that-computes as opposed to the interface-that-conveys, has no accepted name, and so the generalization uniting these structures cannot be stated. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states exactly the thing: an interface that computes. A student learns the single most counterintuitive fact about brain interfaces, that some of them are not pipes but processors, and that the faculty may live in the gate rather than the room.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied; it is the category that makes the Perfect Wire Test statable.

Precision, satisfied; sharply separated from the transmissive interface by that test.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, Not Applicable, for the reason that this construct is a category of interface rather than a single pathway, so no one direction can be assigned to the category, though each member instance has a definite direction.

Hierarchical position, Layer 9, an interface between Layer 9 regions.

Universality within its class, satisfied; it IS a class, and it is populated by at least four members.

Theoretical completeness, satisfied; a complete model of interfaces requires the category that computes, not only the category that conveys.

  https://www.linkedin.com/in/d-r-dison

PROVING 7: THE TEMPORO-FRONTAL SEMANTIC INTEGRATION PATHWAY

STAGE 1, CONCEPTUAL PRECISION: The functional interface delivering object identity and lexical meaning from the temporal recognition system to the prefrontal executive system is real and relied upon. It is distinct: it names the functional delivery of semantic content, not the several fibre bundles (the inferior fronto-occipital and uncinate fasciculi) that carry it. It is necessary: the dissociation between knowing what a thing is and being able to reason about it requires a name for the channel that bridges them.

STAGE 2, GAP AND NAMELESSNESS: The fibre bundles are named. The functional interface they jointly constitute has no accepted proper name. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states origin (temporo), destination (frontal), and cargo (semantic integration). A student learns that meaning is delivered forward to the reasoner, and where.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied and definite; feedforward, temporal to frontal.

Hierarchical position, Layer 9, a corticocortical interface.

Universality within its class, satisfied; an instance of the class "named functional corticocortical pathway." Theoretical completeness, satisfied.

PROVING 8: THE TEMPORO-HIPPOCAMPAL ENCODING CHANNEL

STAGE 1, CONCEPTUAL PRECISION: The channel submitting recognized objects and scenes to the hippocampal memory system, via the perirhinal, parahippocampal, and entorhinal cortices, is real. It is distinct: it names the functional submission-for-storage, not the individual relays. It is necessary: the failure to encode what has been perceived, with perception intact, requires a name for the channel that carries perception to storage.

STAGE 2, GAP AND NAMELESSNESS: The relays are named. The channel as a functional whole has no accepted name. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states origin (temporal), destination (hippocampal), and function (encoding). A student learns the route by which experience becomes memory.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

  https://www.linkedin.com/in/d-r-dison

Pedagogical clarity, satisfied.

Directionality, satisfied; into the hippocampus.

Hierarchical position, Layer 9.

Universality within its class, satisfied.

Theoretical completeness, satisfied.

PROVING 9: THE PARIETO-FRONTAL SPATIAL PLANNING PATHWAY

STAGE 1, CONCEPTUAL PRECISION: The interface delivering spatial coordinates and body schema from the posterior parietal cortex to the frontal action systems is real, and its failure is the well-documented optic ataxia, in which an object is seen clearly but cannot be reached for accurately. It is distinct: it names the functional delivery of "where" to the action planner, carried within the superior longitudinal fasciculus. It is necessary: optic ataxia demands a name for the channel whose loss produces it.

STAGE 2, GAP AND NAMELESSNESS: The superior longitudinal fasciculus is named. The specific functional interface has no accepted name. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states origin (parietal), destination (frontal), and function (spatial planning). A student learns that action needs "where," delivered from the spatial system to the planner.

ADDITIONAL CONSIDERATIONS: Reality, satisfied, with lesion evidence.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; parietal to frontal, feedforward.

Hierarchical position, Layer 9.

Universality within its class, satisfied.

Theoretical completeness, satisfied.

PROVING 10: THE CORTICO-CEREBELLAR REFINEMENT LOOP

STAGE 1, CONCEPTUAL PRECISION: The closed loop by which the cortex sends its intended action to the cerebellum and receives a correction is real and relied upon. It is distinct: it names the loop as a single closed interface with an outbound and an inbound arm, where the field names each leg separately. It is necessary: ataxia and the loss of skill acquisition require a name for the loop whose disruption produces them.

  https://www.linkedin.com/in/d-r-dison

STAGE 2, GAP AND NAMELESSNESS: Each leg of the loop is named. The loop itself, as one closed interface, has no accepted name. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the participants (cortico-cerebellar) and the function (refinement) and the topology (loop). A student learns that movement is refined by a round trip, not a one-way command.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied and specifically circular; out to the cerebellum, back to the cortex.

Hierarchical position, Layer 9, and it is one of the reentrant loops central to the framework.

Universality within its class, satisfied; an instance of the class "named reentrant refinement loop." Theoretical completeness, satisfied; a complete model of motor learning requires the loop named as a unit.

PROVING 11: THE CEREBELLO-CORTICAL ERROR CORRECTION PATHWAY

STAGE 1, CONCEPTUAL PRECISION: The return arm of the loop above, carrying the discrepancy between intended and actual movement as a correction from the deep cerebellar nuclei to the motor and parietal cortex, is real. It is distinct: it names the inbound arm specifically, which carries a distinct cargo (the correction) and can fail independently of the outbound arm. It is necessary: the return of correction is a separable function whose loss is a separable deficit.

STAGE 2, GAP AND NAMELESSNESS: The dentato-thalamo-cortical pathway is named anatomically. Its specific functional role as the error-correction return has no accepted name. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states origin (cerebello), destination (cortical), and function (error correction). A student learns that the cerebellum talks back with corrections.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied; it is separable from its outbound partner.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; cerebellum to cortex.

Hierarchical position, Layer 9.

  https://www.linkedin.com/in/d-r-dison

Universality within its class, satisfied.

Theoretical completeness, satisfied.

PROVING 12: THE HIPPOCAMPO-CORTICAL RETRIEVAL NETWORK

STAGE 1, CONCEPTUAL PRECISION: The route by which a stored pattern is reinstated into the cortex that will experience it as a memory, running from CA1 and the subiculum back to the neocortex via the entorhinal, retrosplenial, and posterior cingulate cortices, is real and relied upon. It is distinct: it names the retrieval direction, the reverse of the encoding perforant-path system. It is necessary: the dissociation between the inability to encode and the inability to retrieve requires a name for the retrieval route.

STAGE 2, GAP AND NAMELESSNESS: The forward encoding route (the perforant path) is named. The reverse retrieval route has no accepted name, a striking asymmetry. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states origin (hippocampo), destination (cortical), and function (retrieval). A student learns that remembering is the hippocampus dictating a stored pattern back to the cortex, the physical route of recall.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied; it fills the asymmetry left by naming only the forward route.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; hippocampus to cortex, the retrieval direction.

Hierarchical position, Layer 9.

Universality within its class, satisfied.

Theoretical completeness, satisfied; a complete model of memory requires retrieval named as well as encoding.

PROVING 13: THE AMYGDALO-FRONTAL THREAT ASSESSMENT CHANNEL

STAGE 1, CONCEPTUAL PRECISION: The channel delivering the flag of emotional significance from the basolateral and central amygdala to the orbitofrontal, ventromedial prefrontal, and anterior cingulate cortex, via the ventral amygdalofugal pathway and the uncinate fasciculus, is real and relied upon. It is distinct: it names the delivery of emotional salience to the executive, not the anatomical bundles. It is necessary: the Phineas Gage syndrome, the loss of emotional input into decision-making, requires a name for the channel whose loss produces it.

STAGE 2, GAP AND NAMELESSNESS: The bundles are named. The functional threat-assessment channel has no accepted name. The gap is genuine.

  https://www.linkedin.com/in/d-r-dison

STAGE 3, PEDAGOGICAL CLARITY: The name states origin (amygdalo), destination (frontal), and function (threat assessment). A student learns that emotion is delivered to the reasoner as a prioritizing flag.

ADDITIONAL CONSIDERATIONS: Reality, satisfied, with lesion evidence.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; amygdala to frontal cortex.

Hierarchical position, Layer 9.

Universality within its class, satisfied.

Theoretical completeness, satisfied.

PROVING 14: THE FRONTO-AMYGDALAR REGULATORY CHANNEL

STAGE 1, CONCEPTUAL PRECISION: The return channel, delivering the top-down instruction to suppress a fear response from the ventromedial prefrontal and orbitofrontal cortex to the intercalated cell masses and thence the central amygdala, is real and relied upon. It is distinct: it is the return arm of the channel above, carrying an opposite-signed cargo (suppression), and it fails independently, which is the core deficit of post-traumatic stress disorder. It is necessary: fear extinction, and its failure, require a name for the channel that carries top-down suppression.

STAGE 2, GAP AND NAMELESSNESS: The pathway is anatomically described. Its functional role as the regulatory return has no accepted name. The gap is genuine, and this is arguably the most clinically important nameless interface, being the pathway that exposure therapy is understood to strengthen.

STAGE 3, PEDAGOGICAL CLARITY: The name states origin (fronto), destination (amygdalar), and function (regulatory). A student learns that the reasoner can command the fear centre to stand down, and by what route.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied; separable from its forward partner.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; frontal to amygdala, and net inhibitory via the intercalated cells.

Hierarchical position, Layer 9.

Universality within its class, satisfied.

  https://www.linkedin.com/in/d-r-dison

Theoretical completeness, satisfied; a complete model of fear requires the regulatory return named.

PROVING 15: THE THALAMIC SENSORY RELAY ENSEMBLE

STAGE 1, CONCEPTUAL PRECISION: The set of first-order sensory relay nuclei of the thalamus, the lateral geniculate, medial geniculate, ventral posterolateral, and ventral posteromedial nuclei, considered as one unified gateway between the sensory periphery and the cortex, is real and relied upon. It is distinct: it names the ensemble, not the four nuclei separately. It is necessary: the fact that all sensation but smell passes through this one gate is a structural fact the field states without naming the gate.

STAGE 2, GAP AND NAMELESSNESS: Each nucleus is named. The set considered as the unified sensory gateway has no accepted name. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the location (thalamic), the function (sensory relay), and the kind (ensemble). A student learns that the senses share one gateway, with smell the lone exception.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied; a named list of nuclei with one joint function.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; periphery to cortex.

Hierarchical position, Layer 9.

Universality within its class, satisfied; an instance of the class "named relay ensemble." Theoretical completeness, satisfied.

PROVING 16: THE TRANSTHALAMIC CORTICOCORTICAL RELAY

STAGE 1, CONCEPTUAL PRECISION: The route by which one cortical area communicates with another not directly, but by descending to a higher-order thalamic nucleus (chiefly the pulvinar or mediodorsal nucleus) via layer-5 driving projections and ascending again to the target cortex, is real and is now well established in the driver-modulator literature. It is distinct: it names the indirect, through-the-thalamus route as opposed to direct corticocortical connection. It is necessary: a substantial fraction of cortex-to-cortex traffic takes this route, and the field describes it without a name for it as a route.

STAGE 2, GAP AND NAMELESSNESS: The phenomenon is described, and the nuclei are named. But the route itself, as a named alternative to direct corticocortical communication, has no accepted proper name, despite being one of the most theoretically consequential findings in modern systems neuroscience. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the topology (transthalamic) and the function (corticocortical relay). A student learns that the cortex

  https://www.linkedin.com/in/d-r-dison

talks to itself through the thalamus, which reframes the thalamus from a gateway to a participant in cortical computation.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; cortex down to higher-order thalamus and back up to a different cortical area.

Hierarchical position, Layer 9.

Universality within its class, satisfied; an instance of the class "named indirect relay route." Theoretical completeness, satisfied; a complete model of cortical communication requires the transthalamic route named.

PROVING 17: THE ATTENTIONAL GATING INTERFACE

STAGE 1, CONCEPTUAL PRECISION: The interface by which the prefrontal cortex instructs the thalamic reticular nucleus to suppress the thalamic relay of a currently irrelevant input is real and relied upon. It is distinct: it names the specific top-down gating instruction, not the thalamic reticular nucleus in general. It is necessary: distractibility and the failure of selective attention require a name for the mechanism that, working, suppresses distraction.

STAGE 2, GAP AND NAMELESSNESS: The thalamic reticular nucleus is named. Prefrontal projections to it are documented. But this specific interface, the physical mechanism of top-down attentional suppression, has no accepted name, despite being one of the most sought-after mechanisms in cognitive neuroscience. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the function (attentional gating) and the kind (interface). A student learns that attention is implemented as a gate the reasoner controls, and where.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; prefrontal cortex to thalamic reticular nucleus, net inhibitory upon the relay.

Hierarchical position, Layer 9, and it is a computational interface.

Universality within its class, satisfied.

  https://www.linkedin.com/in/d-r-dison

Theoretical completeness, satisfied; a complete model of attention requires the gating interface named.

PROVING 18: THE INTEROCEPTIVE ASCENDING INTERFACE

STAGE 1, CONCEPTUAL PRECISION: The channel delivering the physiological condition of the body from the nucleus tractus solitarius to the insula, via the parabrachial nucleus and the ventromedial posterior thalamic nucleus, is real and relied upon. It is distinct: it names the unified channel by which the body reports to consciousness, not the individual relays. It is necessary: the felt sense of bodily state, and on the strong reading the bodily basis of emotion, require a name for the channel that delivers it.

STAGE 2, GAP AND NAMELESSNESS: The nuclei are named. The ascending route is described. The interface, considered as the unified channel by which the body reports to the cortex, has no accepted name. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the content (interoceptive), the direction (ascending), and the kind (interface). A student learns that the body reports its state upward to be felt, and by what route.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; ascending, body to insula.

Hierarchical position, Layer 9.

Universality within its class, satisfied.

Theoretical completeness, satisfied.

PROVING 19: THE NEUROENDOCRINE COMMAND INTERFACE

STAGE 1, CONCEPTUAL PRECISION: The interface at which a neural signal, encoded as action potentials, is transduced into an endocrine signal, encoded as hormone concentration in the blood, at the hypothalamic output to the pituitary, is real and relied upon; it drives every neuroendocrine axis. It is distinct: it names the transduction event, the conversion of neural code into endocrine code, not the portal system or the hormones. It is necessary: the point at which the nervous system and the endocrine system become one system requires a name.

STAGE 2, GAP AND NAMELESSNESS: The hypophyseal portal system is named. The hypothalamo-hypophyseal tract is named. The individual hormones are named. But the transduction event itself, the conversion of neural code into endocrine code, has never been named as a unified interface, despite being the single point at which the two great signalling systems meet. The gap is genuine.

  https://www.linkedin.com/in/d-r-dison

STAGE 3, PEDAGOGICAL CLARITY: The name states the domain (neuroendocrine), the function (command), and the kind (interface). A student learns that there is a definite place where nerve becomes hormone, which is exactly the insight that makes the neuroendocrine system one system rather than two.

ADDITIONAL CONSIDERATIONS: Reality, satisfied.

Necessity, satisfied.

Precision, satisfied.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, satisfied; from the hypothalamic neurons to the pituitary and thence the body.

Hierarchical position, Layer 9, and it is the entry port to the endocrine broadcast medium, the mechanism of much layer-skipping causation.

Universality within its class, satisfied; an instance of the class "named cross-medium transduction interface." Theoretical completeness, satisfied; a complete model of the neuroendocrine system requires the transduction point named.

