Opentheory.net

Principles of Vasocomputation: A Unification of Buddhist Phenomenology, Active Inference, and Physical Reflex (Part I)

Posted on July 12, 2023 by Michael Edward Johnson

A unification of Buddhist phenomenology, active inference, and physical reflexes; a practical theory of suffering, tension, and liberation; the core mechanism for medium-term memory and Bayesian updating; a clinically useful dimension of variation and dysfunction; a description of sensory type safety; a celebration of biological life.

Michael Edward Johnson, Symmetry Institute, July 12, 2023.

I. What is tanha?

By default, the brain tries to grasp and hold onto pleasant sensations and push away unpleasant ones. The Buddha called these 'micro-motions' of greed and aversion *taṇhā*, and the Buddhist consensus seems to be that it accounts for an amazingly large proportion (~90%) of suffering.

Romeo Stevens suggests translating the original Pali term as "fused to," "grasping," or "clenching," and that the mind is trying to make sensations feel stable, satisfactory, and controllable. Nick Cammarata suggests "fast grabby thing" that happens within ~100ms after a sensation enters awareness; Daniel Ingram suggests this 'grab' can occur as quickly as 25-50ms (personal discussion). Uchiyama Roshi describes tanha in terms of its cure, "opening the hand of thought"; Shinzen Young suggests "fixation"; other common translations of tanha are "desire," "thirst," "craving." The vipassana doctrine is that tanha is something the mind instinctively does, and that meditation helps you see this process as it happens, which allows you to stop doing it. Shinzen estimates that his conscious experience is literally 10x better due to having a satisfying meditation practice.

Tanha is not yet a topic of study in affective neuroscience but I suggest it should be. Neuroscience is generally gated by soluble important mysteries: complex dynamics often arise from complex mechanisms, and complex mechanisms are difficult to untangle. The treasures in neuroscience happen when we find *exceptions* to this rule: complex dynamics that arise from elegantly simple core mechanisms. When we find one it generally leads to breakthroughs in both theory and intervention. Does "tanha" arise from a simple or complex mechanism? I believe Buddhist phenomenology is very careful about what it calls dependent origination — and this makes items that Buddhist scholarship considers to be 'basic building-blocks of phenomenology' particularly likely to have a simple, elegant implementations in the brain — and thus are exceptional mysteries to focus scientific attention on.

I don't think tanha has 1000 contributing factors; I think it has one crisp, isolatable factor. And I think if we find this factor, it could herald a reorganization of systems neuroscience similar in magnitude to the past shifts of cybernetics, predictive coding, and active inference.

Core resources:

1. Anuruddha, Ā. (n.d.). A Comprehensive Manual of Abhidhamma.
2. Stevens, R. (2020). (mis)Translating the Buddha. Neurotic Gradient Descent.
3. Cammarata, N. (2021-2023). [Collected Twitter threads on tanha].
4. Markwell, A. (n.d.). Dhamma resources.

II. Tanha as unskillful active inference (TUAI)

The first clue is what tanha is trying to do for us. I'll claim today that tanha is a side-effect of a normal, effective strategy our brains use extensively, active inference. Active inference suggests we impel ourselves to action by first creating some predicted sensation ("I have a sweet taste in my mouth" or "I am not standing near that dangerous-looking man") and then holding it until we act in the world to make this prediction become true (at which point we can release the tension). Active inference argues we store our to-do list as predictions, which are equivalent to untrue sensory observations that we act to make true.

Formally, the "tanha as unskillful active inference" (TUAI) hypothesis is that this process commonly goes awry (i.e. is applied unskillfully) in three ways:

- First, the rate of generating normative predictions can outpace our ability to make them true and overloads a very finite system. Basically we try to control too much, and stress builds up.

- Second, we generate normative predictions in domains that we cannot possibly control; predicting a taste of cake will linger in our mouth forever, predicting that we did not drop

our glass of water on the floor. That good sensations will last forever and the bad did not happen. (This is essentially a "predictive processing" reframe of the story Romeo Stevens has told on his blog, Twitter, and in person.)[1]

- Third, there may be a context desynchronization between the system that represents the world model, and the system that maintains predictions-as-operators on this world model. When desynchronization happens and the basis of the world model shifts in relation to the basis of the predictions, predictions become nonspecific or nonsensical noise and stress.

