The Neuroscience of Meditation: Four Models

Michael Edward Johnson
Qualia Research Institute
San Francisco, CA
MIKE@QUALIARESEARCHINSTITUTE.ORG

Background: I'm a philosopher at the Qualia Research Institute (QRI) working on the intersection of neuroscience and phenomenology. As part of this research and to develop my practice, I recently did a 7-day vipassana meditation retreat. The following are some perspectives, models, and hypotheses I had on how some of the 'Western' ideas we're working with at QRI could connect to 'Eastern' contemplative practices. (Yes, I know I wasn't supposed to think during a retreat, but enlightenment will just have to wait, I have things to do. . .)

Buddhism is the start of something really important

I think Buddhism does a *really good job* at telling a coherent and useful story about how the mind works—and the difficulty of telling stories about the mind that are both *coherent* and *useful* is, I think, drastically under-appreciated. Major credit to Buddha here. But Buddhism is also incomplete: it talks about *what* the big-picture dynamics of subjective experience are, but is silent about *how* these dynamics are implemented. Put simply, Buddhism says a lot about the mind, but nothing about the brain, and this is ultimately limiting. We need ways to connect what's going on in the mind during meditation to neuroscience and information theory; we need more frames for what's going on, we need better and more quantitative frames, we need meta frames for how everything fits together.

One thing I *really* like about Buddhism is how it talks about suffering, and how most of the weight of suffering isn't on the immediate pain / negative valence, but the secondary effects as these negative events get fixated on by the mind. We *do* mostly *dukkha* [8] ourselves. I also agree with the Buddhists that this sort of amplification and propagation of suffering revolves around the self, the way we *identify* with our experiences, stories, craving, pain, *etc.* One of Buddhism's core insights is that internal conditions generate the self, and that there are standard methods (*i.e.*, meditative practices) to predictably change these internal conditions, which will then change what kind of (and *how much*) 'self' is generated. This seems important to study, both because it's interesting and because it seems like one of the most highly-leveraged approaches to reducing suffering.

There are plenty of efforts to unify science and meditation, but they tend to use dated neuroscience

The core tool of Buddhism is meditation. Empirically, it seems to work for many people. But *how* does it work? There are a lot of good 'generalist' books in this space– Robert Wright's *Why Buddhism is True*, Culadasa's *The Mind Illuminated* [11], Richard Davidson's *Altered Traits*. My favorite frame for unifying Eastern and Western thought is from Shinzen Young's *The Science of Enlightenment*—Shinzen's core aim is to:

> [R]eformulate the path to enlightenment in a modern, secular, and science-based vocabulary. I wanted to create a system completely free from the cultural trappings and doctrinal preconceptions of traditional Buddhism yet capable of bringing people to classical enlightenment. [...] Most likely, there are things that are true and important about enlightenment that neither the Buddha nor any of the great masters of the past knew, because to know them requires an understanding of modern science.

Shinzen's intuition seems to be that suffering can be modeled as 'turbulence' in brain activity, and enlightenment as the cessation of this turbulence, and that somehow this cessation could be helped along with technology.

It's an intriguing frame, though to my knowledge Shinzen is still looking for an empirical paradigm for this vision. Most others seem to be building on simpler models, where lack of enlightenment is 'too much activity in some specific brain region' or 'not enough activity in some other brain region'. This style of explanation is based on the 'functional localization' paradigm of neuroscience, where we can cleanly associate specific functions to specific regions of the brain—*e.g.*, "this region handles emotional processing, *that* region handles executive function", *etc.* Unfortunately, this assumption increasingly looks like an illusion [1], an artifact of the way we study the brain with fMRI, rather than a core feature of how the brain itself works—and so any interventions based on this simple paradigm will probably fail.