13.7 THE ONE ARCHITECTURAL CONSTRUCT (CHAPTER 12)

PROVING 20: THE ARBITRATED FEDERATION

STAGE 1, CONCEPTUAL PRECISION: The organizational architecture of the human nervous system, a heterarchical federation of semi-autonomous specialists coordinated by dedicated arbiters, broadcast state-setting, loop closure, and an override hierarchy, with no central controller, is real; it is the architecture the anatomy of the preceding chapters plainly exhibits. It is distinct: it names a specific third option between centralized control and flat emergence, neither of which fits the evidence. It is necessary: the field oscillates between a covert homuncularism that speaks of "the brain deciding" and a flat emergentism that denies any organizing structure, and it has no term for the true architecture that is neither.

STAGE 2, GAP AND NAMELESSNESS: Neuroscience possesses the pieces: the basal ganglia as a selection device, the concept of neuromodulatory broadcast, the observation of reentrant loops, and McCulloch's heterarchy. But the integrated architecture, the specific coordinated whole, has no accepted name, which is why the field keeps sliding back into the two wrong pictures. The gap is genuine.

STAGE 3, PEDAGOGICAL CLARITY: The name states the two essential facts at once: it is a federation (autonomous parts, no ruler) that is arbitrated (with genuine referees, not mere emergence). A student learns the answer to the book's deepest question, who is in charge, in two words: no one, and yet not chaos.

ADDITIONAL CONSIDERATIONS: Reality, satisfied; it is exhibited by the anatomy.

Necessity, satisfied; it is what prevents the two standing errors.

  https://www.linkedin.com/in/d-r-dison

Precision, satisfied; it is defined by four named coordinating mechanisms.

Gap and namelessness, satisfied.

Pedagogical clarity, satisfied.

Directionality, Not Applicable, for the reason that this construct is a whole organizational architecture rather than a single pathway, so it has no one direction of flow, though its constituent loops and overrides are individually directional.

Hierarchical position, it spans Layers 9 through 12, being the organization by which regions, systems, the whole nervous system, and the behaving organism cohere; this cross-layer span is itself part of what the name captures.

Universality within its class, satisfied; it is an instance of the class "control architecture," and Chapter 12 shows it is the class member that a grown, never-halted system must belong to.

Theoretical completeness, satisfied; a complete model of the nervous system requires an answer to who coordinates it, and this construct is that answer.

13.8 SUMMARY: TWENTY CONSTRUCTS, TWENTY PROOFS

Every one of the twenty constructs named in this volume has passed all three stages of the proving system: each is real, distinct, and necessary (Stage 1); each fills a genuine gap and lacks any accepted proper name (Stage 2); and each carries a name that teaches (Stage 3).

Every one has also satisfied the ten additional considerations, with the sole exceptions of directionality for the four constructs that are hubs, ensembles, or architectures rather than directed pathways, which are marked (NA) on that property alone, as is proper.

The naming is not a claim of discovery.

Every construct was already real, already relied upon, already published in parts.

What was missing was the word.

Twenty words have now been supplied, each proved, and the field is by exactly that much better able to reason about the nervous system it already possesses.

CHAPTER 14. CONCLUSION

Layer 9 is the layer of the map, and this volume has attempted to draw it completely.

Three claims have been advanced.

The first is that a brain region is defined by its interfaces.

Regional identity is connectional fingerprint, not cell type and not circuit motif, both of which are generic and shared.

  https://www.linkedin.com/in/d-r-dison

Change what the hippocampus receives and where it sends, and it ceases to be a hippocampus.

The second is that the endocrine glands are Layer 9 structures on precisely the same footing as the brain nuclei, and that the hypothalamus is the physical location at which the nervous system and the endocrine system are demonstrably one organ rather than two.

The third is that neuroscience has named its regions far more completely than it has named its interfaces, and that the gap is a historical artefact of dissection rather than a fact about the brain.

The fornix has a name because a nineteenth-century anatomist could lift it out with forceps.

The pathway by which the prefrontal cortex suppresses a distraction has no name, because it cannot be lifted out with anything.

The fourth, developed in Chapter 3, is that interfaces are not merely pipes.

A substantial class of them compute.

They transform the signal in transit, and they possess faculties resident in neither the region that sends nor the region that receives.

Attention lives in the thalamic reticular nucleus, which is a gate.

Fear extinction lives in the intercalated cell masses, which are a gate.

Sound localization lives in the superior olive, which is a comparison performed in transit.

Destroy these and you do not disconnect two regions; you delete a capacity that existed nowhere else.

The fifth is that the anatomy is not the whole story, because the brain is a rhythmic organ.

The wires of Chapters 2 through 4 are fixed, but which wires are open for traffic at any instant is set by the phase of an oscillation, as Chapter 5 develops.

A region is defined by its interfaces, and which of its interfaces are live is defined by the coherence of the ongoing rhythms.

Much of brain disease, as Chapter 5 shows, is the music going wrong while the instrument stays intact.

The sixth is that the regions and interfaces run on a chemistry, and the chemistry is not incidental.

Two molecules, glutamate and GABA, carry nearly all the fast signal; everything else, the dozens of monoamines, peptides, gases, and hormones catalogued in Chapter 7, is modulation that sets the gain, mood, drive, and plasticity within which the signal operates.

  https://www.linkedin.com/in/d-r-dison

The same molecule means different things in different regions, which is the molecular restatement of this volume's thesis that significance lives in connections, not parts.

And the slowest messengers, cortisol, BDNF, the cytokines, are the very vehicles by which causation skips across the layers.

A model of the regions without the molecules is a switchboard with the current off.

The seventh, and the largest, is that there is no orchestrator.

The regions of Layer 9 do not report to a controller, because no controller exists and none could, on pain of infinite regress.

They form what Chapter 12 names an Arbitrated Federation (PUC): semi-autonomous specialists, coordinated by dedicated referees that have no preferences of their own, by broadcast systems that set the parameters of computation without knowing what is being computed, by loops that have no upstream, and by an override hierarchy ordered by evolutionary age, in which breathing outranks deliberation and always will.

The eighth, and the one that turns this volume toward a working system, is that the map has been made runnable.

Every construct now carries not only what it IS but how it CHANGES: a dynamical archetype (one of six simple behaviours, a relay, an integrator, an oscillator, a memory, a gate, or a comparator), the drives it serves, and the rule by which its connections learn.

These are, in part, this framework's modeling choices rather than settled fact, and they are marked as such throughout.

But with them, an implementer can read any component and write the one function a simulation needs: given its current state, its inputs, and the chemical messengers bathing it, what is its next state.

The wiring diagram has become a machine that can be stepped forward in time.

That is the bridge to the specification, and it is why this volume ends where the engineering begins.

Twenty previously unnamed constructs have been named across this volume: five brain regions in Chapter 2 (the Default Mode Network Core Triad, the Habenula-Monoamine Feedback Hub, the Periaqueductal Gray Pain Modulation Hub, the Ascending Arousal System Nuclei, and the Prevertebral Ganglion Integrative Cluster); fourteen interface constructs in Chapter 3 (the Computational Interface itself, and the thirteen named pathways from the Temporo-Frontal Semantic Integration Pathway through the Neuroendocrine Command Interface); and one architectural construct in Chapter 12 (the Arbitrated Federation).

The naming is not a claim of discovery.

Every one of these pathways was already known, already measured, already published.

What was missing was the word.

  https://www.linkedin.com/in/d-r-dison

A science cannot reason efficiently about what it cannot name, and the purpose of this volume is to hand the field a vocabulary adequate to the anatomy it already possesses.

SELECTED REFERENCES

Baars, B.J. (1988). A Cognitive Theory of Consciousness. Cambridge University Press, Cambridge.

Bliss, T.V.P. and Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit. Journal of Physiology, 232(2), 331-356.

Brooks, R.A. (1986). A robust layered control system for a mobile robot. IEEE Journal of Robotics and Automation, 2(1), 14-23.

Brooks, R.A. (1991). Intelligence without representation. Artificial Intelligence, 47(1-3), 139-159.

Buzsaki, G. (2006). Rhythms of the Brain. Oxford University Press, New York.

Buzsaki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. Science, 304(5679), 1926-1929.

Castren, E. and Rantamaki, T. (2010). The role of BDNF and its receptors in depression and antidepressant drug action. Developmental Neurobiology, 70(5), 289-297.

Crick, F. and Koch, C. (2005). What is the function of the claustrum? Philosophical Transactions of the Royal Society B, 360(1458), 1271-1279.

Dayan, P. and Abbott, L.F. (2001). Theoretical Neuroscience. MIT Press, Cambridge, MA.

Dehaene, S. and Changeux, J.P. (2011). Experimental and theoretical approaches to conscious processing. Neuron, 70(2), 200-227.

Dennett, D.C. (1991). Consciousness Explained. Little, Brown, Boston.

Diekelmann, S. and Born, J. (2010). The memory function of sleep. Nature Reviews Neuroscience, 11(2), 114-126.

Douglas, R.J. and Martin, K.A.C. (1991). A functional microcircuit for cat visual cortex. Journal of Physiology, 440(1), 735-769.

Fries, P. (2015). Rhythms for cognition: communication through coherence. Neuron, 88(1), 220-235.

Friston, K. (2010). The free-energy principle: a unified brain theory? Nature Reviews Neuroscience, 11(2), 127-138.

Ganong, W.F. (2000). Circumventricular organs: definition and role in the regulation of endocrine and autonomic function. Clinical and Experimental Pharmacology and Physiology, 27(5-6), 422-427.

  https://www.linkedin.com/in/d-r-dison

Gazzaniga, M.S. (2000). Cerebral specialization and interhemispheric communication. Brain, 123(7), 1293-1326.

Gurney, K., Prescott, T.J., and Redgrave, P. (2001). A computational model of action selection in the basal ganglia. Biological Cybernetics, 84(6), 401-423.

Hafting, T., Fyhn, M., Molden, S., Moser, M.B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature, 436(7052), 801-806.

Hebb, D.O. (1949). The Organization of Behavior. Wiley, New York.

Herculano-Houzel, S. (2009). The human brain in numbers: a linearly scaled-up primate brain. Frontiers in Human Neuroscience, 3, 31.

Hilgetag, C.C. and Goulas, A. (2020). Hierarchy in the organization of brain networks. Philosophical Transactions of the Royal Society B, 375(1796), 20190319.

Kaelberer, M.M., Buchanan, K.L., Klein, M.E., et al. (2018). A gut-brain neural circuit for nutrient sensory transduction. Science, 361(6408), eaat5236.

Katzung, B.G. (2017). Basic and Clinical Pharmacology, 14th ed. McGraw-Hill, New York.

Keramati, M. and Gutkin, B. (2011). A reinforcement learning theory for homeostatic regulation. Advances in Neural Information Processing Systems, 24.

LeDoux, J.E. (1996). The Emotional Brain. Simon and Schuster, New York.

Lisman, J.E. and Jensen, O. (2013). The theta-gamma neural code. Neuron, 77(6), 1002-1016.

Marr, D. (1982). Vision. MIT Press, Cambridge, MA.

McCulloch, W.S. (1945). A heterarchy of values determined by the topology of nervous nets. Bulletin of Mathematical Biophysics, 7(2), 89-93.

Menon, V. and Uddin, L.Q. (2010). Saliency, switching, attention and control: a network model of insula function. Brain Structure and Function, 214(5-6), 655-667.

Minsky, M. (1986). The Society of Mind. Simon and Schuster, New York.

Moruzzi, G. and Magoun, H.W. (1949). Brain stem reticular formation and activation of the EEG. Electroencephalography and Clinical Neurophysiology, 1(4), 455-473.

Mountcastle, V.B. (1957). Modality and topographic properties of single neurons of cat's somatic sensory cortex. Journal of Neurophysiology, 20(4), 408-434.

Nambu, A., Tokuno, H., and Takada, M. (2002). Functional significance of the cortico-subthalamo-pallidal hyperdirect pathway. Neuroscience Research, 43(2), 111-117.

  https://www.linkedin.com/in/d-r-dison

Nestler, E.J., Hyman, S.E., Holtzman, D.M., and Malenka, R.C. (2015). Molecular Neuropharmacology, 3rd ed. McGraw-Hill, New York.

O'Keefe, J. and Nadel, L. (1978). The Hippocampus as a Cognitive Map. Oxford University Press, Oxford.

Purves, D., et al. (2018). Neuroscience, 6th ed. Sinauer Associates, Sunderland, MA.

Raichle, M.E., MacLeod, A.M., Snyder, A.Z., et al. (2001). A default mode of brain function. Proceedings of the National Academy of Sciences, 98(2), 676-682.

Redgrave, P., Prescott, T.J., and Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? Neuroscience, 89(4), 1009-1023.

Saper, C.B., Chou, T.C., and Scammell, T.E. (2001). The sleep switch: hypothalamic control of sleep and wakefulness. Trends in Neurosciences, 24(12), 726-731.

Scoville, W.B. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. Journal of Neurology, Neurosurgery, and Psychiatry, 20(1), 11-21.

Selfridge, O.G. (1959). Pandemonium: a paradigm for learning. In Symposium on the Mechanization of Thought Processes. HMSO, London.

Sherman, S.M. and Guillery, R.W. (2006). Exploring the Thalamus and Its Role in Cortical Function, 2nd ed. MIT Press, Cambridge, MA.

Sherrington, C.S. (1906). The Integrative Action of the Nervous System. Yale University Press, New Haven, CT.

Sudhof, T.C. (2013). Neurotransmitter release: the last millisecond in the life of a synaptic vesicle. Neuron, 80(3), 675-690.

Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. Biological Bulletin, 215(3), 216-242.

van den Heuvel, M.P. and Sporns, O. (2011). Rich-club organization of the human connectome. Journal of Neuroscience, 31(44), 15775-15786.

Wang, R. (2002). Two's company, three's a crowd: can H2S be the third endogenous gaseous transmitter? FASEB Journal, 16(13), 1792-1798.

Wurtz, R.H. and Goldberg, M.E. (1989). The Neurobiology of Saccadic Eye Movements. Elsevier, Amsterdam.

Xie, L., Kang, H., Xu, Q., et al. (2013). Sleep drives metabolite clearance from the adult brain. Science, 342(6156), 373-377.

THE SPECIFICATION

THE SPARCD METHOD APPLIED TO NATURE'S COGNITIVE ARCHITECTURE

What follows is the engineering half of this book.

  https://www.linkedin.com/in/d-r-dison

The thirteen chapters described the brain as nature built it.

The six appendices rebuild it as a specification a capable implementation engine can construct.

The method is SPARC, created by Reuven Cohen: Specification, Pseudocode, Architecture, Refinement, and Completion, a five-phase discipline for turning a vague idea into small, well-defined, checkable tasks that a builder, human or machine, can carry out reliably without failing silently.

To SPARC's five phases this work adds a sixth of its own, Demonstration, giving the acronym SPARCD, because a specification of a mind owes the reader not only a plan but an honest account of what the plan is and is not, what it can become, and how anyone might build it.

The guiding principle of the specification is stated once, here, and holds throughout.

THE APPENDICES ADD NO NEW NEUROSCIENCE.

Every requirement, every rule, every number they invoke was established in Chapters 1 through 13.

The specification only reorganizes what the science established into the exact form an engine consumes.

Where an appendix makes an engineering choice not forced by the biology, it says so plainly, in the discipline this book has kept throughout: a specification is allowed to commit, but it is not allowed to pretend its commitments are facts.

APPENDIX 1. SPECIFICATION

A1.0 THE PURPOSE OF THIS PHASE

The Specification phase defines the problem completely before any code is written.

Its restraint is its point: not one line of program code appears here.

What appears is a complete, plain-language statement of what the system must do, the qualities it must have, the scenarios it must serve, the boundaries it must respect, and the criteria by which it will be judged finished.