- We may also include a catch-all fourth category for when the prediction machinery becomes altered outside of any semantic context, for example metabolic insufficiency leading to impaired operation.

Core resources:

5. Safron, A. (2020). An Integrated World Modeling Theory (IWMT) of Consciousness: Combining Integrated Information and Global Neuronal Workspace Theories With the Free Energy Principle and Active Inference Framework; Toward Solving the Hard Problem and Characterizing Agentic Causation. Frontiers in Artificial Intelligence, 3. https://doi.org/10.3389/frai.2020.00030
6. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G. (2017). Active inference: A Process Theory. Neural Computation, 29(1), 1-49.
7. Sapolsky, R.M. (2004). Why Zebras Don't Get Ulcers: The Acclaimed Guide to Stress, Stress-Related Diseases, and Coping. Holt Paperbacks. [Note: link is to a video summary.]
8. Pyszczynski, T., Greenberg, J., Solomon, S. (2015). Thirty Years of Terror Management Theory. Advances in Experimental Social Psychology, 52, 1-70.

III. Evaluating tanha requires a world model and cost function

There are many theories about the basic unit of organization of the brain; brain regions, functional circuits, specific network topologies, etc. Adam Safron describes the nervous system's basic building block as Self-Organized Harmonic Modes (SOHMs); I like this because the math of harmonic modes allows a lot of interesting computation to arise 'for free.' Safron suggests these modes function as autoencoders, which I believe are functionally identical to symmetry detectors. It's increasingly looking like SOHMs are organized around physical brain resonances at least as much as connectivity, which been a surprising result.

At high frequencies these SOHMs will act as feature detectors, at lower frequencies we might think of them as wind chimes: by the presence and absence of particular SOHMs and their interactions we obtain a subconscious feeling about what kind of environment we're in and where its rewards and dangers are. We can expect SOHMs will be arranged in a way that optimizes differentiability of possible/likely world states, minimizes crosstalk, and in aggregate constitutes a world model, or in the Neural Annealing/REBUS/ALBUS framework, a belief landscape.

To be in tanha-free "open awareness" without greed, aversion, or expectation is to feel the undoctored hum of your SOHMs. However, we doctor our SOHMs *all the time* — when a nice sensation enters our awareness, we reflexively try to 'grab' it and stabilize the resonance; when something unpleasant comes in, we try to push away and deaden the resonance. Likewise society puts expectations on us to "act normal" and "be useful"; we may consider all such SOHM adjustments/predictions as drawing from the same finite resource pool. "Active SOHM management" is effortful (and unpleasant) in rough proportion to how many SOHMs need to be actively managed and how long they

need to be managed.

But how can the brain manage SOHMs? And if the Buddhists are right and this creates suffering, why does the brain even try?

Core resources:

9. Safron, A. (2020). An Integrated World Modeling Theory (IWMT) of Consciousness: Combining Integrated Information and Global Neuronal Workspace Theories With the Free Energy Principle and Active Inference Framework; Toward Solving the Hard Problem and Characterizing Agentic Causation. Frontiers in Artificial Intelligence, 3. https://doi.org/10.3389/frai.2020.00030
10. Safron, A. (2020). On the varieties of conscious experiences: Altered beliefs under psychedelics (ALBUS). PsyArxiv. Retrieved July 7, 2023, from the PsyArxiv website.
11. Safron, A. (2021). The radically embodied conscious cybernetic bayesian brain: From free energy to free will and back again. Entropy, 23(6), 783. MDPI.
12. Bassett, D. S., & Sporns, O. (2017). Network neuroscience. Nature Neuroscience, 20(3), 353-364.
13. Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. Science, 304(5679), 1926-1929.
14. Johnson, M. (2016). Principia Qualia. opentheory.net.
15. Johnson, M. (2019). Neural Annealing: Toward a Neural Theory of Everything. opentheory.net.
16. Johnson, M. (2023). Qualia Formalism and a Symmetry Theory of Valence. opentheory.net.
17. Carhart-Harris, R. L., & Friston, K. J. (2019). REBUS and the Anarchic Brain: Toward a Unified Model of the Brain Action of Psychedelics. Pharmacological Reviews, 71(3), 316-344.
18. Dahl, C. J., Lutz, A., & Davidson, R. J. (2015). Reconstructing and deconstructing the self: cognitive mechanisms in meditation practice. Trends in Cognitive Sciences, 19(9), 515-523.

