A minimal real-time brain architecture

Bill Softky, Telluride Neuromorphic Engineering Conference 2011 (with Kevin Yang)

These notes and diagrams show a very high-level overview of a plausible architecture for real-time aspects of generic, brain-like sensory processing, including distributed active sensation and control. Because I personally believe that we know far more about the generic mathematical constraints on sensory input, processing, and distributed computation than we do about neural micro-anatomy (hypercolumns, neurons, dendrites, spines), this approach is explicitly **non-neural**: it relies on what *must* be true of brain processing, not on what we think we have measured.

The basic idea : many interconnected algorithmic modules, each of which (as in the first picture) should somehow perform the same core functions across space and time: compress, expand, predict, and control. Exactly how each module does this is "left to the reader"; an initial illustration (minus the de-reverb and control pieces) is outlined in the companion presentation "Generic Sensory Prediction," but I'm sure many people could fill in the details much better.

Each such "Compressor" module should be responsible for its own learning, resource management, and execution. It accepts inputs from its "south gateway": a routing interface which concatenates input from raw sensors and/or other modules and returns to them feedback both about predictions of future inputs and about demands or requests for modified future input.

In effect, except for a few control signals sent from each Compressor directly to muscles/actuators, most muscle-control signals are sent in the language of **desired input** ("go to this point") rather than in the language of actuators ("move this muscle"). The core principle both for learning the muscles-to-state-space control mapping, and for using already-learned inputs to maintain active control in real time, is that of **optimal dithering**: inject just the right pseudo-random signals into the control stream so that reverse-correlated reconstruction can distinguish input changes caused by **this** module from changes caused by other modules or by the outside world. In effect, dithering allows each of many modules to "label" its outputs by its own independent temporal variability, so that it can identify and learn from their effects.

Likewise, each Compressor digests the current best-estimate state-space location and trajectory into a minimal form (e.g. *{timestamp, state, direction}*), which can be distributed event-style from the North Gateway to other modules' South Gateways for state-estimation at arbitrary points in the near future. In this scheme each module knows about and operates only on 10-100 input dimensions at a time, and the Gateways have full responsibility for managing direct inter-module connections and event-driven updates. (Not yet shown, or even understood, is how each module knows which *muscles* to connect to….unfortunately, muscles do not necessarily share the same grouping as the sensors they affect).

Many such modules should be wired together hierarchically, as on the left side of the second diagram. After learning, the topmost level should contain modules which operate on the most abstract, multi-sensory, long-time-scales possible, while the lowermost modules deal with near-term primary sensory input and effector output. The bottom-most layer of the network could contain a predictive-comparator filter, to maximize processing efficiency by removing predictable, redundant inputs (see *Modeling Thalamus as a non-rectifying predictive comparator*).

The "output" of such a scheme not only includes specific, high-dimensional, high-bandwidth predictions of raw sensory signals (as proposed in my 1995 NIPS paper *Unsupervised Pixel-prediction,* in Jeff Hawkins' 2004 book *On Intelligence*, and probably other places as well). That output also implements active perception and control of the effectors associated with sensory input, like creating head movements and saccades for visual input, balance and gait management for the feet, palpation for hands/lips/tongue, proprioceptive feedback from skeletal muscles (by dithering muscle-tremors and/or muscle-fiber spikes), and perhaps even acoustic enhancement in the cochlea. A strong prediction of this model is that active control should exist in every sensory modality possible, at the highest bandwidths possible.

Each individual module on the left is born as a generic tabula-rasa, initially agnostic about both sensory modality and level of abstraction. But an actual brain, to survive in the world, requires at least a few hints about desirable inputs and situations, so it can know what to pursue or avoid. Those hints come from two sources, one of them obscure. The obvious, well-known reward/aversion inputs are themselves sensory inputs like pain, temperature, sugar, and such; those inputs drive a pre-determined suite of "cognitive/behavioral eigenstates," a shorthand for various preset resource- and processing-management strategies (e.g. "get food", "explore", "hide," "investigate"), a concept broadly similar to animal "moods" and "attitudes." Because these strategies are global and typically change slowly, they can be implemented by sending global prioritization-signals to large numbers of Compressors in parallel, telling each how to do its job best in the moment.