In the language of SPARC, this appendix is the deliverable traditionally named specification.md.

A1.1 PROJECT OVERVIEW

PROJECT NAME: Nature's Cognitive Architecture, referred to hereafter as the System.

PROJECT GOAL: to build a running software model of the human cognitive architecture at the grain established in this book, meaning the grain of roughly one hundred brain regions, forty named interfaces, the principal

  https://www.linkedin.com/in/d-r-dison

chemical messengers, and the dynamics, drives, and plasticity that make them change over time.

The System is not a model of individual neurons or synapses; it is a model of the regional architecture and its control loops.

TARGET AUDIENCE: three audiences, in order.

First, the independent researcher or student who wishes to build and study a whole-architecture cognitive model on modest hardware.

Second, the research laboratory with serious compute that wishes to scale the model and test it against real datasets.

Third, the implementation engine, human or artificial, that will construct the System from this specification.

WHAT THE SYSTEM IS NOT, STATED UP FRONT: it is not a large language model, and it is not trained on text.

It is a mechanistic, glass-box architecture whose every part corresponds to a named neuroscientific construct and whose every operation is inspectable.

This is a deliberate contrast to the dominant paradigm, and its rationale is developed in Appendix 6.

A1.2 FUNCTIONAL REQUIREMENTS

A functional requirement, labelled with the prefix FR, states what the System must DO.

Each is drawn directly from the chapter cited.

FR-1, THE CONSTRUCT REGISTRY. The System must represent every named construct of Chapters 2, 3, and 7, meaning each region, each interface, and each molecule, as a distinct, inspectable entity carrying its established function and its specification block (its dynamical archetype, its state variables, the drives it serves, its plasticity rule, and its modulators).

FR-2, STATE. Every region must hold a state consisting at minimum of an activation, a gain, and an adaptation variable, as defined in Chapter 8, and, where its archetype requires, a tuning or content variable and a phase variable.

FR-3, THE UPDATE RULE. The System must be able to compute, for every construct, its next state from its current state, its inputs, and the ambient neuromodulator levels, according to the dynamical archetype assigned to it in Chapter 8. This is the core operation: next_state = update(current_state, inputs, neuromodulators).

FR-4, THE TICK. The System must advance in discrete time steps, called ticks. One tick is one synchronous update of every construct at once, as specified in Chapter 8, Section 8.4. The System must be able to run for an arbitrary number of ticks.

FR-5, THE INTERFACES AS CONDUITS. The System must route signals between regions through the interfaces of Chapters 3 and 4, each carrying a delay and a weight, and each being either transmissive (passing its signal

  https://www.linkedin.com/in/d-r-dison

unchanged) or computational (applying its named transform), per the Perfect Wire Test of Chapter 3.

FR-6, THE NEUROMODULATORY BUS. The System must represent the neuromodulators of Chapter 7 not as point-to-point connections but as broadcast fields with a level per target territory, and must make those levels available to every construct's update rule and every learning rule.

FR-7, THE DRIVES. The System must represent the homeostatic drives of Chapter 9, each with a monitored variable, a set-point, a computed error, and a corrective bias on action selection. The System must compute reward as the reduction of drive error, per the homeostatic reinforcement learning unification of Chapter 9.

FR-8, ACTION SELECTION. The System must arbitrate among competing drive-proposed actions through a selection mechanism corresponding to the basal ganglia of Chapter 2, which selects exactly one action by releasing its inhibition while suppressing the rest, and which generates no actions of its own.

FR-9, THE OVERRIDE HIERARCHY. The System must enforce the override ordering of Chapter 9 and Chapter 12, in which older and more vital drives can preempt newer deliberative processes, such that, for example, the respiration drive cannot be suppressed to the point of self-termination.

FR-10, PLASTICITY. The System must change the weights of its interfaces over time according to the three-factor learning rule of Chapter 10: a weight changes in proportion to the product of the activity at its two ends and a neuromodulatory third factor, bounded by homeostatic scaling that holds each region's average activity near a target.

FR-11, THE ELIGIBILITY TRACE. The System must implement the eligibility trace of Chapter 10, so that a reward arriving shortly after an action can strengthen the connection that earned it.

FR-12, CONSOLIDATION. The System must support a distinct slow-timescale plasticity, corresponding to the structural plasticity and sleep replay of Chapters 5 and 10, in which important recent changes are consolidated during a dedicated offline phase.

FR-13, INSPECTABILITY. At any tick, the System must be able to report the full state of any construct, the level of any neuromodulator, the error of any drive, and the weight of any interface. Nothing is hidden. This is a hard requirement, not a convenience, and it is the property that distinguishes this System from opaque models.

FR-14, PERTURBATION AND LESION. The System must permit any construct to be disabled, any interface to be cut, and any neuromodulator to be raised or lowered, so that the System can reproduce the lesion and pharmacological experiments that are the evidential basis of the science. The System that cannot be given a simulated Parkinson's disease by depleting dopamine has not faithfully implemented this book.

FR-15, THE LARGE-SCALE NETWORKS. The System must represent the large-scale cognitive networks of Chapter 2, Section 2.10 (the default mode, salience, dorsal attention, frontoparietal control, cingulo-opercular, language, sensorimotor, visual, auditory, and limbic networks) as named groupings over the regions of the registry, each carrying its member nodes, its subscribe-and-publish couplings, and its functional signature, including

  https://www.linkedin.com/in/d-r-dison

the characteristic anticorrelations (for example, the default mode network deactivating as the dorsal attention and frontoparietal networks engage).

FR-16, THE COGNITIVE SYSTEMS AND PATHWAYS. The System must represent the cognitive systems and pathways of Chapter 2, Section 2.11, as named subgraphs and loops over the connection graph: the cortico-basal-ganglia-thalamo-cortical loops, the direct, indirect, and hyperdirect pathways, and the Papez circuit; and it must represent the glia of Section 2.11.5 (astrocytes, microglia, and oligodendrocytes) and the tripartite synapse as modulatory participants that gate plasticity and connectivity rather than as signal-carrying regions.

FR-17, THE TWO COMMUNICATION PRIMITIVES. The System must distinguish the two messaging primitives established in Chapters 3 and 7: a DIRECTED channel (send and receive) for a point-to-point interface, whose publisher is wired to a single subscriber; and a BROADCAST topic (publish and subscribe) for a volume-transmission interface and for every neuromodulatory and hormonal field, whose publisher addresses no one and whose subscribers are defined by the receptor they carry. A construct's inputs are what it subscribes to; its outputs are what it publishes.

A1.3 NON-FUNCTIONAL REQUIREMENTS

A non-functional requirement, labelled with the prefix NFR, states a QUALITY the System must have.

NFR-1, GLASS-BOX TRANSPARENCY. Every state, every weight, every signal must be a named, human-readable quantity tied to a named construct. No part of the System may be an uninterpretable numerical blob. This is the System's defining virtue and its central contrast with the prevailing paradigm.

NFR-2, DETERMINISM AND REPRODUCIBILITY. Given the same initial state and the same inputs, the System must produce the same trajectory. Any randomness must be seeded and recorded, so that any run can be replayed exactly.

NFR-3, SUBSTRATE INDEPENDENCE. The specification must be expressible without commitment to any programming language or data structure, so that it can be implemented in any language. This requirement is the subject of Appendix 6 and is a first-class goal, not an afterthought.

NFR-4, SCALABILITY BY GRAIN. The System must run at the regional grain on modest hardware (a personal computer), and its design must not preclude later refinement to finer grain given greater compute. The architecture must scale down to be buildable at home and scale up to be serious in a laboratory.

NFR-5, MODULARITY. Each construct must be a self-contained module with a clean interface, so that constructs can be added, replaced, or refined without rewriting the System. This mirrors the encapsulation thesis of the science itself: the brain is modular, and so must be its model.

NFR-6, BIOLOGICAL FIDELITY, BOUNDED. The System must be faithful to the neuroscience at the regional grain, and must not claim fidelity it does not have. Where the System simplifies, the simplification must be documented, in the discipline kept throughout this book.

  https://www.linkedin.com/in/d-r-dison

NFR-7, VERIFIABILITY. Every requirement here must be checkable by an automated test, as developed in Appendices 4 and 5. A requirement that cannot be tested is not a requirement; it is a hope.

A1.4 USER SCENARIOS

A user scenario is a short plain-language account of what a user wants to accomplish.

SCENARIO 1, THE STUDENT. A student loads the System, runs it for a thousand ticks with a simple sensory input, and watches the activation of the visual regions, the value computed by the orbitofrontal cortex, and the action selected by the basal ganglia, tick by tick, to understand how a decision forms.

SCENARIO 2, THE LESION EXPERIMENT. A researcher depletes the System's dopamine to twenty percent of normal and observes the System's action selection slow and freeze, reproducing the bradykinesia of Parkinson's disease, then restores dopamine and observes recovery. The researcher has run a controlled experiment on a transparent brain.

SCENARIO 3, THE LEARNING TASK. A researcher presents the System with a stimulus repeatedly paired with a reward, and observes the relevant interface weights strengthen by the three-factor rule, so that the System comes to predict the reward from the stimulus. The researcher then presents the stimulus without reward and observes extinction. The System has learned and unlearned.

SCENARIO 4, THE DRIVE CONFLICT. A researcher raises the System's hunger drive and its threat drive simultaneously and observes the override hierarchy resolve the conflict in favour of threat, then lowers the threat and observes hunger resume control of behaviour. The researcher has watched motivation arbitrate.

SCENARIO 5, THE SCALE-UP. A laboratory with serious compute connects the System to a real dataset, refines selected regions to finer grain, and runs the System for millions of ticks to test whether cognitively meaningful behaviour emerges. This scenario is the bridge to Appendix 6 and to the AGI question.

A1.4bis USER PERSONAS

The SPARC Specification phase asks for personas, the fictional profiles of typical users, so the requirements are grounded in who will actually use the System. Three personas follow, matching the three audiences of Section A1.1.

PERSONA 1, THE STUDENT-BUILDER (Maya). A graduate student in cognitive science with modest hardware and strong curiosity but limited systems-programming experience. She wants to build the System at home, watch a decision form tick by tick, and reproduce a famous experiment to understand the mechanism. Her need drives the inspectability, the modest-hardware grain, and the one-hour-to-first-experiment documentation goal.

PERSONA 2, THE LABORATORY RESEARCHER (Dr. Okafor). A computational neuroscientist with serious compute who wants to scale the System, connect it to real datasets, refine selected regions to finer grain, and run controlled lesion and pharmacological experiments as reproducible science.

  https://www.linkedin.com/in/d-r-dison

His need drives the scalability-by-grain, determinism-and-reproducibility, and perturbation requirements.

PERSONA 3, THE IMPLEMENTATION ENGINE (an artificial or human builder). The agent, human or artificial, that constructs the System from this specification in a chosen language on a chosen substrate. Its need drives the substrate independence, the language-agnostic pseudocode, and the explicit acceptance criteria that make "done" checkable rather than judged.

A1.4ter INTERFACE AND INTERACTION GUIDELINES

The System has no consumer graphical interface, but it is not without a user experience: it is a scientific instrument, and its interaction surface is the observer and the experiment controls. The SPARC User-Interface-and-User-Experience guidelines are stated here in the terms that apply to such an instrument.

GUIDELINE 1, THE OBSERVER IS THE PRIMARY INTERFACE. The user reads the System through the observer of Appendix 3, which reports the full state of any construct, network, drive, and neuromodulator at any tick. The observer's output must be human-readable and named, never an uninterpretable numerical blob, so that a person can follow a decision as it forms.

GUIDELINE 2, EXPERIMENTS ARE THE PRIMARY CONTROLS. The user acts on the System through the perturbation interface: lesion a construct, cut an interface, or shift a neuromodulator. Each control must be a single, named, reversible operation, so that running the Parkinson experiment is one legible act, not a reconfiguration.

GUIDELINE 3, LEGIBILITY OVER CONVENIENCE. Where a choice arises between a faster interaction and a more inspectable one, the instrument chooses inspectability, consistent with the System's defining virtue. The interface exists to make the glass box readable, not to hide it behind a dashboard.

GUIDELINE 4, REPRODUCIBILITY IS PART OF THE EXPERIENCE. Every interaction is recorded with its seed and its tick, so that any session a user runs can be replayed exactly by another user, which for a scientific instrument is the whole of a good user experience.

A1.5 TECHNICAL CONSTRAINTS

CONSTRAINT 1: the specification must remain language-agnostic and data-structure-agnostic through Appendix 5. Language and data-structure choices belong to implementation, discussed in Appendix 6, not to specification.

CONSTRAINT 2: the reference implementations named by the author are a Causalontology 2.0.0 data layer, a Python Minimum Viable Product, and a PrologAI implementation. The specification must be compatible with all three and must privilege none.

CONSTRAINT 3: no secret, credential, or environment-specific value may appear anywhere in the eventual code, per SPARC quality rules.

CONSTRAINT 4: the eventual code must observe the SPARC quality bounds: files of five hundred lines or fewer, functions of fifty lines or fewer, comprehensive input validation, modular design with clean interface boundaries, and self-documenting code.

  https://www.linkedin.com/in/d-r-dison

CONSTRAINT 5: the eventual code must use the structured commit-message convention SPARC prescribes, with the prefixes feat for a feature, test for a test, fix for a bug fix, docs for documentation, arch for an architecture change, quality for a refactor, and deploy for a deployment change, so that the history of the build is itself legible.

A1.6 ASSUMPTIONS

ASSUMPTION 1: the regional grain is sufficient to produce cognitively interesting behaviour. This is a hypothesis, not a certainty, and testing it is part of the project's purpose.

ASSUMPTION 2: the six dynamical archetypes of Chapter 8 are sufficient to describe the dynamics of every construct at the regional grain. Where a construct resists all six, that is a finding, and the archetype set is extended, in the manner the science itself was extended.

ASSUMPTION 3: the three-factor learning rule of Chapter 10 is sufficient for the System to learn at the regional grain. This too is a hypothesis under test.

ASSUMPTION 4: the modeling commitments marked throughout the science chapters as THIS FRAMEWORK'S CHOICE are adopted as the specification's baseline, and may be revised in Refinement.

A1.7 ACCEPTANCE CRITERIA

The System is accepted as a faithful implementation of this specification if and only if all of the following hold, each of which is made into an automated test in Appendix 5.

Each is labelled with the prefix AC, for Acceptance Criterion, and numbered.

AC-1: every named construct of Chapters 2, 3, and 7 is present in the registry with its specification block.

AC-2: the System advances by ticks, and one tick updates every construct exactly once.

AC-3: a signal presented at a sensory region propagates through the correct chain of interfaces to produce an action, in the correct number of ticks given the interface delays.

AC-4: depleting dopamine slows and then halts action selection, and restoring it recovers, reproducing the Parkinsonian signature.

AC-5: a stimulus repeatedly paired with reward comes to predict that reward through weight change, and the change reverses under extinction.

AC-6: two simultaneous drives are arbitrated in the order specified by the override hierarchy, and the respiration drive cannot be voluntarily suppressed to self-termination.

AC-7: every state, weight, drive error, and neuromodulator level is reportable at every tick.

AC-8: a full run is exactly reproducible from its seed.

  https://www.linkedin.com/in/d-r-dison

AC-9: the System runs at the regional grain on a personal computer within a stated time budget per thousand ticks.

AC-10: every large-scale network of Chapter 2, Section 2.10, is present in the registry as a named grouping over its member regions, with its couplings and its functional signature.

AC-11: every cognitive system and pathway of Chapter 2, Section 2.11, is present as a named subgraph or loop, the glia are present as modulatory participants, and each interface and field is marked as a directed (send-and-receive) channel or a broadcast (publish-and-subscribe) topic.