IV. Tanha as artifact of compression pressure

I propose reframing tanha as an artifact of the brain's *compression pressure*. I.e. tanha is an artifact of a continual process that subtly but systematically pushes on the complexity of 'what is' (the neural patterns represented by undoctored SOHMs) to collapse it into a more simple configuration, and sometimes holds it there until we act to make that simplification true. The result of this compression drive conflates "what is", "what could be", "what should be", and "what will be," and this conflation is the source of no end of moral and epistemological confusion.

This reframes tanha as both the pressure which collapses complexity into simplicity, and the ongoing stress that comes from maintaining the counterfactual aspects of this collapse (*compression stress*). We can think of this process as balancing two costs: on one hand, applying compression pressure has metabolic and epistemic costs, both immediate and ongoing. On the other hand, the brain is a finite system and if it doesn't continually "compress away" patterns there will be unmanageable sensory chaos. The right amount of compression pressure is not zero.[2]

Equivalently, we can consider tanha as an excessive forcefulness in the metabolization of uncertainty. Erik P. Hoel has written about energy, information, and uncertainty as equivalent and conserved quantities (Hoel 2020): much like literal digestion, the imperative of the nervous system is to extract value from sensations then excrete the remaining information, leaving a low-information, low-uncertainty, clean slate ready for the next sensation (thank you Benjamin Anderson for discussion). However, we are often unskillful in the ways we try to extract value from sensations, e.g. improperly assessing

context, trying to extract too much or too little certainty, or trying to extract forms of certainty inappropriate for the sensation.

We can define a person's personality, aesthetic, and a large part of their phenomenology in terms of how they metabolize uncertainty — their library of motifs for (a) initial probing, (b) digestion and integration, and (c) excretion/externalization of any waste products, and the particular reagents for this process they can't give themselves and must seek in the world.

So far we've been discussing brain dynamics on the computational level. But *how* does the brain do all this — what is the *mechanism* by which it attempts to apply compression pressure to SOHMs? This is essentially the question neuroscience has been asking for the last decade. I believe evolution has coupled two very different systems together to selectively apply compression/prediction pressure in a way that preserves the perceptive reliability of the underlying system (undoctored SOHMs as ground-truth perception) but allows near-infinite capacity for adjustment and hypotheticals. One system focused on perception; one on compression, judgment, planning, and action.

The traditional neuroscience approach for locating these executive functions has been to associate them with particular areas of the brain. I suspect the core logic is hiding much closer to the action.

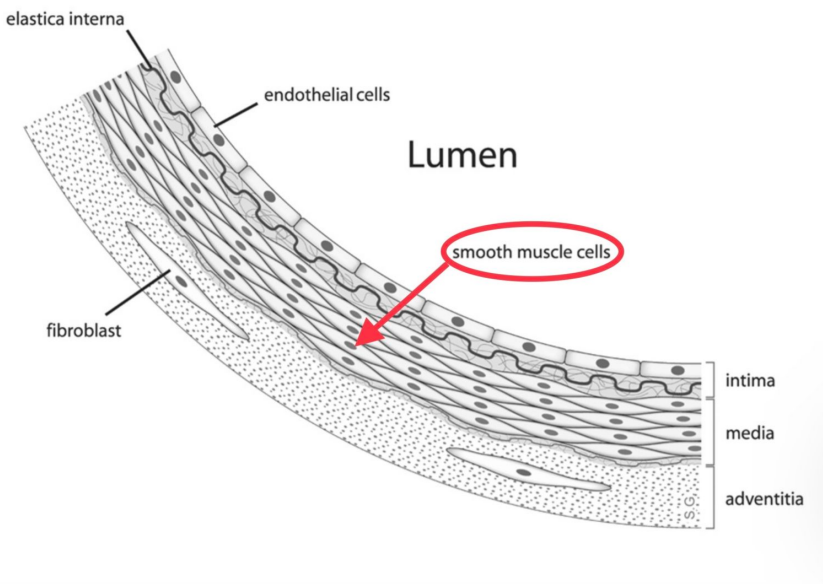
Core resources:

19. Schmidhuber, J. (2008). Driven by Compression Progress: A Simple Principle Explains Essential Aspects of Subjective Beauty, Novelty, Surprise, Interestingness, Attention, Curiosity, Creativity, Art, Science, Music, Jokes. Arxiv.