But many crucial situations are *not* coded by specific pain/pleasure receptors, and obviously cannot be coded at birth by intermediate Compressors whose mappings haven't yet even been learned. The only other source for such assessments would be at the very *topmost* level of "trans-sensory statistics, " which would reflect global statistical answers to generic questions like "am I in control," "is this desirable", "is the situation improving," or "am I safe." Any such global "reward" signal would clearly complement the raw sensory rewards we already know about.

This tentative outline—not yet even prototyped!—illustrates the integration of several disparate tricks: dimensionality reduction, distributed processing, distributed control via optimal dithering, trans-sensory reward signals, and learning to predict the future. Because all these functions MUST be taking place in even small brains, we can focus on their simplest-possible interactions before we add human-scale sophistication like episodic memory, strategy, language, and consciousness.

Compress

dim-reduction
de-noise
generalize
de-reverb
segment
estimate/infer
temporal rescaling

Expand

pseudo-invert
sparsify
superpose

Predict

extrapolate
validate

Control

dither
invert impulse-response
pursue
delegate
execute
confirm

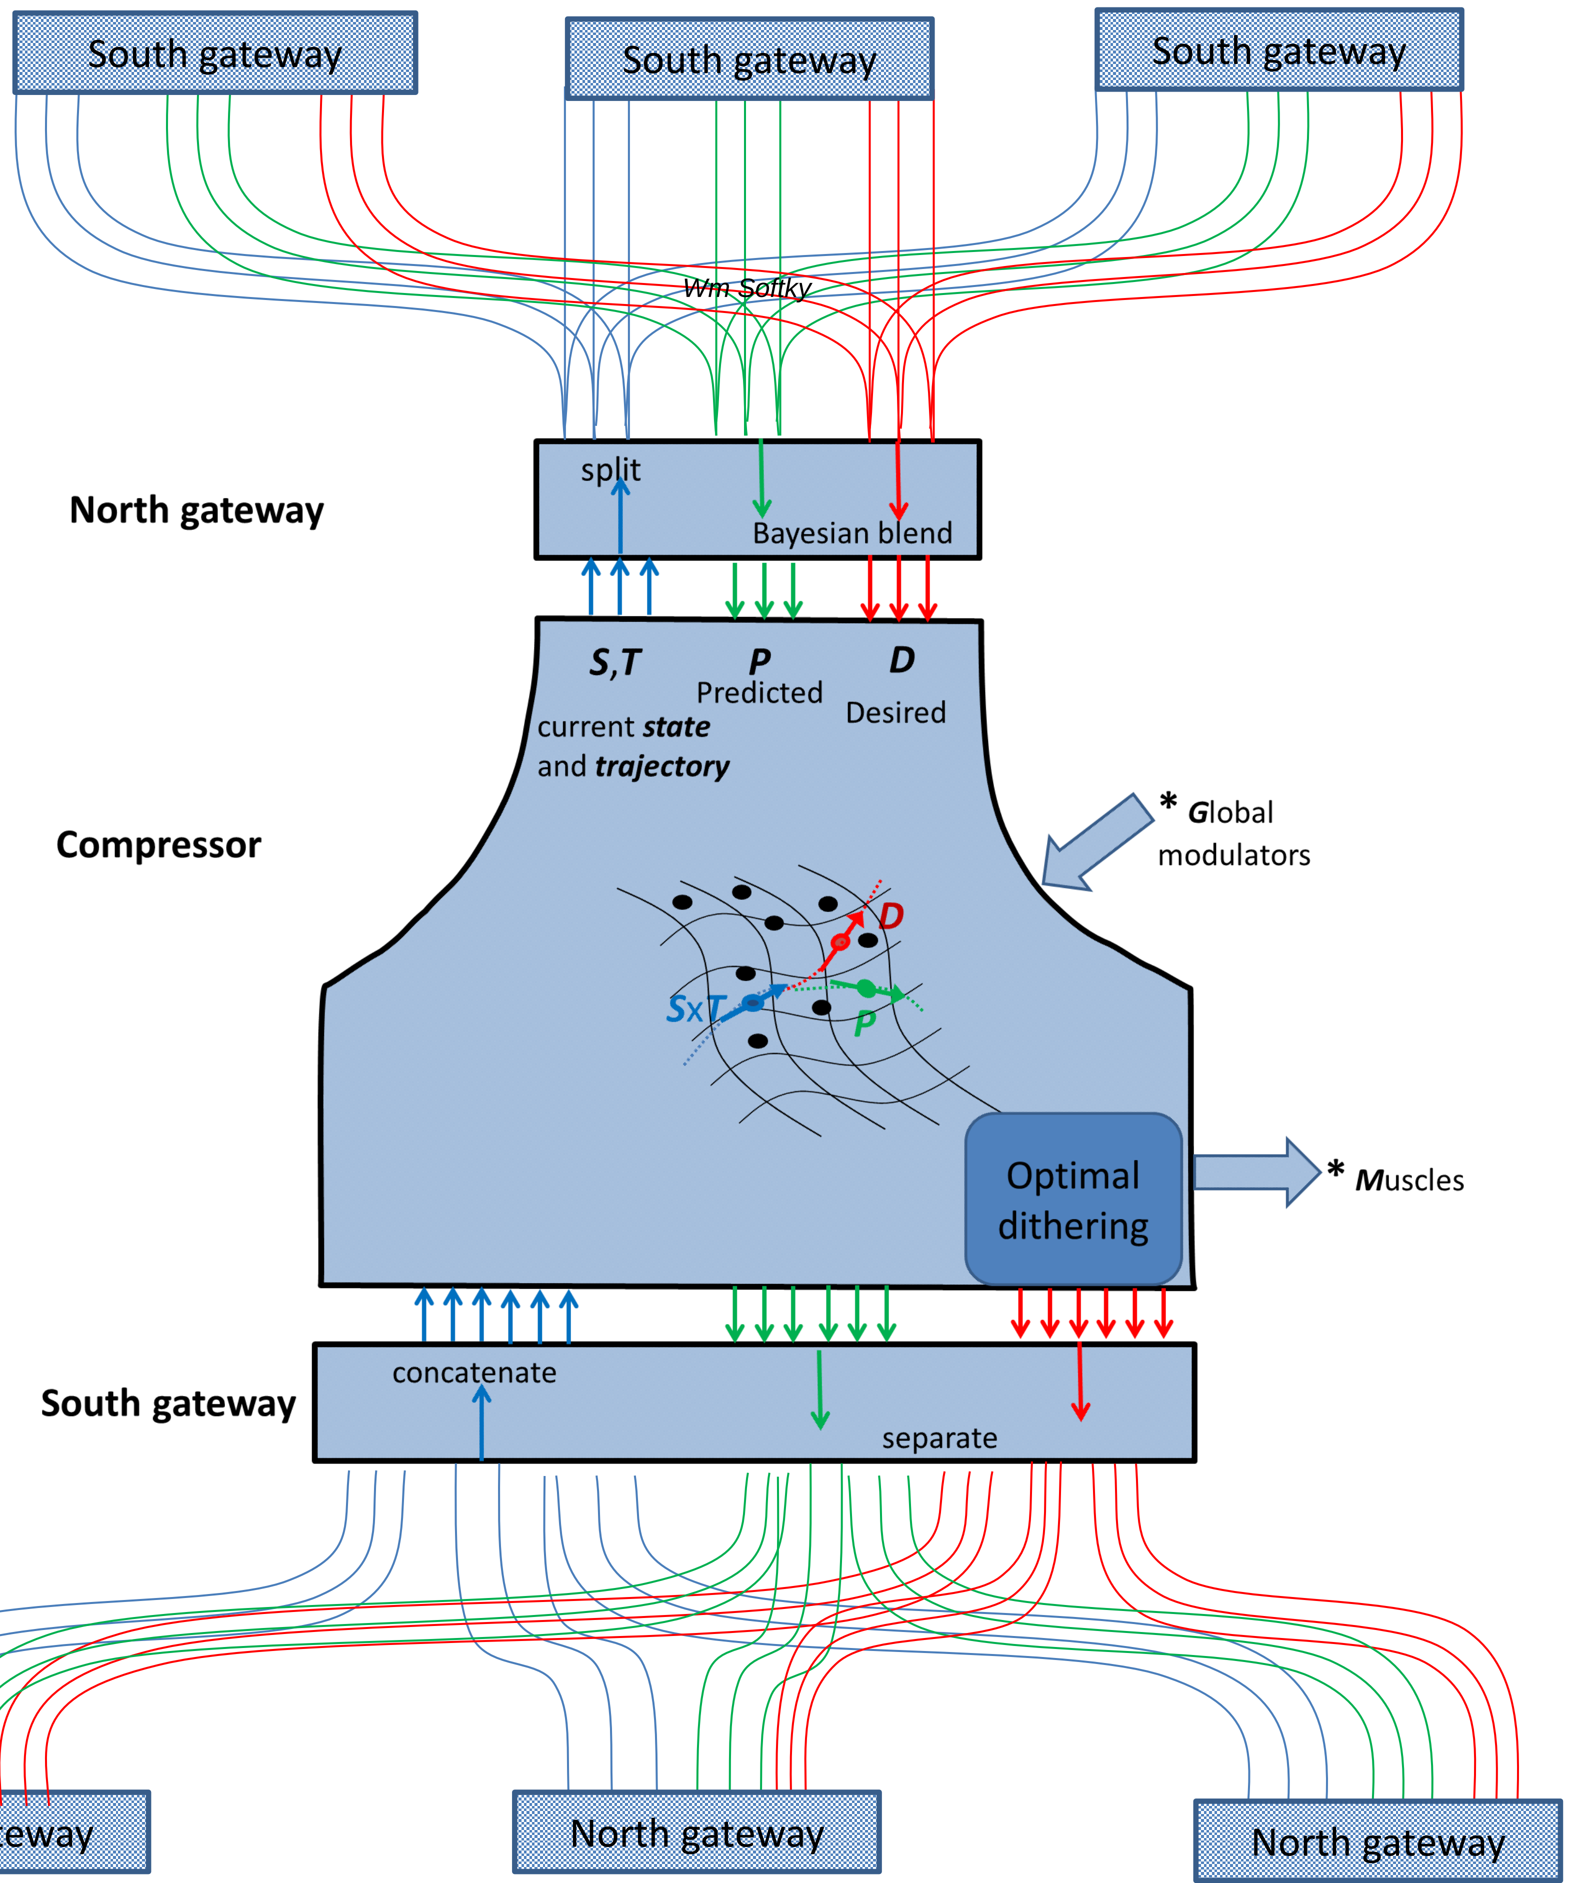

1. Map manifold “pearls”

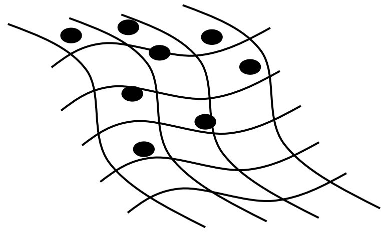

2. Map manifold sequences

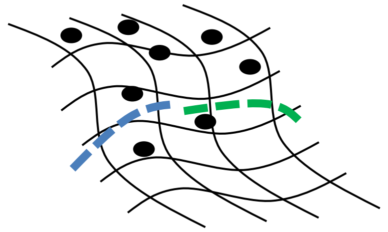

3. Extrapolate sequences

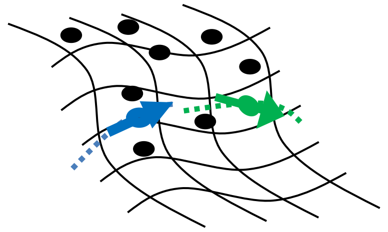

4. Publish (S,T) northward

5. Incorporate predictions

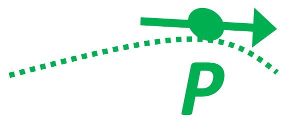