Parallel description is a good meta-frame

I think a really powerful way to keep track of all of this is parallel description. In other words, we can attempt to describe what's going on during suffering & during meditation at multiple levels of abstraction [21], and the more stories we can identify and weave together and cross-validate, the better our understanding will get. In particular, if we get 'stuck' on describing what's happening on one level, we can hop to another level and try to see what's going on from there. I also believe we should be *neuroscience snobs* and only deal with neuroscience models that the *very best* neuroscientists are currently excited about, since the difference between an 'industry standard' neuroscience paradigm and a 'best in the world' paradigm is really enormous. Mostly people talk about better neuroscience being more elegant and predictive, but I think it's even more important to not import old confusions, outdated and wrong ways of looking at the brain. As the saying goes, 'It ain't what you don't know that gets you into trouble. It's what you know for sure that just ain't so.'

Today I have **four stories** (including Buddhism's) about the self, suffering, meditation, and enlightenment. I'm going to set aside questions of metaphysics for the time being, and just focus on identifying or speculating on common themes in mechanisms and dynamics.

Buddhism

Buddhism claims the self arises through the presence of craving and *identifying* with this craving, and this delusion, or 'defilement,' propagates through and infects our entire experience. Meditation helps because by 'noting and knowing' experiences which arise, we can notice their impermanence, and notice that what we call the 'self' is an illusion and our sensations don't really have an 'owner.' Over time as we keep doing this, we slowly generate the inference space to build better intuitive perspectives on the *real* dynamics of our minds, and we feel less of a compulsion to reflexively cling to our objects of craving or aversion (or the craving/aversion itself). This 'spaciousness', or freedom from the usual web of intentionality, allows us to develop the seven enlightenment factors (mindfulness, wisdom, energy, rapture, relaxation, concentration, equanimity), and ultimately the conditions which sustain the self/craving/suffering can drop away.[1]

1. Predictive coding

Predictive coding is a formal framework which says that the brain's core drive is to *minimize surprise*, and that it does this by constantly creating, testing, and adjusting *stories* about the world. This is a two-tier system: the first 'tier' is subconscious prediction, which tries to filter out the 'easy stuff' using simple algorithms. You don't feel the weight of your shirt against your skin, or the pressure of your shoes, or hear the traffic in the background, because this stuff is easy for your brain to categorize then ignore. But if your brain's subconscious can't predict something—say you get bitten by a snake, or you fail a midterm you thought you passed—then this stuff gets sent up for the *conscious mind* to deal with. Essentially, the mind is a story-telling machine, and we make our stories out of the 'unusual' signals the unconscious brain can't explain away. The more surprising/salient something is, the more the brain thinks it's probably important and should be a part of whatever story we're telling ourselves, and the more 'sticky' it feels.

Evolution has not only primed us to *tell* such stories to help us predict our world, but has biased them toward certain topics: our stories are often about how we *want* things and how we'll be *sad if we don't get them*. Evolution didn't have any incentive to make us tell ourselves satisfying stories where we can just relax, where we're doing enough, where we *have enough*, where suffering is not always close; quite the opposite. The unsatisfied human is an active human, less likely to get eaten by tigers and more likely to outwork lazy competitors. This tends to make our stories *stressful*, full of *grasping*, *restlessness*, and *suffering*. Also unfortunately, evolution doesn't really care about higher-order *logical coherence*, and allows us to tell ourselves multiple conflicting stories, where we want several contradictory things. At any rate, under this model the self is the 'grand meta-narrative',

1. Links: The original Pāli texts (*suttas*) of Buddhism [5]; Orientation on the Contemplative Path.[26]

the most *interconnected/antifragile/causally-upstream* story. The story which contextualizes all other stories.

Under the predictive coding model, I'd describe the process of meditation as attempting to 'tag' sensations early in the prediction pipeline as "okay/nothing to worry about/not anomalous/not something to update on/doesn't have to be part of our story", before the sensation becomes high-confidence and sticky and *needs* to be part of the story. In the short term, this helps avoid filling up the mind with fixed points (the more fixed points we have to predict with our stories, the larger and more stressful the stories become). Iterated enough in the long-term, this can help tune down the *overall mechanism* which drives this storytelling.[2] Our stories become smaller, less sticky, less unpleasant, until we have no stories at all, just our raw sense-data. (Relatedly, if you find yourself suffering, try to make your story smaller. It may help.)[3]

2. Connectome-specific harmonic waves

All systems with periodic activity have natural modes, frequencies they 'like' to resonate at. Wineglasses, tuning forks, and guitars have them; the brain has them too. Connectome-specific harmonic waves (CSHW) is a new but promising paradigm for *defining* and *measuring* these natural harmonic modes in brains.