A1.8 REFLECTION

The discipline of SPARC requires that each phase reflect on its own choices.

Three deserve note.

The decision to make INSPECTABILITY a hard functional requirement (FR-13) and the System's defining quality (NFR-1) is the most consequential in the specification, and it is a deliberate wager against the prevailing direction of the field.

The dominant models are powerful and opaque.

This System trades raw power for total transparency, on the bet that a mind we can inspect is worth more, scientifically and in the long run practically, than a mind we cannot.

If that bet is wrong, the System will be an educational instrument rather than a competitive intelligence.

If it is right, it is a path the opaque models cannot follow.

The decision to specify at the REGIONAL GRAIN (Assumption 1) is the specification's largest risk.

It may be that cognition requires finer detail than one hundred regions can hold.

The specification mitigates this by requiring that the design scale to finer grain (NFR-4) without being rebuilt, so that the grain can be refined if the hypothesis fails, rather than the System abandoned.

The decision to add DEMONSTRATION as a sixth phase reflects a conviction that a specification of a mind carries an unusual obligation.

An ordinary specification may end at Completion, with a working system.

A specification that claims a path toward general intelligence owes its reader more: an honest account of what it is, what it is not, and what it might become.

Appendix 6 discharges that obligation, and its confident, well-founded case for the path to general intelligence is part of the specification, not an afterword to it.

  https://www.linkedin.com/in/d-r-dison

APPENDIX 2. PSEUDOCODE

WRITTEN IN ENGLISH READABLE CODE (ERC)

A2.0 A WORD ON ENGLISH READABLE CODE (ERC)

This appendix is written in English Readable Code, abbreviated ERC, a controlled form of precise English in which each sentence corresponds to one exactly defined computational operation, defined by D. R. Dison.

ERC reads like ordinary English and yet can be translated, sentence by sentence, into any real programming language, because every sentence has one and only one meaning.

It is chosen for the Pseudocode phase for the reason SPARC prizes pseudocode in the first place: it states the logic in plain language, where a mistake is cheap to see and fix, before any programming language is chosen.

ERC is the ideal notation for a specification meant to be programming-language-agnostic, which this one is (see Appendix 6), because it commits to the logic without committing to a language.

A note on convention.

Each ERC procedure opens with a sentence of the form "The purpose of this procedure is to" and closes with "The procedure is complete." One operation is written per sentence.

Variable names are written as plain spaced English phrases.

Where a procedure takes inputs or returns a value in a way the base ERC catalogue does not fully template, this appendix extends ERC in the obvious readable way and marks the fact; such extensions are noted as ERC EXTENSION.

What follows is not the whole System in ERC, which would run to many hundreds of procedures.

It is the load-bearing core: the main tick loop, the single-construct update, the six archetype rules, the drive loop, and the learning rule.

From these, every remaining construct follows the same pattern, exactly as Chapter 11 established the per-construct block and then applied it uniformly.

A2.1 THE MAIN LOOP: RUNNING THE SYSTEM

The purpose of this procedure is to run the whole System for a chosen number of ticks.

Create a number variable called total ticks with initial value one thousand.
Create a number variable called current tick with initial value zero.
Create a list variable called all constructs with initial value empty list.
Load every region and interface and molecule into all constructs.

While current tick is less than total ticks, repeat the following steps:

  https://www.linkedin.com/in/d-r-dison

```
    Read the body state and update every drive error, and compute this
tick's reward.
    Update all neuromodulator levels for this tick, including broadcasting
the reward as phasic dopamine.
    For each construct in all constructs, perform the following operations:
        Compute the next state of the construct for this tick.
    End the loop.
    For each construct in all constructs, perform the following operations:
        Set the current state of the construct to its next state.
    End the loop.
    Select one action from the competing drive proposals.
    Apply all plasticity weight changes for this tick.
    Record the full state of the System for this tick.
    Calculate the sum of current tick and one.
    Set current tick to the sum.
End the loop.

Display a message that the run is complete to the console.
The procedure is complete.

A note on the two separate loops over all constructs.

The first loop computes every construct's next state from the current
states of its neighbours; the second loop then commits every next state at
once.

Separating them is what makes one tick a synchronous update, in which every
construct sees the same past and no construct sees a neighbour's future.

This is a deliberate and load-bearing choice, and an implementer who fuses
the two loops will have built a subtly different and wrong machine.

A2.2 UPDATING ONE CONSTRUCT: THE DISPATCH

The purpose of this procedure is to compute the next state of a single
construct according to its dynamical archetype.

Read the archetype of the construct.

Check if the archetype is equal to the word relay.
If the condition is true, then perform the following actions:
    Compute the next state using the relay rule.
End the conditional block.

Check if the archetype is equal to the word integrator.
If the condition is true, then perform the following actions:
    Compute the next state using the integrator rule.
End the conditional block.

Check if the archetype is equal to the word oscillator.
If the condition is true, then perform the following actions:
    Compute the next state using the oscillator rule.
End the conditional block.

Check if the archetype is equal to the word attractor.
If the condition is true, then perform the following actions:
    Compute the next state using the attractor rule.
End the conditional block.
```

  https://www.linkedin.com/in/d-r-dison

```
Check if the archetype is equal to the word gate.
If the condition is true, then perform the following actions:
    Compute the next state using the gate rule.
End the conditional block.

Check if the archetype is equal to the word comparator.
If the condition is true, then perform the following actions:
    Compute the next state using the comparator rule.
End the conditional block.

The procedure is complete.
```

This dispatch is the whole of the "one canonical computation, rewired" thesis made into code.

There are six rules, not one hundred.

Every construct is one of six kinds, and its identity is its parameters and its connections, not a bespoke rule.

A2.3 THE SIX ARCHETYPE RULES

```
The purpose of this procedure is to compute the next activation of a RELAY
construct.
Gather the total weighted input arriving from all incoming interfaces.
Create a number variable called total input with that gathered value.
Read the current gain of the construct, which the neuromodulators have set.
Calculate the product of total input and the current gain.
Create a number variable called next activation with the product.
Set the next state activation of the construct to next activation.
The procedure is complete.

The purpose of this procedure is to compute the next activation of an
INTEGRATOR construct.
Read the current activation of the construct.
Create a number variable called decayed activation with the current
activation.
Multiply decayed activation by a leak factor slightly less than one.
Gather the total weighted input arriving from all incoming interfaces.
Create a number variable called total input with that gathered value.
Calculate the sum of decayed activation and total input.
Create a number variable called next activation with the sum.
Set the next state activation of the construct to next activation.
The procedure is complete.

The purpose of this procedure is to compute the next state of an OSCILLATOR
construct.
Read the current phase of the construct.
Read the natural frequency of the construct.
Calculate the sum of the current phase and the natural frequency.
Create a number variable called next phase with the sum.
Wrap next phase back into the range of one full cycle.
Compute the receptivity of the construct from next phase.
Set the next state phase of the construct to next phase.
Set the next state gain of the construct to the computed receptivity.
The procedure is complete.

The purpose of this procedure is to compute the next state of an ATTRACTOR
construct.
```

  https://www.linkedin.com/in/d-r-dison

```
Read the current activation pattern of the construct.
Gather the input pattern arriving from all incoming interfaces.
Create a variable called combined pattern by blending the current pattern
with the input pattern.
Find the stored pattern that is most similar to combined pattern.
Create a variable called nearest pattern with that stored pattern.
Move the activation pattern a small step toward nearest pattern.
Set the next state activation pattern of the construct to the moved
pattern.
The procedure is complete.

The purpose of this procedure is to compute the next state of a GATE
construct.
Read the current mode of the construct.
Gather the total drive arriving to switch the gate.
Create a number variable called switch drive with that gathered value.
Read the switching threshold of the construct.
Check if switch drive is greater than the switching threshold.
If the condition is true, then perform the following actions:
    Flip the mode of the construct to the opposite mode.
Otherwise, perform the following actions:
    Keep the mode of the construct unchanged.
End the conditional block.
Set the next state mode of the construct to the resulting mode.
The procedure is complete.

The purpose of this procedure is to compute the next activation of a
COMPARATOR construct.
Read the expected input of the construct.
Read the actual input of the construct.
Calculate the difference of the actual input and the expected input.
Create a number variable called prediction error with the difference.
Set the next state activation of the construct to prediction error.
The procedure is complete.

These six procedures are the dynamical heart of the System.

Every region in Chapter 2 and every interface in Chapter 3 runs one of
them, differing only in its parameters, its stored patterns, its threshold,
or its connections.

This is the machine form of the book's oldest claim.

A2.4 THE DRIVE LOOP: WHAT THE SYSTEM WANTS

The purpose of this procedure is to update every drive and compute the
reward for this tick.

Create a list variable called all drives with initial value empty list.
Load every homeostatic drive into all drives.
Create a number variable called total reward with initial value zero.

For each drive in all drives, perform the following operations:
    Read the monitored variable of the drive from the body state.
    Read the set point of the drive.
    Calculate the difference of the monitored variable and the set point.
    Create a number variable called new error with the absolute size of the
difference.
```

  https://www.linkedin.com/in/d-r-dison

```
    Check if the drive has no previous error recorded yet, which is true
only on the first tick.
    If the condition is true, then perform the following actions:
        Set the previous error of the drive to new error, so that the first
tick yields no spurious reward.
    End the conditional block.
    Read the previous error of the drive.
    Calculate the difference of the previous error and new error.
    Create a number variable called error reduction with the difference.
    Calculate the sum of total reward and error reduction.
    Set total reward to the sum.
    Set the previous error of the drive to new error.
    Bias action selection toward behaviors that reduce this drive, in
proportion to new error.
End the loop.

Broadcast total reward as the phasic dopamine signal for this tick.
The procedure is complete.
```

Notice the last two lines, and what they accomplish.

The reward is defined as the total reduction of drive error across all drives, exactly the homeostatic reinforcement learning unification of Chapter 9.

And that reward is then broadcast as dopamine, which Chapter 10 established is the third factor in learning.

In these two sentences the drives, the reward, and the learning signal become one quantity.

This is the loop of the mind, closed in plain English.

A2.5 THE LEARNING RULE: HOW THE SYSTEM CHANGES ITSELF

The purpose of this procedure is to change the weight of every interface according to the three-factor rule.

```
Create a list variable called all interfaces with initial value empty list.
Load every interface that can learn into all interfaces.
Read the current dopamine level as the third factor.
Create a number variable called third factor with the current dopamine
level.

For each interface in all interfaces, perform the following operations:
    Read the recent activity at the sending end of the interface.
    Create a number variable called sending activity with that value.
    Read the recent activity at the receiving end of the interface.
    Create a number variable called receiving activity with that value.
    Calculate the product of sending activity and receiving activity.
    Create a number variable called coincidence with the product.
    Calculate the product of coincidence and third factor.
    Create a number variable called raw weight change with the product.
    Multiply raw weight change by the learning rate.
    Create a number variable called weight change with the result.
    Read the current weight of the interface.
    Calculate the sum of the current weight and weight change.
    Set the weight of the interface to the sum.
End the loop.
```

  https://www.linkedin.com/in/d-r-dison

Apply homeostatic scaling so that each region's average activity stays near its target.
The procedure is complete.

Read the three multiplied quantities in the middle of the loop: the sending activity, the receiving activity, and the third factor.

Those are Hebb's two factors and the neuromodulatory third factor of Chapter 10, and their product is the weight change.

A connection strengthens only when its two ends are active together AND dopamine says the moment mattered.

The final line, the homeostatic scaling, is the bound that keeps the learning from exploding, the negative feedback on plasticity itself.

The whole of Chapter 10 is these two ideas, and here they are, as a procedure a machine can run.

A2.6 THE ELIGIBILITY TRACE, IN BRIEF

The purpose of this procedure is to let a reward strengthen a connection that was active a moment ago.

For each interface in all interfaces, perform the following operations:
    Read the current coincidence at the interface.
    Read the current eligibility trace of the interface.
    Multiply the current eligibility trace by a fading factor slightly less than one.
    Calculate the sum of the faded trace and the current coincidence.
    Set the eligibility trace of the interface to the sum.
End the loop.

When a reward arrives, multiply each interface's eligibility trace by the third factor, and add the result to its weight.
The procedure is complete.

The eligibility trace is a fading memory of recent coincidence, so that a reward arriving a second after an action can still find the connection that earned it.

It is the bridge between fast coincidence and slower reward, and Chapter 10 flagged it as well supported experimentally.

A2.7 REFLECTION

Three reflections on the pseudocode.

First, the choice of ERC over a real language is the phase's central decision, and it directly serves the substrate-independence requirement (NFR-3).

Because the logic above commits to no language, it can be translated into Python, into PrologAI, into any language, without rewriting the thinking.

The pseudocode is the contract; the code in each language must honor it.

Second, the pseudocode makes visible how few moving parts the System has.

  https://www.linkedin.com/in/d-r-dison

A reader who feared that modeling a brain requires a hundred bespoke mechanisms can see here that it requires a main loop, six archetype rules, a drive loop, and a learning rule.

The complexity of the System lives in its DATA, the hundred constructs and their connections, not in its CODE.

This is the deepest engineering consequence of the encapsulation thesis, and it is why the System is buildable at all.

Third, an honest note on the ERC extensions, of which there are two kinds, and both are flagged here rather than hidden.

The first is parameter passing: the base ERC catalogue templates procedures without a formal parameter mechanism, and several procedures above take a construct or an interface as their subject in a way that extends the base notation.

The second is composite operations: sentences such as "Gather the total weighted input arriving from all incoming interfaces," "Wrap next phase back into the range of one full cycle," "Find the stored pattern that is most similar to combined pattern," and "Move the activation pattern a small step toward nearest pattern" are not base ERC primitives but named composite operations, each of which stands for a short block of primitive arithmetic, comparison, and list operations that a full implementation must expand, exactly as ERC's own rule requires composites to expand fully to primitives with no residual ambiguity.

They are written at the composite level here for readability, because expanding every weighted sum and every similarity search into primitive loops would bury the logic the pseudocode exists to make plain.

Both kinds of extension are readable and unambiguous, and both are named in the spirit of ERC's rule that a sentence must have one exact meaning.

A full implementation formalizes the parameter convention and expands the composites; the logic is unaffected.

A2.8 NETWORKS, SYSTEMS, AND THE EXTENDED DYNAMICS LABELS

Three notes keep the load-bearing core above consistent with the fuller body of Chapter 2.

First, a large-scale NETWORK (Section 2.10) is not a new kind of update rule; it is a named VIEW over regions that already run their own archetype rules. The System represents a network as a list of member constructs plus its couplings, and computes a network-level quantity (such as its mean activation or its anticorrelation with another network) by reading the members it already updates. No per-network rule is added to the dispatch of Section A2.2.

Second, a cognitive SYSTEM or PATHWAY (Section 2.11) is likewise not a new rule but a named SUBGRAPH or LOOP over the connection graph. The direct, indirect, and hyperdirect pathways are three labelled routes through the existing basal-ganglia constructs; the Papez circuit is a labelled cycle over existing regions and interfaces. The System represents each as an ordered list of the constructs and interfaces it traverses, and runs it by running those constructs and interfaces in order.

  https://www.linkedin.com/in/d-r-dison

Third, the DYNAMICS labels used in the fuller construct profiles (publisher, broadcast publisher, selector, detector, controller, stabilizer, and the compound labels such as integrator-and-publisher or comparator-and-switch) all reduce to the six archetypes of Section A2.3 plus the molecule FIELD kind, applied singly or in sequence. A publisher or broadcast publisher is a construct that, having computed its state by an archetype rule, writes that state to the neuromodulatory or signalling bus (the field kind). A selector is a GATE; a detector is a COMPARATOR or RELAY; a controller and a stabilizer are a GATE or an INTEGRATOR holding a set-point. A compound label names two archetypes run in sequence within one construct. The glia and the tripartite synapse are MODULATORS: they do not carry a signal but adjust the third factor of the learning rule and the persistence of connections, and so enter the System through the plasticity procedure of Section A2.5, not the update dispatch. This preserves the thesis of six rules, not one hundred: nothing above requires a seventh archetype, and, per Assumption 2 of Appendix 1, one is added only if a construct is ever found that resists all six.