6. Dither to learn control

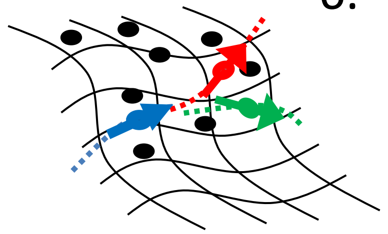

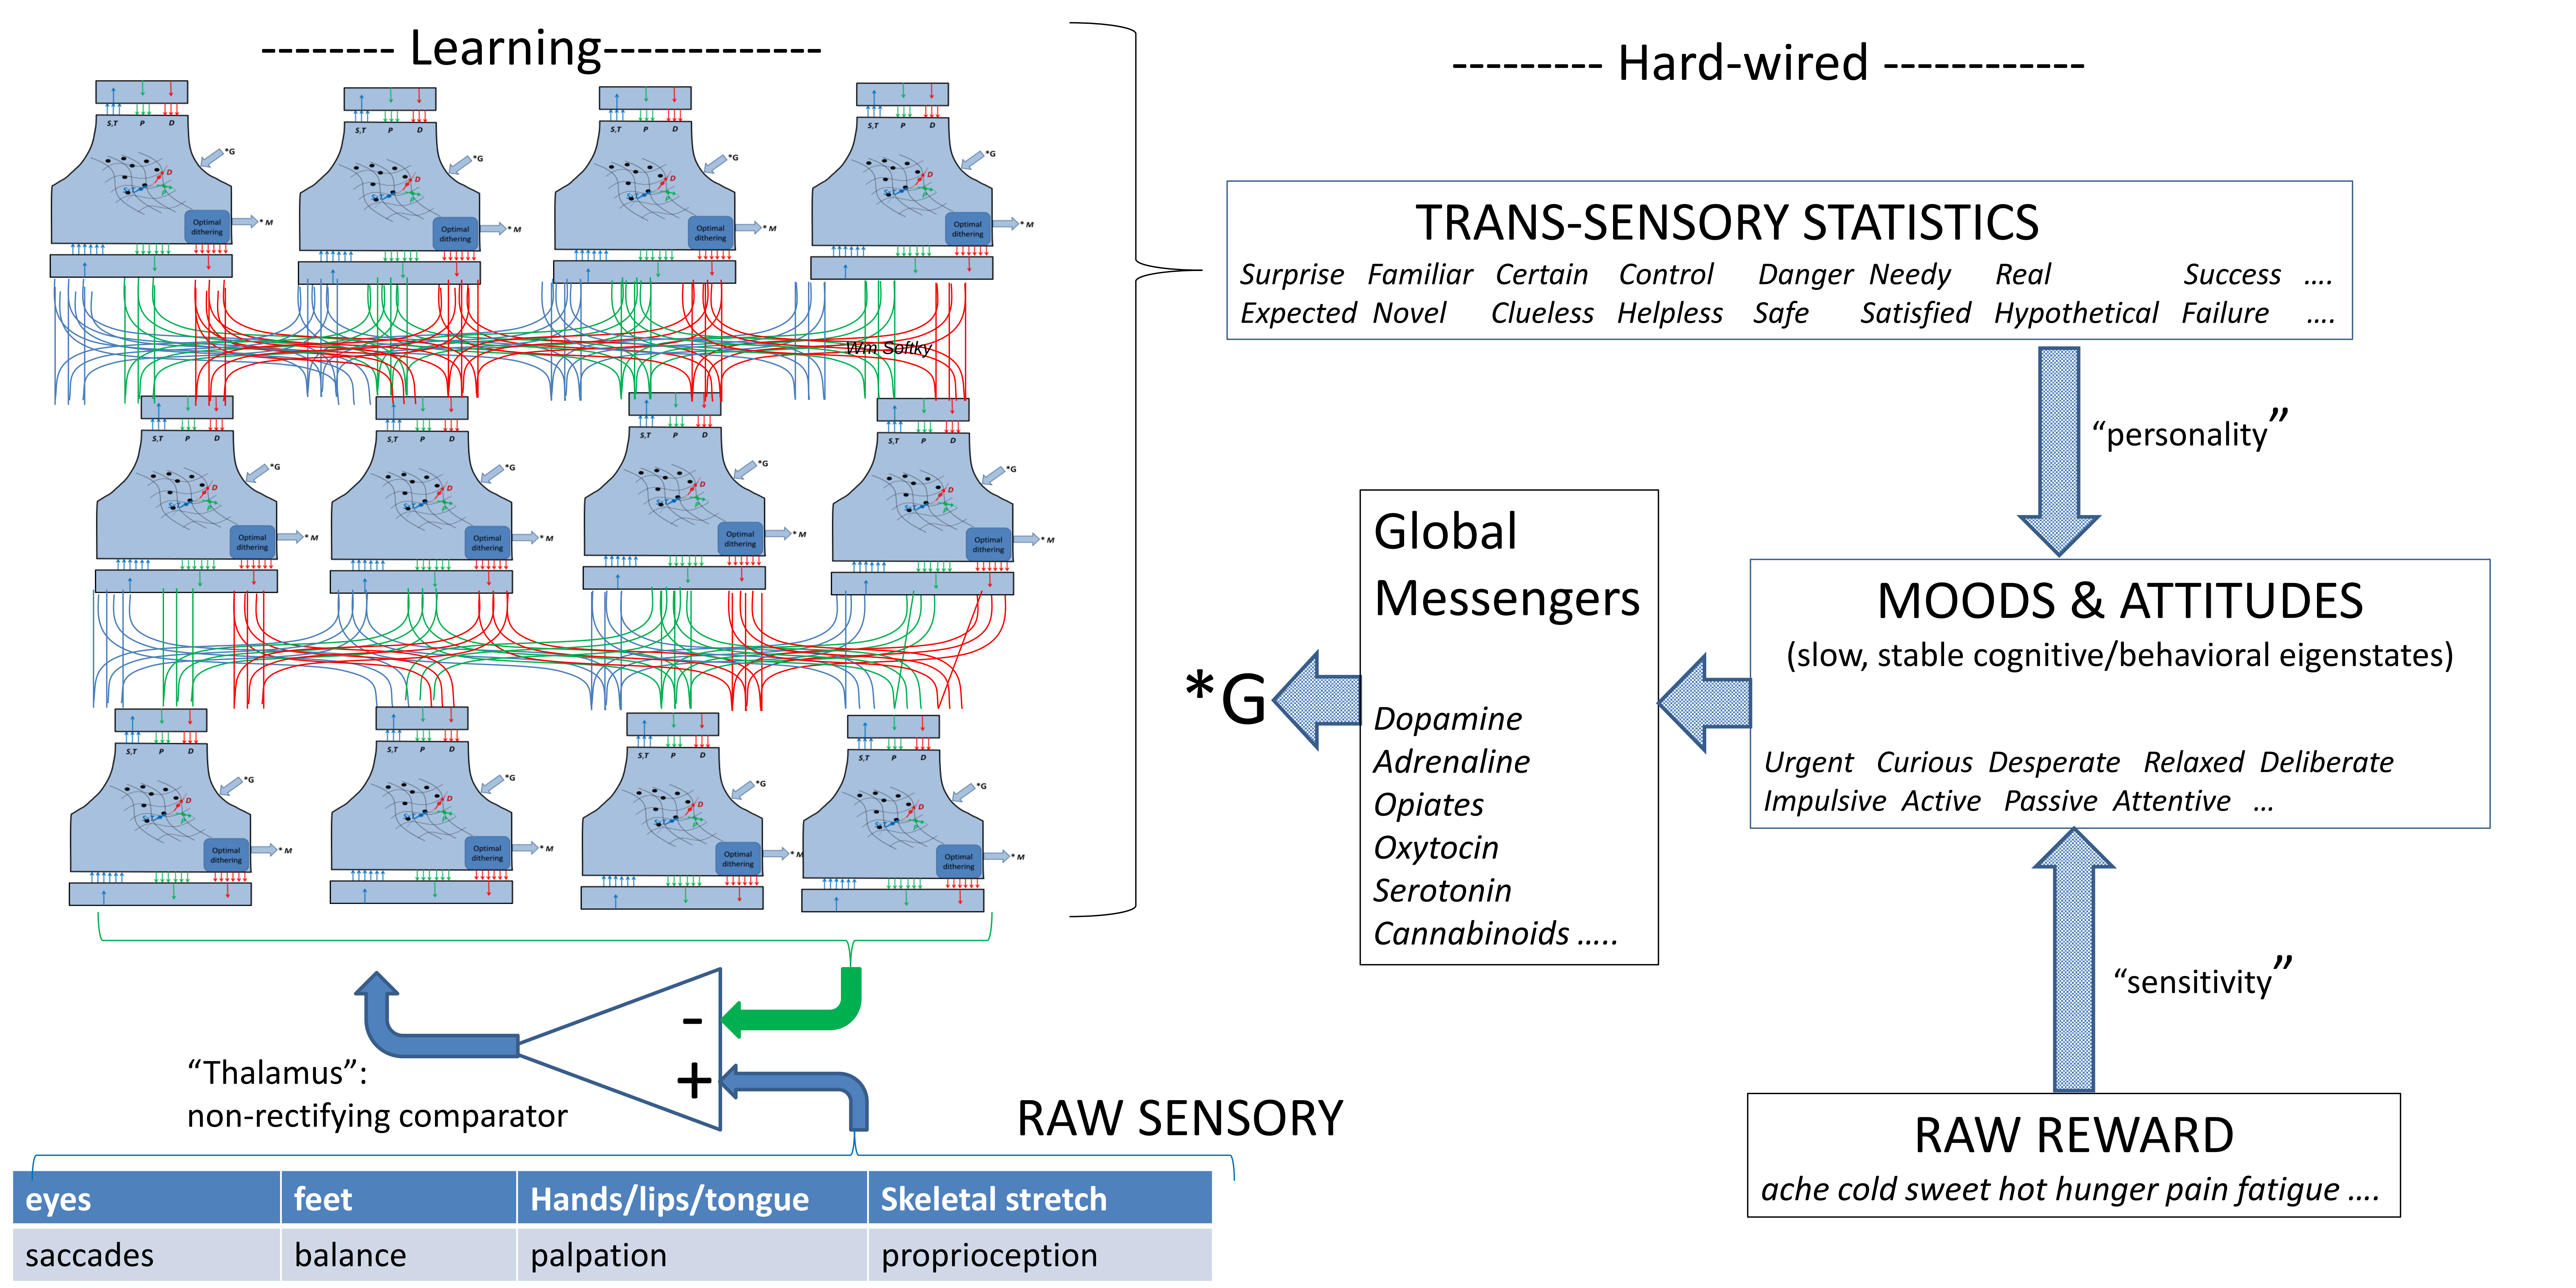

-------- Learning------------
--------- Hard-wired -----------
Optimal dithering
*G
*M
Wm Softky
TRANS-SENSORY STATISTICS
Surprise Familiar Certain Control Danger Needy Real Success ....
Expected Novel Clueless Helpless Safe Satisfied Hypothetical Failure ....
"personality"
Global Messengers
Dopamine
Adrenaline
Opiates
Oxytocin
Serotonin
Cannabinoids .....
*G
MOODS & ATTITUDES
(slow, stable cognitive/behavioral eigenstates)
Urgent Curious Desperate Relaxed Deliberate
Impulsive Active Passive Attentive ...
"sensitivity"
RAW REWARD
ache cold sweet hot hunger pain fatigue ....
-
+
"Thalamus":
non-rectifying comparator
RAW SENSORY
eyes
saccades
feet
balance
Hands/lips/tongue
palpation
Skeletal stretch
proprioception