But resonance in the brain isn't just passive, like a guitar; rather, which frequencies naturally resonate changes over time, and the brain is more like a *self-tuning musical instrument*. Or to put it in a fancier way, the brain is essentially a *conditional, differential amplifier of its natural harmonics*. A guitar that can change the shape of its box, or a piano that can change the length of its strings. Unfortunately, evolution has primed the brain to stabilize, amplify, and propagate dissonance [25] (unpleasant feeling) within itself, because dissonance information often implies a *problem* somewhere, and paying attention to potential problems has significant survival value.

2. Vipassana training consists of both working on *mindfulness*, careful observation of what enters the mind, and *concentration*, holding attention on some object. These capacities are seen as somehow opposed but also complimentary to each other, and it's important that they develop in tandem, that one doesn't lag behind. This roughly corresponds with predictive coding's idea that perception consists of a negotiation [14] between 'bottom-up' raw sense data, and 'top-down' models of reality which provide context for the raw data and fill in any missing gaps. Likewise, predictive coding suggests that if one of these processes is much stronger than the other, problems occur—*e.g.*, if bottom-up sense-data dominates we'll experience noise and confusion as we struggle to sort things out, and if top-down predictions dominate we'll experience hallucinations, stories not connected to facts. My intuition is that modern life's focus on planning and abstract thought tends to make us a little 'top-heavy' here, so at least as beginners, *most* people would probably benefit more from mindfulness meditation (retraining themselves to focus on their raw sense-data) than they would from concentration meditation (being able to maintain and conceptually purify their stories/predictions about their world). This may be different than in Buddha's time.
3. Links: The Genius Neuroscientist Who Might Hold the Key to True AI [24]; The free-energy principle: a unified brain theory? [18]

I believe that this *amplification* and *propagation* of dissonance (suffering) likely happens through *brain harmonics getting 'coupled' together*, where adding energy to one frequency band spills over into others which are mathematically similar. This 'harmonic coupling' [12] happens in musical instruments too—*e.g.*, if you hit Middle C (C4) on a piano then bring your ear close to the strings for nearby notes like E4, G4, or C5, you'll hear them vibrate too. This idea of frequencies being 'coupled together' could be a core part of the mechanism by which the brain propagates information inside itself, and in fact we might attempt to *define* the self as '*the particular configuration of how semi-stable frequency bands are coupled together in the brain.*'

How does meditation affect harmonic coupling? In the short-term, the 'noting and knowing' of meditation may act to dampen specific harmonics before their activity spills over into others and becomes self-propagating, leading to a quieter mind with a better signal-to-noise ratio. Iterated over the long term, I'd speculate that meditation helps by *permanently de-tuning some of the tight coupling* that so effectively propagates reactive-style behavior/updates through the system. The system becomes less 'locked in' to certain responses, and in general has *much less internal noise*[4] and much more perceptual clarity due to dampening certain forms of feedback.

One of the core strengths of this model is that *all of this stuff about natural brain resonances and harmonic coupling is* ***empirically measurable*** with the CSHW paradigm, and we could potentially use it to track meditation progress, and later on perhaps as the basis for technological interventions.[5],[6] I find it plausible that this model could also be used to try to ground and quantify Shinzen's metaphor about suffering arising from turbulence in the brain: a core piece of the CSHW paradigm is the 'wave propagation equation' [27], the specific parameters which determine how waves of neural activity travel in the brain (or as Shinzen puts it, we can measure the brain's 'viscosity'). Changes in this property of the brain will lead to changes (and perhaps reductions) in whatever turbulence patterns are arising. And again, this is measurable with current technology.[7]