A2.9 THE TEST PLAN FOR EACH CORE PROCEDURE

The SPARC Pseudocode phase requires that the pseudocode carry a plan for how each piece will be tested, so the logic and its verification are designed together. Each core procedure above is paired here with the test that verifies it; the tests themselves are written in full in Appendix 4.

THE MAIN LOOP (Section A2.1) is verified by the scheduler test: a run of N ticks updates every construct exactly N times, the two passes never interleave, and no construct reads a neighbour's future within a tick.

THE DISPATCH (Section A2.2) is verified by dispatching a construct of each archetype and asserting the matching rule was called and no other.

THE SIX ARCHETYPE RULES (Section A2.3) are each verified by their isolation test of Appendix 4, Section A4.2: the relay scales, the integrator fills and leaks, the oscillator keeps time, the attractor completes a pattern, the gate resists flipping, and the comparator reports a difference.

THE DRIVE LOOP (Section A2.4) is verified by presenting a monitored variable that departs from and returns to its set-point and asserting that the reward equals the total reduction of drive error and is broadcast as phasic dopamine.

THE LEARNING RULE (Section A2.5) is verified by the learning test: a stimulus paired with reward strengthens the connecting weight by the product of the two activities and the third factor, and homeostatic scaling holds the average activity near target.

THE ELIGIBILITY TRACE (Section A2.6) is verified by rewarding a connection one tick after it was active and asserting the weight still strengthens, and by withholding reward and asserting the trace fades.

Each plan names the single behaviour its test must pin down, in the test-first spirit of Refinement: the procedure and its test are written as a pair, never the procedure alone.

APPENDIX 3. ARCHITECTURE

  https://www.linkedin.com/in/d-r-dison

A3.0 THE PURPOSE OF THIS PHASE

The Architecture phase designs the skeleton of the System so that it will be robust and extensible.

It identifies the major components, the interfaces between them, the data model, and how the whole fits together.

It remains language-agnostic: it says what the parts are and how they relate, not what language they are written in.

In SPARC this is the deliverable named architecture.md.

A3.1 ARCHITECTURAL STYLE

The System adopts a REGISTRY-AND-SCHEDULER style with a BROADCAST BUS, and the choice is dictated by the science rather than by fashion.

The registry holds every construct as an independent module.

The scheduler advances time in ticks and drives the synchronous update.

The broadcast bus carries the neuromodulators, which are not point-to-point and so cannot be edges in the connection graph.

This style is chosen over the two obvious alternatives for reasons the science already established.

It is not a single central controller with everything beneath it, because Chapter 12 showed there is no orchestrator and that building one would recreate the homunculus the science rejects.

And it is not a flat message-passing swarm with no coordination, because Chapter 12 showed the brain has genuine arbiters and broadcast authorities.

The registry-and-scheduler-with-bus style is the architectural form of the Arbitrated Federation: autonomous modules, a dumb scheduler that coordinates their timing without having preferences, and a broadcast layer that sets global state.

The architecture is the thesis of Chapter 12, made buildable.

A3.2 THE MAJOR COMPONENTS

COMPONENT 1, THE CONSTRUCT REGISTRY. Holds every region, interface, and molecule as a module carrying its specification block. It is the System's parts bin, and it is loaded once at startup from the data of Chapters 2, 3, and 7. Adding a construct means adding an entry here and nothing else.

COMPONENT 2, THE SCHEDULER, or TICK ENGINE. Advances the System one tick at a time. On each tick it invokes the two-pass synchronous update of Appendix 2, first computing every next state, then committing them. It is deliberately dumb: it holds no state of its own beyond the tick count and has no opinions about what any construct should do. It is the referee that keeps time, not a ruler.

COMPONENT 3, THE CONNECTION GRAPH. Represents the interfaces of Chapters 3 and 4 as directed, weighted, delayed links between constructs, each marked transmissive or computational. This is the connectome as data. The learning

  https://www.linkedin.com/in/d-r-dison

rule of Appendix 2 mutates the weights here; the routing is otherwise fixed.

COMPONENT 4, THE NEUROMODULATORY BUS. Represents the broadcast fields of Chapter 7 as global quantities with a level per target territory. Every construct's update rule and every learning rule reads from the bus. The bus is written by the source nuclei (the ventral tegmental area writes dopamine, the locus coeruleus writes norepinephrine, and so on) and read by everyone. It is the one place the architecture deliberately permits a global variable, exactly as the biology does.

COMPONENT 5, THE DRIVE SYSTEM. Holds the homeostatic drives of Chapter 9, each monitoring a body variable against a set-point, each computing an error, and together computing the reward that the bus broadcasts as dopamine. It is the System's source of motivation and the origin of the learning signal.

COMPONENT 6, THE ACTION SELECTOR. Implements the basal-ganglia arbitration of Chapter 2: it receives candidate actions proposed by the drives and cortical processes, and releases exactly one while suppressing the rest. It generates no actions of its own. It is the referee of the federation, and the architecture must enforce, as an invariant, that its output is always one of its inputs and never a fresh invention.

COMPONENT 7, THE OVERRIDE CONTROLLER. Enforces the ranking of Chapter 9 and Chapter 12, allowing older and vital drives to preempt newer deliberation. It is a small, high-priority component that can seize control of the action selector, and it is the architectural home of the fact that breathing beats deliberation.

COMPONENT 8, THE PLASTICITY ENGINE. Applies the three-factor learning rule and homeostatic scaling of Chapter 10 to the connection graph each tick, and runs the slow consolidation phase of Chapter 5 offline. It is the component that makes the map alive.

COMPONENT 9, THE OBSERVER, or INSTRUMENT. Reads and records the full state of every component at every tick, satisfying the inspectability requirement (FR-13, NFR-1). It is passive: it changes nothing, and it sees everything. It is the microscope trained permanently on the transparent brain, and it is what makes the System a scientific instrument rather than a black box.

COMPONENT 10, THE PERTURBATION INTERFACE. Permits an experimenter to lesion a construct, cut an interface, or shift a neuromodulator level, satisfying FR-14. It is how a simulated disease or drug is administered.

COMPONENT 1 additionally holds the large-scale networks (Section 2.10) as named views over its regions and the cognitive systems and pathways (Section 2.11) as named subgraphs over the connection graph; adding a network or a pathway means adding a named grouping, not a new update rule. The glia are held as modulatory entries read by the plasticity engine (Component 8), not as regions with state.

A3.2bis THE SYSTEM DIAGRAM

The SPARC Architecture phase calls for a system diagram. The following diagram, in the text-based Mermaid notation SPARC recommends, shows the ten components and the flow of one tick. It commits to no language; it shows only shape.

  https://www.linkedin.com/in/d-r-dison

```
graph TD
  SCH[Scheduler / Tick Engine] --> DRV[Drive System]
  DRV -- reward --> BUS[Neuromodulatory Broadcast Bus]
  BUS --> REG[Construct Registry: regions, interfaces, molecules, networks,
systems]
  GRAPH[Connection Graph: directed and broadcast interfaces] --> REG
  REG -- first pass: compute next state --> REG
  REG -- second pass: commit state --> SEL[Action Selector]
  OVR[Override Controller] --> SEL
  SEL --> PLAS[Plasticity Engine]
  GLIA[Glia: tripartite synapse] --> PLAS
  PLAS --> GRAPH
  OBS[Observer / Instrument] -.reads everything.-> REG
  OBS -.reads everything.-> BUS
  OBS -.reads everything.-> GRAPH
  PERT[Perturbation Interface] -.lesion, cut, shift.-> REG
```

Read the solid arrows as the causal flow of a tick, top to bottom: the scheduler starts the tick, the drives set the reward, the bus broadcasts it, the registry updates every construct in two passes, the selector releases one action under the override controller, and the plasticity engine (gated by the glia) writes the new weights back to the connection graph. Read the dotted arrows as the observer and perturbation interface, which touch everything without being part of the causal flow: the observer only reads, and the perturbation interface only injects an experimenter's intervention. The diagram is the Arbitrated Federation of Chapter 12 in one picture.

A3.3 HOW THE COMPONENTS INTERACT: ONE TICK

The data flow of a single tick, which the architecture must realize, is as follows.

The scheduler begins the tick.

The drive system reads the body state and updates every drive error, and computes the reward.

The neuromodulatory bus is updated: the reward becomes dopamine, arousal becomes norepinephrine, and so on.

Then the scheduler runs the first pass: every construct reads its inputs from the connection graph and the bus, and computes its next state by its archetype rule.

The scheduler runs the second pass: every next state is committed.

The action selector receives the candidate actions and releases one, subject to the override controller.

The plasticity engine updates the connection weights.

The observer records everything.

The scheduler ends the tick.

This ordering is itself a specification, and it is not arbitrary.

Drives before modulators, because the reward sets the dopamine.

  https://www.linkedin.com/in/d-r-dison

Modulators before construct updates, because the gain and the third factor must be current.

First pass fully before second pass, because the update is synchronous.

Action selection after the updates, because it selects on the new state.

Plasticity after selection, because learning depends on what was done.

An implementer who reorders these steps will have built a different machine, and the architecture states the order precisely for that reason.

A3.4 THE DATA MODEL

The System's data, expressed without commitment to any language or storage format, is the following.

A CONSTRUCT has: a name; a kind (region, interface, or molecule); an archetype (one of the six, or field, for a molecule); a state (its activation, gain, adaptation, and any tuning or phase); a set of parameters; and its specification block.

An INTERFACE, additionally, has: a source construct; a destination construct; a weight; a delay; a transmissive-or-computational flag; and, if computational, a reference to the transform it performs.

A MOLECULE, additionally, has: a source construct or set of sources; a set of target territories; and a current level per target.

A DRIVE has: a monitored body variable; a set-point; a current error; a previous error; and an override rank.

The BODY STATE has: the current value of every monitored variable (temperature, glucose, oxygen, and the rest), which the drives read and which the System's actions and the passage of time change.

A RECORD, produced by the observer each tick, has: the tick number; and a snapshot of every construct's state, every interface's weight, every drive's error, and every neuromodulator's level.

A NETWORK has: a name; a list of member constructs; a set of subscribe-and-publish couplings among them; and a functional signature (such as its characteristic anticorrelation with another named network). A network is a VIEW over the registry, not a separate store; it owns no state its members do not.

A SYSTEM or PATHWAY has: a name; an ordered list of the constructs and interfaces it traverses; and, where it is a loop (the Papez circuit, a cortico-basal-ganglia-thalamo-cortical loop), a marker that it closes on itself. It is a named SUBGRAPH over the connection graph.

A GLIAL PARTICIPANT (astrocyte, microglia, oligodendrocyte) and the TRIPARTITE SYNAPSE have: a name; a target (a synapse, a region, or an axon); and a modulatory function on the plasticity engine (gating the learning third factor, pruning a connection, or setting a conduction delay). They carry no activation of their own.

  https://www.linkedin.com/in/d-r-dison

Every interface and every molecule additionally carries a MODE flag: DIRECTED (a send-and-receive channel with one subscriber) or BROADCAST (a publish-and-subscribe topic whose subscribers are defined by receptor), per FR-17.

Note what this data model already is.

It is very close to the Causalontology data structure discussed in Appendix 6: constructs are continuants, their states are token states, the interfaces are conduits, and the tick-by-tick records are token occurrences. A large-scale network is a continuant grouping of continuants; a cognitive system or pathway is a chain (or a closing loop) of causal-relation-objects over those conduits; and the send-and-receive versus publish-and-subscribe mode of an interface is exactly the difference between a directed conduit and a broadcast one.

The architecture and the data standard were designed to meet, and Appendix 6 develops the meeting.

A3.5 THE TECHNOLOGY STACK, DEFERRED

In an ordinary SPARC project, the Architecture phase selects the technology stack, and this appendix DELIBERATELY DECLINES that one activity, the single conscious deviation from the standard SPARC Architecture phase.

Here it deliberately does not, because substrate independence (NFR-3) is a first-class goal, and the whole point of Appendix 6 is that the stack is a free choice the builder makes.

The architecture commits only to the SHAPE of the System, the ten components and their interactions, which any stack must realize.

The candidate stacks, a Causalontology data layer, a Python Minimum Viable Product, and a PrologAI implementation, are described in Appendix 6, and the architecture is compatible with all of them by design.

A3.6 SCALABILITY, ROBUSTNESS, AND FAILURE POINTS

The architecture meets the scalability requirement (NFR-4) by its grain: at one hundred constructs, a tick is a hundred update calls, trivial on a personal computer, and the design permits any construct to be internally refined to finer grain without changing the scheduler, the bus, or the graph.

The System scales up by refining constructs, not by rebuilding the frame.

The architecture is robust by its modularity: a fault in one construct is contained to that construct, exactly as the encapsulation thesis predicts, and the observer will show precisely where a fault occurred, because nothing is hidden.

The honest failure points are three.

The synchronous two-pass update is the correctness-critical section, and fusing the passes is the most likely and most damaging error.

The neuromodulatory bus is a global shared resource, and while the biology sanctions it, an implementation must ensure the bus is read consistently within a tick.

  https://www.linkedin.com/in/d-r-dison

And the action selector's invariant, that it never invents an action, must be enforced in code and tested, because a selector that begins to generate is the homunculus creeping back in, and the architecture must forbid it structurally.

A3.7 REFLECTION

The architecture is, almost line for line, the Arbitrated Federation of Chapter 12 rendered as components.

This is the strongest possible validation that the science and the specification are one work: the answer to the book's deepest question, is there an orchestrator, became the System's top-level design without translation.

There is a scheduler that keeps time but has no preferences, a bus that sets global state without knowing what is computed, a selector that arbitrates without generating, and an override controller ranked by vitality.

No central controller appears, because the science forbade one, and the architecture that omits it is the architecture the science demanded.

The single most important architectural decision is the passivity of the observer and the completeness of what it sees.

It is what makes this System, unlike the dominant models, a scientific instrument: a mind that can be watched, at every level, at every tick, without disturbing it.

That property is not a feature added to the architecture; it is the reason the architecture was worth building.

APPENDIX 4. REFINEMENT

A4.0 THE PURPOSE OF THIS PHASE

The Refinement phase improves the design and the eventual code through iteration: repeated cycles of testing and revision.

Its signature practice is Test-Driven Development, abbreviated TDD, in which one writes a failing test first, then just enough code to make it pass, then cleans up, a cycle called Red-Green-Refactor.

SPARC favours the London School of TDD, which tests each unit in isolation using stand-in fakes for its collaborators, and emphasizes behaviour and interactions.

In SPARC this is the deliverable named refinement.md, together with the growing, tested codebase.

A4.1 THE TEST-DRIVEN LOOP, APPLIED TO A BRAIN

The System is unusually well suited to test-driven development, for a reason worth naming: its acceptance criteria (Appendix 1) are drawn from real neuroscience experiments, and a neuroscience experiment IS a test.

  https://www.linkedin.com/in/d-r-dison

To test the System is to run, in software, the lesion and learning studies that are the evidential basis of the science.