Retrieved July 7, 2023, from the Arxiv website.

20. Johnson, M. (2023). Qualia Formalism and a Symmetry Theory of Valence. opentheory.net.
21. Hoel, E. (2020). The Overfitted Brain: Dreams evolved to assist generalization. Arxiv. Retrieved July 7, 2023, from the Arxiv website.
22. Friston, K. (2010). The free-energy principle: a unified brain theory? Nature Reviews Neuroscience, 11(2), 127-138.
23. Chater, N., & Vitányi, P. (2003). Simplicity: a unifying principle in cognitive science? Trends in Cognitive Sciences, 7(1), 19-22.
24. Bach, D.R., & Dolan, R.J. (2012). Knowing how much you don't know: a neural organization of uncertainty estimates. Nature Reviews Neuroscience, 13(8), 572-586.

V. VSMCs as computational infrastructure



Above: the vertical section of an artery wall (Wikipedia, emphasis added; video): the physical mechanism by which we grab sensations and make predictions; the proximate cause of 90% of suffering and 90% of goal-directed behavior.

All blood vessels are wrapped by a thin sheathe of vascular smooth muscle cells (VSMCs). The current scientific consensus has the vasculature system as a spiderweb of ever-narrower channels for blood, powered by the heart as a central pump, and supporting systems such as the brain, stomach, limbs, and so on by bringing them nutrients and taking away waste. The sheathe of muscle wrapped around blood vessels undulates in a process called "vasomotion" that we think helps blood keep circulating, much like peristalsis in the gut helps keep food moving, and can help adjust blood pressure.

I think all this is true, but is also a product of what's been easy to measure and misses 90% of what these cells do.

Evolution works in layers, and the most ancient base layers often have rudimentary versions of more specialized capacities (Levin 2022) as well as deep control hooks into newer systems that are built around them. The vascular system actually predates neurons and has co-evolved with the nervous system for hundreds of millions of years. It also has mechanical actuators (VSMCs) that have physical access to all parts of the body and can flex in arbitrary patterns and rhythms. It would be extremely surprising if evolution didn't use this system for something more than plumbing. We can also "follow the money"; the vascular system controls the nutrients and waste disposal for the neural system and will win in any heads-up competition over co-regulation balance.

I expect VSMC contractions to influence nearby neurons through e.g. ephaptic coupling, reducing blood flow, and adjusting local physical resonance, and to be triggered by local dissonance in the electromagnetic field.

I'll offer three related hypotheses about the computational role of VSMCs[3] today that in aggregate constitute a

neural regulatory paradigm I'm calling *vasocomputation*:

1. **Compressive Vasomotion Hypothesis (CVH)**: the vasomotion reflex functions as a compression sweep on nearby neural resonances, collapsing and merging fragile ambivalent patterns (the "Bayesian blur" problem) into a more durable, definite state. Motifs of vasomotion, reflexive reactions to uncertainties, and patterns of tanha are equivalent.
2. **Vascular Clamp Hypothesis (VCH)**: vascular contractions freeze local neural patterns and plasticity for the duration of the contraction, similar to collapsing a superposition or probability distribution, clamping a harmonic system, or pinching a critical network into a definite circuit. Specific vascular constrictions correspond with specific predictions within the Active Inference framework and function as medium-term memory.
3. **Latched Hyperprior Hypothesis (LHH)**: if a vascular contraction is held long enough, it will engage the latch-bridge mechanism common to smooth muscle cells. This will durably 'freeze' the nearby circuit, isolating it from conscious experience and global updating and leading to a much-reduced dynamical repertoire; essentially creating a durable commitment to a specific hyperprior. The local vasculature will unlatch once the prediction the latch corresponds to is resolved, restoring the ability of the nearby neural networks to support a larger superposition of possibilities.

The initial contractive sweep jostles the neural superposition of interpretations into specificity; the contracted state temporarily freezes the result; if the contraction is sustained, the latch bridge mechanism engages and cements this freeze as a hyperprior. With one motion the door of possibility slams shut. And so we collapse our world into something less magical but more manageable, one clench at a time. *Tanha is cringe.*

The claim relevant to the Free Energy Principle – Active Inference paradigm is we can productively understand the motifs of smooth muscle cells (particularly in the vascular system) as "where the brain's top-down predictive models are hiding," which has been an open mystery in FEP-AI. Specific predictions are held as vascular tension, and vascular tension in turn is released by action, consolidated by Neural Annealing, or rendered superfluous by neural remodeling (hold a pattern in place long enough and it becomes the default). Phrased in terms of the Deep CANALs framework which imports ideas from machine learning: the neural weights that give rise to SOHMs constitute the learning landscape, and SOHMs+vascular tension constitute the inference landscape.