4. We might also expect both greater perceptual clarity and less background dissonance if (as Daniel Ingram suggests) meditation can increase the 'sample rate' of experience, because there would be less aliasing of sense-data due to the Nyquist limit [9]. In this case we could look for frequency bands (*e.g.*, alpha vs beta) which are close to a 1:2 ratio in non-meditators, and under a 1:2 ratio in experienced meditators.
5. Very speculatively, perhaps we could use alexithymia [28] ('missing emotions') as a target model for getting rid of the various 'defilements' Buddhism talks about. *E.g.*, if the brain can't resonate to the harmonics which characterize anger, it can't be angry. Or we could go the other way and tune the brain to resonate more easily to the characteristic harmonics of loving kindness.
6. Since the brain has to compute and keep track of many, separate things at once, I suspect that it will have multiple 'functional partitions' of harmonics, where a harmonic will be strongly coupled to other harmonics inside the partition but only weakly to harmonics outside it. I think a good place to begin a quantitative analysis would be to use what we know about the mathematics of harmonic coupling [23] in order to estimate the *expected functional isolation* between partitions. My expectation is that more reactive people have more 'crosstalk' between their partitions.
7. Links: A Future for Neuroscience [20]; Selen Atasoy describing CSHW [3]; Quantifying Bliss [17].

3. Neural annealing

Annealing involves heating a metal above its recrystallization temperature, keeping it there for long enough for the microstructure of the metal to reach equilibrium, then slowly cooling it down, letting new patterns crystallize. This releases the internal stresses of the material, and is often used to restore ductility (plasticity and toughness) on metals that have been 'cold-worked' and have become very hard and brittle—in a sense, annealing is a 'reset switch' which allows metals to go back to a more pristine, natural state after being bent or stressed. I suspect this is a useful metaphor for brains, in that they can become hard and brittle over time with a build-up of internal stresses, and these stresses can be released by periodically entering high-energy states where a more natural neural microstructure can reemerge.[8]

Furthermore, from what I gather from experienced meditators, *successfully entering meditative flow may be one of the most reliable ways to reach these high-energy brain states. I.e.*, it's very common for meditation to produce feelings of high intensity, at least in people able to actually enter meditative flow. Meditation also produces more 'pure' or 'neutral' high-energy states, ones that are free of the intentional content usually associated with intense experiences which may distort or limit the scope of the annealing process. So we can think of intermediate-to-advanced ('successful flow-state') meditation as a reheating process, whereby the brain enters a more plastic and neutral state, releases pent-up structural stresses, and recrystallizes into a more balanced, neutral configuration as it cools. Iterated many times, this will drive an evolutionary process and will produce a very different brain, one which is more unified and anti-fragile, less distorted toward intentionality [15], and in general *structurally optimized against stress.*

An open question is *how* or *why* meditation produces high-energy brain states. There isn't any consensus on this, but with a nod to the predictive coding framework, I'd offer that bottom-up sense-data is generally excitatory, adding energy to the system, whereas top-down predictive Bayesian models are generally inhibitory, functioning as 'energy sinks'. And so by 'noting and knowing' our sensations before our top-down models activate, in a sense we're diverting the 'energy' of our sensations away from its usual counterbalancing force. If we do this long enough and skillfully enough, this energy can build up and lead to 'entropic disintegration' [13], essentially pushing enough energy into the system that existing attractors are disrupted and annealing can occur.[9]

If this model is true, it feels very important for optimizing a meditation practice. *E.g.*, we should try to figure out some rules of thumb for:

- How to identify a high-energy brain state, in yourself and others, and how best to create them;

8. Robin Carhart-Harris [13] frames this as 'entropic disintegration', or the idea that if the brain enters a high-energy state, existing attractors will break down. He doesn't speak about what comes *after* the disintegration– the self-organization toward new attractors and the crystallization of these attractors as the brain 'cools' – aka *annealing*— but I believe it's implicit in his framework.
9. Thanks to Andrés for discussion here.

- Things to do, and things not to do, during an annealing process ('how to anneal the right things');
- Identifying tradeoffs in 'cooling' the brain quickly vs. slowly.