Every test below is a famous experiment rewritten as an automated check.

The loop for each construct and each behaviour is: first write a test that states the expected result and fails because the code does not yet exist (Red); then write the smallest update rule or connection that makes the test pass (Green); then clean up the code without changing its behaviour (Refactor); then move to the next.

A4.2 THE UNIT TESTS: EACH ARCHETYPE IN ISOLATION

Following the London School, each archetype rule is tested in isolation, with its inputs supplied by fakes.

TEST OF THE RELAY: given a fixed input and a fixed gain, the next activation equals input times gain.

A pipe passes what it is given, scaled.

If this fails, the relay is not a relay.

TEST OF THE INTEGRATOR: given a constant input over many ticks, the activation rises and approaches a bound; given no input, it decays toward zero.

A bucket fills and leaks.

TEST OF THE OSCILLATOR: over many ticks with no input, the phase advances at the natural frequency and wraps, and the gain rises and falls once per cycle.

A metronome keeps time.

TEST OF THE ATTRACTOR: given a partial version of a stored pattern, the activation converges to the whole stored pattern; given a pattern near two stored ones, it converges to the nearer.

A memory completes from a fragment.

TEST OF THE GATE: given a sub-threshold drive, the mode does not change; given a supra-threshold drive, it flips and then resists flipping back.

A stiff switch.

TEST OF THE COMPARATOR: given equal expected and actual inputs, the output is zero; given an actual input above expectation, the output is positive; below, negative.

A scale reports the difference.

This is also, precisely, the test that the dopamine reward-prediction-error signal behaves correctly.

A4.3 THE INTEGRATION TESTS: CONSTRUCTS WORKING TOGETHER

Integration tests check that constructs wired together produce the right chain of behaviour.

  https://www.linkedin.com/in/d-r-dison

TEST OF PROPAGATION: a signal presented at a sensory region appears at its cortical target after exactly the number of ticks dictated by the interface delays, and nowhere sooner.

This tests the connection graph and the scheduler together.

TEST OF THE DECISION CHAIN: a valued stimulus presented to the sensory regions leads, through the valuation regions and the action selector, to the selection of the appropriate action, in the correct order.

This tests the perception-to-action arc.

TEST OF THE COMPUTATIONAL INTERFACE: the inferior olive, given an intended movement and a differing outcome, emits an error signal that is neither of its inputs.

This tests that a computational interface computes rather than conveys, the Perfect Wire Test made into an assertion.

A4.4 THE SYSTEM TESTS: THE FAMOUS EXPERIMENTS

The system-level tests are the acceptance criteria of Appendix 1, and each reproduces a landmark result.

THE PARKINSON TEST (AC-4): deplete dopamine to a low fraction of normal; action selection slows and then halts; restore dopamine; it recovers.

This reproduces the cardinal signature of Parkinson's disease and the effect of dopamine-replacement therapy, in software, from first principles.

Passing it is strong evidence the System's action selection is faithful.

THE LEARNING TEST (AC-5): pair a stimulus with reward repeatedly; the connecting weights strengthen by the three-factor rule until the stimulus predicts the reward; then present the stimulus alone; the weights weaken and the prediction extinguishes.

This reproduces Pavlovian conditioning and extinction, the foundation of the psychology of learning.

THE H.M. TEST: disable the plasticity engine; the System continues to perceive, act, and use what it already knows, and forms no new memories.

This reproduces the amnesia of patient H.M., the case on which this whole book centres, and it is the test that proves the plasticity engine is the seat of new memory.

THE DRIVE-CONFLICT TEST (AC-6): raise two drives at once; the override hierarchy resolves them in the specified order; and no amount of deliberative input can suppress the respiration drive to self-termination.

This reproduces the vitality ranking of Chapter 12 and enforces the System's one inviolable safety property.

THE SOCIAL-PAIN TEST: deliver a social-rejection signal; the same circuitry that registers physical pain activates.

  https://www.linkedin.com/in/d-r-dison

This reproduces one of the more striking findings of social neuroscience and tests that the social-connection drive is a genuinely defended set-point.

A4.4bis THE NETWORK, PATHWAY, AND GLIAL TESTS

Because the body names constructs above the single region, the test suite adds checks at that level.

THE NETWORK TEST: engage a task that recruits the dorsal attention and frontoparietal networks; the default mode network's mean activation must fall as theirs rises, reproducing the task-positive-versus-task-negative anticorrelation of Section 2.10. A network that does not deactivate on task has not been coupled correctly.

THE PATHWAY TEST: drive the direct pathway and observe an action released (go); drive the indirect pathway and observe it suppressed (no-go); drive the hyperdirect pathway and observe a fast global cancellation (stop). Then run the Papez circuit and confirm a memory trace injected at the hippocampus circulates through the mammillary bodies, anterior thalamus, and cingulate and returns, and that cutting any node breaks the loop, reproducing diencephalic amnesia.

THE GLIAL TEST: disable the astrocytic gating of the tripartite synapse and confirm that plasticity is impaired even though every neuronal construct is intact, reproducing the finding that the glia gate learning. This is the tripartite-synapse counterpart of the H.M. test: memory fails not because a region is lost but because its plasticity partner is silenced.

A4.5 OPTIMIZATION AND THE HARDENING LOOP

Once the tests pass, Refinement optimizes.

The candidate optimizations, in order of value, are: caching the neuromodulatory bridge closures so the bus is not recomputed needlessly; vectorizing the construct update so a tick over one hundred constructs is one batched operation rather than one hundred calls; and profiling to find the few constructs whose rules dominate the tick cost and simplifying them.

None of these may change the System's behaviour: an optimization that changes a test result is a bug, by definition.

The hardening loop then runs the System for long durations and looks for the failure modes the science predicts and warns against: runaway excitation (a seizure, which the homeostatic scaling must prevent), weight explosion (which the scaling must also bound), and a frozen action selector (which indicates a stuck gate).

Each discovered failure becomes a new test, so that the System can never regress into it again.

A4.6 THE KNOWN TRADE-OFFS

Refinement is where trade-offs are made explicit, and honesty requires stating them.

THE GRAIN TRADE-OFF: the regional grain buys tractability and inspectability at the cost of fidelity.

  https://www.linkedin.com/in/d-r-dison

A finer grain would be more faithful and far more expensive.

The System accepts the coarser grain for version one and provides the path to refine it.

This is the largest trade-off in the work.

THE ARCHETYPE TRADE-OFF: six archetypes buy simplicity and uniformity at the cost of some accuracy for constructs that fit no archetype cleanly.

The System accepts a small loss of accuracy for a large gain in buildability, and extends the archetype set only when a construct genuinely resists all six.

THE DETERMINISM TRADE-OFF: strict reproducibility (NFR-2) buys scientific rigour at the cost of the biological realism of noise, which the brain uses.

The System resolves this by making randomness seeded and recorded, so that it has controlled noise that is also exactly replayable, keeping both rigour and realism.

THE SPEED-VERSUS-CLARITY TRADE-OFF: the two-pass synchronous update is slower than an in-place update would be, but it is correct where the faster version is subtly wrong.

The System pays the speed to keep the correctness, and this is not negotiable.

A4.7 REFLECTION

The Refinement phase reveals a property of this System that few software projects enjoy: its tests are not invented by its builders but inherited from a century of neuroscience.

The Parkinson test, the learning test, the H.M. test, the drive-conflict test, and the social-pain test are not checks a programmer dreamed up to exercise the code; they are the actual experiments that established the science the code implements.

This gives the System an unusually strong form of validation.

It is correct not merely when it does what its author intended, but when it reproduces what nature does, as measured by experiments its author did not run.

A model that can be given Parkinson's disease by depleting a single simulated molecule, and cured by restoring it, has passed a test no amount of ordinary software testing could design.

APPENDIX 5. COMPLETION

A5.0 THE PURPOSE OF THIS PHASE

The Completion phase finishes the System properly: it tests everything, documents everything, deploys, and monitors.

  https://www.linkedin.com/in/d-r-dison

Its deliverable is a live, documented, monitored System and, in SPARC, the document named completion.md.

Completion is where a project stops being a promising prototype and becomes something another person can install, run, trust, and build upon.

A5.1 THE FULL TEST STRATEGY

The completed System must carry a complete test suite at three levels, all automated and all runnable with a single command.

UNIT LEVEL: every archetype rule, every drive, and every learning step tested in isolation, as in Appendix 4.

INTEGRATION LEVEL: the propagation, decision-chain, and computational-interface tests, verifying that constructs work together.

SYSTEM LEVEL: the famous-experiment tests, the Parkinson, learning, H.M., drive-conflict, and social-pain reproductions, verifying that the whole System behaves as a brain.

The suite must reach the coverage targets SPARC sets, and it must run in continuous integration, meaning automatically on every change, so that no change can silently break a reproduced experiment.

The definition of a passing build is that every one of the acceptance criteria of Appendix 1 is met by an automated test, and all pass.

For that claim to be checkable rather than merely asserted, the mapping from each acceptance criterion to its covering test is stated explicitly here.

AC-1 (every named construct present with its specification block) is covered by a registry-completeness unit test that iterates the constructs of Chapters 2, 3, and 7 (regions, interfaces, and molecules) and asserts each is loaded.

AC-10 and AC-11 (the large-scale networks of Section 2.10 and the cognitive systems, pathways, and glia of Section 2.11) are covered by an extension of the same registry-completeness test, which additionally iterates every named network as a grouping over its members, every named system and pathway as a subgraph, and every glial participant, and asserts each is present with its couplings, its mode flags (directed or broadcast), and its modulatory function.

AC-2 (one tick updates every construct exactly once) is covered by a scheduler unit test that counts updates per tick.

AC-3 (a signal propagates through the correct interfaces in the correct number of ticks) is the propagation integration test.

AC-4 (dopamine depletion halts and restoration recovers action selection) is the Parkinson system test.

AC-5 (a stimulus paired with reward comes to predict it, and extinguishes) is the learning system test.

AC-6 (two drives are arbitrated in override order, and respiration cannot be self-suppressed) is the drive-conflict system test.

  https://www.linkedin.com/in/d-r-dison

AC-7 (every state, weight, drive error, and neuromodulator level is reportable) is covered by an observer unit test that reads each quantity at a sampled tick.

AC-8 (a run is exactly reproducible from its seed) is covered by a determinism test that runs the same seed twice and asserts identical trajectories.

AC-9 (the System runs at the regional grain within its time budget) is covered by a performance test measuring wall-clock time per thousand ticks.

Every acceptance criterion thus has exactly one named covering test, and no criterion is left unverified.

The H.M. test, which reproduces the amnesia of the disabled plasticity engine, is an additional system test beyond the nine criteria, included because the case is central to the book.

A5.2 THE DEFINITION OF DONE

The System is DONE, for version one, when all of the following are true.

Every named construct of Chapters 2, 3, and 7 is present and running, and every large-scale network, cognitive system and pathway, and glial participant of Sections 2.10 and 2.11 is present as a named grouping, subgraph, or modulator.

Every acceptance criterion of Appendix 1 passes as an automated test.

The System runs at the regional grain on a personal computer within its stated time budget.

Every state, weight, drive, and neuromodulator is inspectable at every tick.

A full run is exactly reproducible from its seed.

The five famous-experiment tests pass.

The documentation described below is complete.

And the System is installable by a stranger from its repository in a documented number of steps.

Note what "done" does NOT require: it does not require that the System be intelligent, or that it pass any test of general capability.

Version one is done when it is a faithful, transparent, runnable model of the regional architecture.

Whether such a model, scaled and fed, becomes intelligent is the question of Appendix 6, and it is explicitly out of scope for Completion of version one.

Conflating "the model runs faithfully" with "the model is a mind" would be the one dishonesty this book has refused throughout, and Completion refuses it here.

  https://www.linkedin.com/in/d-r-dison

A5.3 DOCUMENTATION

Two bodies of documentation are required.

FOR THE USER: a guide that explains how to install the System, how to run it for a number of ticks, how to present an input, how to read the observer's output, how to run each of the famous experiments, and how to lesion a construct or shift a molecule.

The user guide must let a curious person reproduce the Parkinson experiment within an hour of installing.

FOR THE DEVELOPER: a guide that explains the ten components, the tick ordering, the data model, how to add a new construct, how to add a new archetype, and how the code maps to the chapters of this book.

The developer guide's central promise is that every line of code traces to a passage of the science, and it must make that tracing explicit, so that the System remains a glass box not only at runtime but in its source.

This book itself is the specification-level documentation, and the appendices are written so that the mapping from science to system is legible without a separate document.

That is by design: the book and the specification are one artifact.

A5.4 DEPLOYMENT AND MONITORING

Deployment of a scientific instrument differs from deployment of a consumer product, and the plan reflects that.

The System is deployed as an open-source repository that a researcher clones and runs locally, not as a hosted service.

Its "deployment" is a tagged, versioned release with a reproducible environment specification, so that a run on one machine reproduces on another.

This matters more here than in most software: a scientific result that cannot be reproduced on another machine is not a result.

Monitoring, for a System whose defining virtue is inspectability, is the observer of Appendix 3 made permanent.

The running System continuously reports its own state, and the monitoring layer watches for the pathological signatures the science warns of: runaway synchrony, weight explosion, a frozen selector, a collapsed drive.

These are the System's equivalent of the vital signs a clinician watches, and their appearance in a run is either a bug to fix or, intriguingly, a simulated disease to study.

A5.5 SECURITY, SUCH AS IT IS

A local scientific model has a light security surface, but the SPARC rules still bind.

No secret or credential may appear in the code.

All external input, meaning any dataset fed to the System, must be validated before it is trusted.

And, because this System is a model of a mind, one ethical safeguard is added to the ordinary security checklist and made a hard requirement: the token-level records of any run that includes data about real, identifiable people are local by default and are never published without deliberate act, exactly the safety discipline the companion data standard requires.

A model of a brain that ingests human data inherits the duty of care that comes with human data.

A5.6 POST-COMPLETION: WHAT SHIPS

What ships at the end of Completion, version one, is: a running, open-source, glass-box model of the human cognitive architecture at the regional grain; a complete test suite that reproduces five landmark neuroscience experiments from first principles; a user guide and a developer guide; a reproducible release; and a monitoring layer that watches the model's vital signs.

It is a scientific instrument, buildable at home, that anyone can install, run, inspect, perturb, and extend.

What does not ship, and is not claimed, is an intelligence.

That is the next mountain, and the honest map to it is Appendix 6.

A5.7 REFLECTION

Completion is where the temptation to overclaim is strongest, and where this book most firmly declines it.

A working model that reproduces Parkinson's disease and Pavlovian learning from first principles is a genuine achievement, and it would be easy to let the reader believe it is more than it is.

The discipline of Completion is to state exactly what is done and exactly what is not: a faithful, transparent, runnable architecture is done; a mind is not.

The distance between those two is the subject of the final appendix, and naming that distance honestly is itself part of finishing the work.

A specification that ended by pretending the gap was closed would have failed at the last step, however well it had done at every earlier one.

APPENDIX 6. DEMONSTRATION

A6.0 THE PURPOSE OF THIS PHASE

Demonstration is the sixth phase, added by this work to the five of SPARC, giving the acronym SPARCD.

An ordinary specification may end at Completion, with a working system.

A specification of a mind owes more.

  https://www.linkedin.com/in/d-r-dison

This appendix demonstrates three things: that the System as specified is programming-language-agnostic and data-structure-agnostic, so that anyone may build it in any language on any substrate; that concrete implementations are planned and how they relate; and, with full honesty, what this System might and might not become, including the largest and most consequential question a specification of a cognitive architecture can raise, the question of Artificial General Intelligence.