The claim relevant to Theravada Buddhism is we can productively understand the motifs of the vascular system as the means by which we attempt to manipulate our sensations. Vasomotion corresponds to an attempt to 'pin down' a sensation (i.e. tanha); muscle contractions freeze patterns; smooth muscle latches block out feelings of possibility and awareness of that somatic area. Progress on the contemplative path will correspond with both using these forms of tension less, and needing them less. I expect cessations to correspond with a nigh-complete absence of vasomotion (and EEG may measure vasomotion moreso than neural activity).

The claim relevant to practical health is that smooth muscle tension, especially in VSMCs, and especially latched tension, is a system science knows relatively little about but is involved in an incredibly wide range of problems, and understanding this system is hugely helpful for knowing how to take care of yourself and others. The "latch-bridge" mechanism is especially important, where smooth muscle cells have a discrete state where they attach their myosin heads to actin in a way that "locks" or "latches" the tension without requiring ongoing energy. Latches take between

seconds to minutes to form & dissolve — a simple way to experience the latch-bridge cycle releasing is to have a hot bath and notice waves of muscle relaxation. Latches can persist for minutes, hours, days, months, or years (depending on what prediction they're stabilizing), and the sum total of all latches likely accounts for the majority of bodily suffering. If you are "holding tension in your body" you are subject to the mechanics of the latch-bridge mechanism. Migraines and cluster headaches are almost certainly inappropriate VSMC latches; all hollow organs are surrounded by smooth muscle and can latch. A long-term diet of poor food (e.g. seed oils) leads to random latch formation and "lumpy" phenomenology. Sauna + cold plunges are an effective way to force the clench-release cycle and release latches; likewise, simply taking time to feel your body and put your attention into latched tissues can release them. Psychedelics can force open latches. Many issues in neuropathy & psychiatry are likely due to what I call "latch spirals" — a latch forms, which reduces blood flow to that area, which reduces energy available to those tissues, which prevents the latch from releasing (since releasing the latch requires activation energy and returning to a freely cycling state also increases the cell's rate of energy expenditure).

Core resources:

25. Levin, M. (2022). Technological Approach to Mind Everywhere: An Experimentally-Grounded Framework for Understanding Diverse Bodies and Minds. Frontiers in Systems Neuroscience, 16. https://doi.org/10.3389/fnsys.2022.768201
26. Watson, R., McGilchrist, I., & Levin, M. (2023). Conversation between Richard Watson, Iain McGilchrist, and Michael Levin #2. YouTube.
27. Wikipedia contributors. (2023, April 26). Smooth muscle. In *Wikipedia*, *The Free Encyclopedia*. Retrieved 22:39, July 7, 2023, from https://en.wikipedia.org/w/index.php?

title=Smooth_muscle&oldid=1151758279

28. Wikipedia contributors. (2023, June 27). Circulatory system. In *Wikipedia*, *The Free Encyclopedia*. Retrieved 22:41, July 7, 2023, from https://en.wikipedia.org/w/index.php?title=Circulatory_system&oldid=1162138829
29. Johnson, M., GPT4. (2023). [Mike+GPT4: Latch bridge mechanism discussion].
30. Juliani, A., Safron, A., & Kanai, R. (2023, May 18). Deep CANALs: A Deep Learning Approach to Refining the Canalization Theory of Psychopathology. https://doi.org/10.31234/osf.io/uxmz6
31. Moore CI, Cao R. The hemo-neural hypothesis: on the role of blood flow in information processing. J Neurophysiol. 2008 May;99(5):2035-47. doi: 10.1152/jn.01366.2006. Epub 2007 Oct 3. PMID: 17913979; PMCID: PMC3655718 **Added 11-17-23; recommended priority reading**
32. Jacob M, Ford J and Deacon T (2023) Cognition is entangled with metabolism: relevance for resting-state EEG-fMRI. *Front. Hum. Neurosci.* 17:976036. doi: 10.3389/fnhum.2023.976036 **Added 1-19-24**