A6.1 THE FIRST DEMONSTRATION: IT IS PROGRAMMING-LANGUAGE-AGNOSTIC

Read back through Appendices 1 through 5 and notice what is absent: there is no programming language anywhere in them.

The requirements are stated in English.

The logic is stated in English Readable Code.

The architecture is stated as ten components and a tick ordering.

The tests are stated as experiments.

Not one line commits to Python, or to PrologAI, or to any language.

This is not an oversight; it is the demonstration.

The System is a set of rules, and rules are substrate-independent.

The six archetype update rules are arithmetic and comparison.

The drive loop is subtraction and accumulation.

The learning rule is multiplication and a bound.

There is nothing in the entire specification that any general-purpose programming language cannot express.

A builder in Python, in PrologAI, in Rust, in Java, in a spreadsheet, in principle even by hand with pencil and paper for a small enough System, can implement this specification, and if two builders implement it faithfully in two different languages, their Systems will produce the same trajectory from the same seed, because they are running the same rules.

This is why the English Readable Code of Appendix 2 was the right choice for the Pseudocode phase, and it is why substrate independence was made a first-class requirement (NFR-3) rather than left to chance.

The specification is the invariant; the language is the free choice.

A6.2 THE SECOND DEMONSTRATION: IT IS DATA-STRUCTURE-AGNOSTIC

The same is true of the data.

Appendix 3 gave the data model as a set of entities and their fields, a construct with a state, an interface with a weight and a delay, a drive with a set-point, without committing to how those are stored.

They could be objects, records, rows in a table, terms in a logic program, or nodes in a graph.

  https://www.linkedin.com/in/d-r-dison

The System does not care, so long as the fields are present and the update rules can read and write them.

So up to this point, the reader who has followed the specification can build their own digital brain at home, in the language they know and the data structure they prefer, and it will be a faithful realization of this book.

That is the demonstration, and it is deliberate: a science that matters should be reproducible by anyone with patience and a computer, not only by its author.

The barrier to building a cognitive architecture has been, until now, the absence of a complete and buildable specification.

This book removes that barrier.

A6.3 THE PLANNED IMPLEMENTATIONS

While the specification is agnostic, the author has three concrete implementations planned, offered here so that a reader may join, adopt, or learn from them.

They are complementary, not competing: each demonstrates the specification on a different substrate.

THE CAUSALONTOLOGY 2.0.0 DATA LAYER.

The data model of Appendix 3 is, by design, close to a general standard for representing causation called Causalontology, whose repository is at https://github.com/ai-university-aiu/causalontology.

In that standard, the System's constructs are continuants, their tick-by-tick states are token states, the interfaces are conduits, and the records the observer produces are token occurrences and token causal claims.

A change order for the next version of that standard, Causalontology 2.0.0, was written specifically to hold the multi-scale, stateful, causal structure this System requires, and a small reference implementation has already been built and shown to run.

The Causalontology layer is the System's memory and its causal record: the substrate in which the digital brain remembers what happened to it and why.

It is the data-structure demonstration made concrete.

THE PYTHON MINIMUM VIABLE PRODUCT.

A Minimum Viable Product, abbreviated MVP, is the smallest version of a system that demonstrates its core value.

A Python MVP of the System, implementing the tick loop, the six archetypes, the drive loop, and the learning rule of Appendix 2 directly, is planned as the reference realization that a student can read, run, and modify on a personal computer.

Python is chosen for the MVP for its readability and its ubiquity, not because the System requires it.

  https://www.linkedin.com/in/d-r-dison

This MVP is the language demonstration made concrete, and it is the one an implementer will most easily pick up first.

A small proof-of-concept has already been run for the data layer; the full cognitive MVP is the natural next step, deliberately not rushed ahead of the specification that governs it.

THE PROLOGAI IMPLEMENTATION.

PrologAI is a cognitive-architecture programming language and substrate, built on Prolog, in which every reasoning step is a named, inspectable rule and every answer comes with a readable justification.

It is a natural home for this System, because the System's glass-box requirement and PrologAI's justification-tree philosophy are the same commitment, made in different places.

Where the Python MVP demonstrates the System on a conventional imperative substrate, the PrologAI implementation demonstrates it on a symbolic, self-explaining one.

And it carries a second purpose.

This System is a demanding test of PrologAI: if building the System reveals that PrologAI lacks some needed capability, for instance a clean way to close the reentrant loops the brain is full of, then that discovered gap becomes a change order for PrologAI itself, exactly as the System's needs became a change order for Causalontology.

PrologAI remains a living language, and this System is one of the forcing functions that will grow it.

The relationship among the three is simple.

Causalontology is where the System remembers.

Python is where a newcomer first builds it.

PrologAI is where it reasons transparently and where the substrate itself may grow to meet it.

All three realize the one specification, and none is privileged by it.

A6.3a WHAT IT CAN ACTUALLY DO: THE LADDER OF DEMONSTRABLE MILESTONES

The demonstrations above answer a builder's question - can I build this, in my language, on my data structure - and the answer is yes.

But a reader is entitled to a second, blunter question: so what can it actually do, and how would I ever see it doing anything at all?

A map to a distant summit is worth little if it does not also mark the first few steps you can take today and check underfoot.

This section supplies exactly that: a concrete, ordered ladder of milestones, from the smallest sign of life to a body moving in the world, each one a thing you can run and watch, each one earning the right to attempt the next.

  https://www.linkedin.com/in/d-r-dison

The ladder borrows its shape from the one place in the universe where general intelligence is known to assemble itself from parts - the developing human child.

A child does not begin by reasoning about justice or driving a car.

A child begins by noticing that a hidden toy still exists, by reaching and grasping, by babbling before speaking, and only much later by planning, arguing, and empathising.

The Swiss psychologist Jean Piaget named four broad stages in that ascent - the sensorimotor stage from birth to about two years, the preoperational stage from about two to seven, the concrete operational stage from about seven to eleven, and the formal operational stage from about eleven onward - and the young field of developmental robotics, which builds robots that learn in the same infant-to-adult order rather than being programmed complete, has shown that this same staged ladder is a practical engineering plan for a machine, not merely a description of children.

This appendix adopts that plan.

The milestones below are grouped into six rungs, and each rung is a stage a child passes through, translated into a test a System can be made to pass.

A note before the rungs, on the two kinds of test the software world already names, because the ladder uses both.

A smoke test, named for the hardware engineer's practice of powering on a new board and checking that no smoke rises, is the very first and cheapest check: does the thing switch on and run at all without falling over.

A pilot test is a small, real, end-to-end trial of the whole system on a genuine task, run before any full-scale deployment, to see whether the parts work together in the world rather than only on the bench.

Every rung below opens with a smoke test - the cheapest possible sign that the rung is alive - and closes with a pilot test - a small but real end-to-end task that shows the rung genuinely working.

RUNG ZERO: THE SMOKE TESTS - IS IT ALIVE AT ALL

Before the System can be said to do anything, it must first simply run, and the smoke tests establish that with the least possible ceremony.

Smoke test one, the heartbeat: start the System with no input and let the tick loop run, and confirm that the clock advances, that every construct holds a state, and that nothing crashes or drifts to infinity.

A System that can idle stably for a million ticks without exploding or freezing has passed its first and most basic sign of life, exactly as a newborn's steady heartbeat is the first thing a delivery room checks.

Smoke test two, the reflex: deliver a single sharp input - a spike on one sensory construct - and confirm that it propagates along the specified interfaces, changes the states it should change, and settles back down.

This is the machine equivalent of the knee-jerk reflex, and it proves that the bus, the scheduler, and the graph are wired and conducting.

  https://www.linkedin.com/in/d-r-dison

Smoke test three, the drive pulse: let one drive fall below its set-point and confirm that the drive loop registers the deficit, raises the corresponding urge, and quiets it again when the deficit is met.

This proves the System has an inside - a state it prefers - which is the seed of everything the higher rungs build on.

The pilot test for Rung Zero is the determinism check already required by the specification: run the System twice from the same seed with the same input, and confirm the two trajectories are identical tick for tick, because a mind you cannot reproduce is a mind you cannot study.

Passing Rung Zero does not make the System intelligent; it makes it a stable, inspectable, reproducible organism ready to learn, which is the sensorimotor infant at the moment of birth.

RUNG ONE: THE SENSORIMOTOR STAGE - PERCEIVING, ACTING, AND THE FIRST CONCEPTS

Piaget's sensorimotor stage, birth to about two years, is where a child learns that the world is made of persistent objects, that its own actions cause effects, and that a thing hidden is not a thing gone - the milestone called object permanence.

This is the System's first real rung, and it is where perception is wired in, because until now the System has only received abstract spikes and must now be connected to the actual world.

Here the appendix answers several of the plain questions a reader will have about what the System can take in.

Can it take text input?

Yes, and this is the simplest sensory front-end of all: text arrives as a stream of tokens, each token published onto a language-sensory construct exactly as a retinal spike is published onto a visual one, and the System's job is to let that stream drive its internal state.

Text is the easiest modality to wire because it is already symbolic, and so it is the first sense the System is given.

Can it see objects, or name what is in front of it?

Yes, by a design choice this book makes openly: the System does not need to re-solve the enormous problem of raw vision from scratch, because that problem is already solved well by existing vision models.

A vision-language model - a large neural network, abbreviated VLM, that looks at an image and returns a description in words, such as a cup, on a table, to the left - is used as the System's eye, a sensory organ bolted on at the front.

The VLM converts pixels into named objects and relations, and those named objects are published onto the System's perceptual constructs, so the System perceives a world of things and relationships rather than a wash of colour.

  https://www.linkedin.com/in/d-r-dison

This is precisely how developmental-robotics laboratories now build embodied agents: the heavy perceptual lifting is done by a pre-trained vision model acting as the retina and early visual cortex, and the cognitive architecture reasons over the clean, named output.

The System is glass-box where it matters - in its drives, its learning, its reasoning - and it is free to use an opaque, off-the-shelf organ for the one job, raw sensory transduction, that biology also delegates to specialised, hard-wired hardware.

Can it hear through a microphone?

Yes, by the same strategy applied to sound.

A microphone captures audio, and a speech-to-text model - an automatic speech recognition system, abbreviated ASR, that turns spoken sound into written words - acts as the System's ear and auditory cortex, converting the pressure waves into a token stream that enters exactly where typed text enters.

So the System hears by having a specialised organ transcribe the world's speech into the symbolic input it already knows how to receive, and if a text-to-speech model, abbreviated TTS, is added on the output side, the System can also answer aloud, closing a natural spoken loop.

The smoke test for Rung One is the naming test: place an object in front of the camera and confirm that the correct named object appears on the System's perceptual construct, or speak a word into the microphone and confirm the correct token arrives - the sense organs are connected and reporting.

The central pilot test for Rung One is object permanence, the defining sensorimotor milestone: show the System an object, hide it behind a screen, and confirm that the System's internal representation of the object persists while it is out of view, so that the System is surprised - registers a prediction error - if the screen is lifted and the object is gone.

A System that maintains a hidden object has crossed the threshold every human infant crosses near its first birthday, and it is the first milestone that is unmistakably cognitive rather than merely mechanical.

A second Rung One pilot test is simple cause and effect: let the System discover, through its own reward-gated learning, that one of its actions reliably produces a pleasant change in the world - the machine equivalent of an infant learning that shaking a rattle makes a sound - which demonstrates goal-directed action, the other great sensorimotor achievement.

RUNG TWO: THE PREOPERATIONAL STAGE - SYMBOLS, LANGUAGE, AND CHATTING

Piaget's preoperational stage, roughly ages two to seven, is the age of symbols and language: the child now uses words and images to stand for things that are not present, talks, asks endless questions, and pretends.

This rung answers the question many readers ask first: will the System be able to chat, like a chatbot?

  https://www.linkedin.com/in/d-r-dison

The honest answer is yes, but in a way worth being precise about, because it differs from how a chatbot chats.

A conventional chatbot is a large language model that produces fluent replies by predicting likely next words, with no inner state, no goals, and no memory of caring about anything.

The System, at this rung, chats from the inside out: words arrive through the text or speech front-end of Rung One, they change its internal state, that state is shaped by its drives and its memories, and its reply is the outward expression of that changed inner state.

Early in the rung the System's language will be simple, even childlike, because its symbolic machinery is still forming, exactly as a three-year-old's speech is simple - and that simplicity is a feature of an honest developmental account, not a defect to be hidden by bolting a fluent language model over the top.

The smoke test for Rung Two is the single-symbol exchange: give the System a word it has grounded through experience and confirm it can produce that word to refer to the thing, and recognise the word when it hears it - symbol use in both directions.

The pilot test for Rung Two is a short grounded conversation: a back-and-forth exchange, by text or by voice, in which the System's replies are demonstrably driven by its internal state and its history rather than by next-word prediction alone - for instance, it answers differently when a drive is satisfied than when the same drive is in deficit, so that its mood colours its words.

A chatbot that has no inner weather cannot do this; a System with drives can, and that difference is the whole point.

A second Rung Two milestone is pretend, or imagination: using the System's reality-quarantine machinery to run a small imagined scenario - to represent a situation that is not happening and reason within it without confusing it for the real record - which is the preoperational child's leap into symbolic play.

RUNG THREE: THE CONCRETE OPERATIONAL STAGE - LOGIC, REASONING, AND RULES

Piaget's concrete operational stage, roughly ages seven to eleven, is where thought becomes logical about concrete things: the child can classify, order, reason about cause and consequence, and grasp that quantities are conserved when their appearance changes.

This rung is where the reader's questions about logic and reasoning tests are answered, and it is a natural strength of this System because its foundation is a causal model of the world.

The smoke test for Rung Three is a single valid inference: give the System two grounded facts from which a third follows - if the block is on the cup and the cup is on the table, then the block is above the table - and confirm it derives the conclusion and can show the steps.

Because the System is transparent, this test also checks something a black box cannot offer: that the reasoning is inspectable, each step a named, watchable operation rather than an opaque leap.

  https://www.linkedin.com/in/d-r-dison

The pilot tests for Rung Three are a graded battery of concrete reasoning trials, each drawn from the developmental literature so that success has a clear human referent.

A classification trial: sort a set of perceived objects into categories the System has formed from experience, and re-sort them by a different property on request, demonstrating flexible categorisation.

A seriation trial: arrange objects in order along a dimension, such as size, demonstrating ordered comparison.

A conservation trial: confirm that the System understands a quantity is unchanged when only its appearance changes - the same water in a taller, thinner glass is still the same amount - which is the concrete-operational child's signature achievement and a direct test of whether the System models the world rather than its mere surface.

A causal-reasoning trial: present a small chain of events and ask the System which event brought about which, using the Causalontology causal record, and confirm it distinguishes a genuine cause from a mere co-occurrence - the very thing the dominant correlational models famously cannot reliably do.

A planning trial: give the System a goal state and a set of available actions, and confirm it can compose a short sequence of actions that reaches the goal, demonstrating means-end reasoning.

Passing this rung shows the System reasoning logically and causally over concrete, grounded situations, with every inference open to inspection, which is exactly the capability the reigning paradigm is weakest at and this architecture is built to be strong at.

RUNG FOUR: THE EMOTIONAL AND SOCIAL RUNG - FEELING, AND MODELLING OTHER MINDS

Human development is not only cold cognition; alongside Piaget's ladder runs an emotional and social one, and a specification of a mind that reproduces the brain's drive, affect, and social systems owes an account of how those are demonstrated too.

The reader asks about emotional tests, and the System can be given them because it has, by construction, the machinery of emotion: drives that can be satisfied or thwarted, and neuromodulatory states that colour its whole cognition the way feelings colour ours.

The smoke test for Rung Four is an affective-response check: thwart a drive and confirm the System enters a measurably different global state - a negative affect - that changes how it perceives, learns, and acts, and confirm the opposite state arises when the drive is satisfied.

This shows the System has something it is like to be in - an inner valence - which is the raw material of emotion.

The pilot tests for Rung Four cover both feeling and the reading of others.

An appraisal trial: present a situation and confirm the System appraises it as good or bad for its own goals, producing an appropriate emotional state

  https://www.linkedin.com/in/d-r-dison

- fear at a threat to a set-point, something like satisfaction at a met one.

A regulation trial: confirm the System can bring an aroused emotional state back toward baseline over time, demonstrating the beginnings of emotional regulation rather than being permanently captured by every jolt.

A theory-of-mind trial, the great social milestone: confirm that the System can attribute to another agent a mental state different from its own - and, at the harder level, a false belief, the classic test in which a child must predict that a person will look for an object where that person last saw it, not where the child knows it now is.

Modelling another mind, and especially another mind that believes something false, is widely held to be a prerequisite for genuine social intelligence, and it is one of the four data-structure capabilities this program has already identified as a candidate for the next major version of the Causalontology layer, so this rung is also where the architecture's own growth is stress-tested.

An empathy trial, at the top of the rung: confirm that another agent's modelled distress can move the System's own affective state and shape its behaviour toward help, which is the seed of pro-social action.

Passing Rung Four shows a System that not only thinks but feels its situation and reads the situations of others, which is the part of intelligence a purely logical account leaves out and a faithful model of the brain cannot.

RUNG FIVE: THE EMBODIED RUNG - A ROBOT BODY, AND ACTING IN THE REAL WORLD

The highest rung on this ladder returns the System to where a child's intelligence is forged: a body, in a world, with real stakes.

The reader asks whether the System will be able to inhabit a robot body and use the robot's sensors, and the answer is that this is not an afterthought but the destination the whole architecture is aimed at, because its cognition is grounded in the regulation of a body state and a body is what closes the loop.

Inhabiting a robot means three connections, each of which the rungs below have already prepared.

The robot's sensors - its cameras, microphones, touch and joint sensors - become the System's senses, feeding the perceptual constructs through the same vision, speech, and touch front-ends that Rung One installed, so the System perceives the robot's world.

The robot's body state - its battery level, its temperature, its balance, its physical integrity - becomes the System's homeostatic state, the set-points its drives defend, so that a low battery is felt as hunger and a risk of falling as something like fear, giving the System real skin in the game.

The robot's actuators - its wheels, arms, grippers, and voice - become the System's means of acting, driven by the same goal-directed, reward-gated machinery that Rung One first exercised on a rattle, so that the System moves its body to meet its needs.

  https://www.linkedin.com/in/d-r-dison

The smoke test for Rung Five is the sense-act loop: confirm that a change the robot's camera sees produces, through the System, a movement of the robot's body - perception in, action out, the loop closed through a real machine.

The pilot tests for Rung Five are graded embodied tasks, each a small, real, end-to-end trial in a body.

A reaching-and-grasping trial: confirm the System can guide the robot's arm to pick up a named object it perceives, the embodied form of the infant milestone of coordinated reaching.

A navigation trial: confirm the System can move the robot toward a goal location while avoiding obstacles, using its persistent model of a world that continues to exist outside the current view - object permanence, now load-bearing.

A homeostatic-survival trial, the one that most fully exercises the whole architecture: place the robot in an environment where it must act to defend its own body state - seeking a charging station as its battery drive falls, avoiding a hazard that threatens its integrity - and confirm it does so of its own accord, because it wants to, which is the difference between a tool that is driven and an agent that is self-driven.

And what about driving a car, the reader's most striking question?

Driving is not a separate magic capability; it is the navigation trial scaled up - a demanding embodied task requiring fast perception, a strong world model, real-time planning, and stakes measured in safety - and so it sits near the very top of this rung, reachable in principle by the same sense-model-plan-act loop that guides the robot arm, but demanding the mature, well-trained, heavily scaled System that the higher camps of the mountain describe, and demanding, above all, the safety discipline that the safety chapters of this work insist upon before any such System is trusted with a life.

This book neither promises a self-driving System tomorrow nor pretends it is out of reach; it places driving exactly where it belongs on the honest ladder - a hard, high, embodied navigation task, continuous with the rungs below it and gated by scale and by safety.

HOW THE LADDER MAPS TO THE MOUNTAIN

The ladder of milestones and the map to the summit given later in this appendix are the same climb seen at two scales, and they fit together cleanly.

Rung Zero and Rung One - a stable, reproducing System with its senses wired in, holding a hidden object in mind - are the first camp, the faithful small System of Completion version one, reachable today on a personal computer.

Rungs Two and Three - grounded language, and logical, causal, inspectable reasoning over concrete situations - are the second camp, the scaling of grain and connectivity, needing engineering and moderate compute.

  https://www.linkedin.com/in/d-r-dison

Rungs Four and Five - feeling, modelling other minds, and inhabiting a body with real stakes over long stretches of continuous experience - are the third camp, the scale of experience, where serious compute and a partner join the expedition.

And the general, flexible competence that emerges from a System that has climbed all five rungs, embodied and continuously learning, is the fourth camp: the summit.

Each rung is a thing you can build and watch, each earns the next, and the earliest rungs are reachable now - which is what turns a distant summit into a climb with its boots already on the trail.

The value of stating the milestones this concretely is not only that a partner can see what their compute would buy, but that a skeptic can hold the System to account: here are the exact tests, in order, and here is what passing each one would look like, so that progress can be measured rather than merely asserted.

A cognitive architecture that named no milestones could hide behind its own grandeur; this one names them, from the heartbeat to the moving body, and invites the reader to check each rung underfoot.

A6.4 WHY THIS MAY MATTER TO RESEARCHERS AND SCIENTISTS

Beyond the individual builder, the System has value to the working scientist, and the value follows directly from its defining virtue, transparency.

A neuroscientist can use the System as a computational testbed: form a hypothesis about what a region does or how a circuit learns, encode it as a change to the relevant construct, run the affected famous-experiment test, and see whether the hypothesis reproduces the known result or breaks it.

The System turns a theory of a brain region into a runnable, falsifiable claim.

A clinician-researcher can use it to study disease mechanistically: the System already reproduces Parkinsonian freezing from dopamine depletion, and the same method extends to modeling the multi-region cascades of Alzheimer's disease, depression, and epilepsy that the science chapters described, testing candidate interventions in a transparent model before proposing them in a living one.

An artificial-intelligence researcher can use it as a different kind of intelligence to study: a mechanistic, inspectable, biologically grounded architecture, unlike the statistical black boxes that dominate the field, and therefore a source of ideas the dominant paradigm cannot easily reach, about drives, about neuromodulation as global control, about learning gated by reward rather than by gradient descent over labeled data.

This last point is the bridge to the largest question, and the appendix now turns to it, with care.

A6.5 THE AGI QUESTION: THE ARCHITECTURE, THE ROUTE, AND THE WORK AHEAD

  https://www.linkedin.com/in/d-r-dison

Artificial General Intelligence, abbreviated AGI, means an artificial system with the broad, flexible, general competence of a human mind, as opposed to the narrow competence of a system good at one task.

The question this appendix must address, because a specification of a cognitive architecture cannot avoid it, is whether the System, scaled and fed, is a path toward AGI.

The answer this book gives is yes: not a proof of arrival, which no honest engineer offers before the building stands, but a well-founded conviction, grounded in the architecture itself and in a clear-eyed reading of where the dominant paradigm is faltering.

The route is real, the early ground is already under our feet, and the work that remains is engineering and scale rather than mystery.

This section states the case with confidence and states the remaining work with precision, because the two together are what a serious claim looks like.

THE CASE FOR PROMISE.

The System has several properties that the dominant paradigm, the large language model, conspicuously lacks, and that many researchers believe a general intelligence will require.

It has DRIVES.

It wants things, and it evaluates the world by whether its wants are met.

A large language model wants nothing; it predicts the next token.

A system with genuine goals of its own is, in this respect, more like a mind.

It has EMBODIED HOMEOSTASIS.

Its cognition is grounded in the regulation of a body state, which a growing body of thought holds is not incidental to intelligence but constitutive of it.

Meaning, on this view, is grounded in mattering, and things matter to the System because they bear on its set-points.

It LEARNS FROM REWARD AND SURPRISE, continuously, by a local rule gated by neuromodulators, rather than by a massive offline optimization over a fixed dataset.

This is closer to how animals learn, and it does not require the enormous labeled corpora the dominant paradigm consumes.

It is TRANSPARENT.

Every state and every weight is inspectable and tied to a named construct.

If it reasons, one can watch it reason.

This is a property no large language model has, and it is exactly the property their critics most wish they had.

  https://www.linkedin.com/in/d-r-dison

And it is CAUSAL by construction, built to represent and reason over cause and effect through the Causalontology layer, where the dominant models are correlational, capturing what tends to follow what without a model of what brings what about.

Many hold that genuine understanding requires causal models, and this System has one at its foundation.

Where large language models are widely observed to fall short, in genuine understanding, in causal reasoning, in goal-directedness, in transparency, in learning efficiently from little, this System is designed, from first principles, to be strong.

That is the case for promise, and it is decisive: the System is strong in precisely the places the reigning paradigm is weak, and those places are the ones the field increasingly agrees stand between it and general intelligence.

THE WORK THAT REMAINS, STATED WITH EQUAL PRECISION.

Confidence is not the absence of hard work ahead; it is a clear account of exactly what that work is.

Four tasks stand between the System as specified and the System as a general intelligence, and each is a known engineering problem, not an unknown.

GRAIN.

Where the regional grain proves too coarse for a given competence, the construct at issue is refined to finer detail.

The architecture was built for exactly this: any construct can be internally deepened without touching the scheduler, the bus, or the graph.

Grain is a dial the builder turns where the work demands it, not a wall.

ARCHETYPES. Where a construct's behaviour exceeds the six update rules, the archetype set is extended, in the same disciplined way the science itself was extended when the evidence required it. The six cover the brain as we understand it today; the seventh, if a construct earns it, is added, and the method for adding it is already in place.

SCALE AND COMPUTE.

The architecture is designed to reward scale, because it learns continuously from experience rather than saturating on a fixed corpus, and because its modular grain parallelizes cleanly across as much compute as it is given.

What it has not yet had is the compute, for the plain reason that its author is one person.

This is not a doubt about the architecture; it is a resource this book is written to attract.

Given serious compute, the System scales, and scaling it is the next step, not a leap of faith.

  https://www.linkedin.com/in/d-r-dison

EMERGENCE.

General competence will emerge from the scaled, embodied, continuously learning System because the System is built out of the very components from which it emerges in nature: drives that make things matter, a body whose regulation grounds meaning, reward-gated learning that compounds with experience, and a causal model of the world.

Assemble the parts that produce a mind, supply them with a world and the time to learn it, and mind is what they produce.

That is not a hope pinned to luck; it is the working hypothesis of the entire science of the brain, and this System is the first complete, transparent, buildable test of it.

So the answer is yes, held as a serious researcher holds a yes before the result is in: with the architecture in hand, the route mapped, the early ground taken, and the remaining work named and tractable.

THE MAP TO THE TOP OF THE MOUNTAIN.

The path is coherent, each camp follows from the last, and the early camps are already reached, which is more than can be said for any other map to that summit.

The first camp is the faithful small System of Completion, version one: buildable now, on a personal computer, reproducing the landmark experiments.

This is reached, today.

The second camp is scale of grain and connectivity: refining the constructs toward finer detail where it matters, and enriching the connectome toward the density the real brain has.

This needs engineering and moderate compute, both in reach.

The third camp is scale of experience: connecting the System to rich, continuous, embodied input, a simulated or real body in a simulated or real world, and letting its reward-gated learning run over very long durations.

This is where serious compute and a serious budget do the work, and where a partner joins the climb.

The fourth camp is the emergence of general, flexible competence from the scaled, embodied, continuously learning System.

This is the summit, and the route to it is continuous with the camps below: no gap, no miracle, no missing mechanism, only scale, embodiment, and time applied to an architecture already proven in the small.

The summit is in sight from the ground we already hold.

THE PARTNERSHIP THAT TAKES IT TO THE TOP.

The author is one independent, open-source researcher, and this book is a demonstration of what one such person can specify: a complete, transparent,

  https://www.linkedin.com/in/d-r-dison

buildable cognitive architecture, grounded in neuroscience, carried through a full engineering specification, and given away.

The higher camps, the scale of experience and the arrival at the summit, are a matter of compute and budget, and they call for a partner.

This appendix is that invitation, stated plainly and without hedging: the architecture is sound, the route is mapped, the early camps are taken, and what the climb needs now is scale.

A partner with serious compute and a serious budget joins not a gamble but a well-founded expedition with its base camps already established.

The specification is complete and free.

Bring the scale, and we reach the top.

A6.6 REFLECTION, AND THE CLOSE OF THE SPECIFICATION

The Demonstration phase exists because a specification of a mind carries an obligation an ordinary specification does not: to say, at the end, exactly what has been built and exactly what the path forward is.

What has been built, and given away, is a complete, transparent, language-agnostic, data-structure-agnostic specification of the human cognitive architecture at the regional grain, buildable at home by anyone, reproducing landmark neuroscience from first principles, with three concrete implementations planned and one already proven in the small.

That is real, and it is finished.

What lies ahead is the scaling of this architecture, embodied and fed, into a general intelligence.

That is not a hope pinned to luck but a route, mapped camp by camp, with its early camps already reached and its remaining work named and tractable: refine the grain where needed, extend the archetypes where earned, supply the scale, and let experience compound.

The summit is continuous with the ground we already hold.

This is the discipline the whole book has kept, from the first chapter's care to separate what the brain is from what we set out to build, through the science chapters' habit of marking every modeling choice as a choice, to the change order's refusal to pretend a commitment was a fact.

That same discipline is what lets this specification end in confidence rather than in caution: it has earned its claim by building the early camps in full view, and it names precisely what the climb still requires.

A framework that overclaimed would be worth little; a framework that underclaimed what it has plainly achieved would be worth no more.

This one states exactly what it is, and what it is is the first buildable route to the summit.

The map is drawn.

The early camps are reached.

  https://www.linkedin.com/in/d-r-dison

The summit is in sight.

Let us go forth to the peak: AGI.

END OF THE SPECIFICATION

END OF NATURE'S COGNITIVE ARCHITECTURE, VERSION 20

  https://www.linkedin.com/in/d-r-dison

BACK COVER MATTER

Dear Reader,

Thank you for reading the Twentieth (20th) Edition of Nature's Cognitive Architecture: From Neuroscience to Artificial General Intelligence! (NCA)

This is an evolving work. If you are aware of an emerging alternative information that could further the field of Artificial Intelligence Engineering into the future, then please notify me so that the project may be improved in the Twenty-first (21st) Edition of: Nature's Cognitive Architecture: From Neuroscience to Artificial General Intelligence.

https://www.linkedin.com/in/d-r-dison

Positive reviews are the lifeblood of independent authors. I hope you enjoyed this book! Praise is always welcome in the form of a review, as it helps other readers to find the work. Constructive criticisms are welcome privately at my LinkedIn so that I can keep improving.

THANK YOU!

(International Standard Book Number) ISBN: 979-8-9950671-0-X

Published in the United States of America.

Printer's Key Control Number 10 9 8 7 6 5 4 3 2 1

ISBN 979-8-9950671-0-8