To summarize the story so far: tanha is a grabby reflex which is the source of most moment-by-moment suffering. The 'tanha as unskillful active inference' (TUAI) hypothesis suggests that we can think of this "grabbing" as part of the brain's normal predictive and compressive sensemaking, but by default it makes many unskillful predictions that can't possibly come true and must hold in a costly way. The vascular clamp hypothesis (VCH) is that we store these predictions (both skillful and unskillful) in vascular tension. The VCH can be divided into three distinct hypotheses (CVH, VCH, LHH) that describe the role of this reflex at different computational and temporal scales. An important and non-obvious aspect of smooth muscle (e.g. VSMCs) is they have a discrete "latch" setting wherein energy usage and flexibility drops significantly, and sometimes these

latches are overly 'sticky'; unlatching our sticky latches is a core part of the human condition.

Concluding Part I: the above work describes a bridge between three distinct levels of abstraction: a central element in Buddhist phenomenology, the core accounting system within active inference, and a specific muscular reflex. I think this may offer a functional route to synthesize the FEP-AI paradigm and Michael Levin's distributed stress minimization work, and in future posts I plan to explore why this mechanism has been overlooked, and how its dynamics are intimately connected with human problems and capacities.

I view this research program as integral to both human flourishing and AI alignment.

Acknowledgements: This work owes a great deal to Romeo Stevens' scholarship on tanha, pioneering tanha as a 'clench' dynamic, intuitions about muscle tension and prediction, and notion that we commit to dukkha ourselves until we get what we want; Nick Cammarata's fresh perspectives on Buddhism and his tireless and generative inquiry around the phenomenology & timescale of tanha; Justin Mares' gentle and persistent encouragement; Andrea Bortolameazzi's many thoughtful comments and observations about the path, critical feedback, and thoughtful support; and Adam Safron's steadfast belief and support, theorizing on SOHMs, and teachings about predictive coding and active inference. Much of my knowledge of Buddhist psychology comes from the work and teachings of Anthony Markwell; much of my intuition around tantra and interpersonal embodiment dynamics comes from Elena Selezneva. I'm also grateful for conversations with Benjamin Anderson about emergence, to Curran Janssens for supporting my research, and to

Ivanna Evtukhova for starting me on the contemplative path. An evergreen thank you to my parents their unconditional support. Finally, a big thank-you to Janine Leger and Vitalik Buterin's Zuzalu co-living community for creating a space to work on this writeup and make it real.

Footnotes:

[1] We might attempt to decompose the Active Inference – FEP term of 'precision weighting' as (1) the amount of sensory clarity (the amount of precision available in stimuli), and (2) the amount of 'grabbiness' of the compression system (the amount of precision we empirically try to extract). Perhaps we could begin to put numbers on tanha by calculating the KL divergence between these distributions.

[2] We can speculate that the arrow of compression points away from Buddhism's three attributes: e.g. the brain tries to push and prod its SOHMs toward patterns that are stable (dissonance minimization), satisfactory (harmony maximization), and controllable (compression maximization) — similar yet subtly distinct targets. Thanks to both Romeo and Andrea for discussion about the three attributes and their opposite.

[3] (Added July 19, 2023) Skeletal muscle, smooth muscle, and fascia (which contains myofibroblasts with actin fibers similar to those in muscles) are all found throughout the body and reflexively distribute physical load; it's likely they do the same for cognitive-emotional load. Why focus on VSMCs in particular? Three reasons: (1) they have the best physical access to neurons, (2) they regulate bloodflow, and (3) they have the latch-bridge mechanism. I.e. skeletal, non-VSMC smooth muscle, and fascia all likely contribute significantly to distributed stress minimization, and

perhaps do so via similar principles/heuristics, but VSMCs seem to be the only muscle with means, motive, and opportunity to finely puppet the neural system, and I believe are indispensably integrated with its moment-by-moment operation in more ways than are other contractive cells. (Thanks to @askyatharth for bringing up fascia.)

Posted in Uncategorized

&

‹ NEW WHITEPAPER: QUALIA FORMALISM AND A SYMMETRY THEORY OF VALENCE

MINDS AS HYPERSPHERES; THE EQUAL-EXTENSION THESIS AND ITS IMPLICATIONS FOR THE FRAMERATE OF CONSCIOUSNESS ›

Proudly powered by WordPress | Theme: Tabut by Naveen Kharwar.