Off the top of my head, I'd suggest that one of the worst things you could do after entering a high-energy brain state would be to fill your environment with distractions (*e.g.*, watching TV, inane smalltalk, or other 'low-quality patterns'). Likewise, it seems crucial to avoid socially toxic or otherwise highly stressful conditions. Most likely, going to sleep as soon as possible [22] without breaking flow would be a good strategy to get the most out of a high-energy state–the more slowly you can 'cool off' the better, and there's some evidence annealing can continue during sleep. Avoiding strong negative emotions during such states seems important, as does managing your associations (psychedelics [16] are another way to reach these high-energy states, and people have noticed there's an 'imprinting' process where the things you think about and feel while high can leave durable imprints on how you feel after the trip). It seems plausible that taking certain nootropics could help strengthen (or weaken [6]) the magnitude of this annealing process.

Finally, to speculate a little about one of the deep mysteries of life, perhaps we can describe *love* as the result of a strong annealing process while under the influence of some pattern. *I.e.*, evolution has primed us such that certain intentional objects (*e.g.*, romantic partners) can trigger high-energy states where the brain smooths out its discontinuities/dissonances, such that *given the presence of that pattern* our brains are in harmony.[10] This is obviously a two-edged sword: on one hand it heals and renews our 'cold-worked' brain circuits and unifies our minds, but also makes us *dependent*: the felt-sense of this intentional object becomes the key which unlocks this state. (I believe we can also anneal to *archetypes* instead of specific people.)

Annealing can produce durable patterns, but isn't permanent; over time, discontinuities creep back in as the system gets 'cold-worked'. To stay in love over the long-term, a couple will need to re-anneal in the felt-presence of each other on a regular basis.[11] From my experience, some people have a natural psychological drive toward reflexive stability here: they see their partner as the source of goodness in their lives, so naturally they work hard to keep their mind aligned on valuing them. (It's circular, but it works.) Whereas others are more self-reliant, exploratory, and restless, less prone toward these self-stable loops or annealing around external intentional objects in general. Whether or not, and within which precise contexts, someone's annealing habits fall into this 'reflexive stability attractor' might explain much about, *e.g.*, attachment style, hedonic strategy, and aesthetic trajectory.[12]

10. Anecdotally, the phenomenology of love-annealing is the object 'feels beautiful from all angles'. This may imply that things (ideas, patterns, people) which are more internally coherent & invariant across contexts can produce stronger annealing effects—*i.e.*, these things are more easy to fall deeply in love with given the same 'annealing budget', and this love is more durable.
11. It's important to note that both intense positive and intense negative experiences can push the brain into high-energy states; repeated annealing to negative emotions may serve many of the same functions as 'positive annealing', but also predispose brains to 'sing in a minor key' (see 'kindling' [4]).
12. Links: Annealing (metallurgy) [7]; The entropic brain [13]

4. On unification

Ultimately, the goal is to not only *collect* parallel descriptions, but also try to *unify* them in a way that paints a richer and more constraining picture of what's going on. A Future for Neuroscience [20] suggests we might be able to partially unify connectome-specific harmonic waves (CSHW) and predictive coding by treating low-frequency harmonics as the *mechanism* by which information about Bayesian priors is propagated:

> High frequency harmonics will tend to stop at the boundaries of brain regions, and thus will be used more for fine-grained and very local information processing; low frequency harmonics will tend to travel longer distances, much as low frequency sounds travel better through walls. This paints a possible, and I think useful, picture of what emotions fundamentally are: semi-discrete conditional bundles of low(ish) frequency brain harmonics that essentially act as Bayesian [*i.e.*, predictive coding] priors for our limbic system. Change the harmonics, change the priors and thus the behavior.

Friston's work [24] on predictive coding may also offer a generative frame for extending the annealing metaphor. Neural dynamics happen across multiple functional scales, and a neuron's behavior should contribute to healthy brain function at each scale–*e.g.*, the micro, meso, and macro. Neurons should contribute to functional circuits, circuits should combine into well-calibrated computational systems, and computational systems should collaborate to drive healthy emotion and behavior. But these dynamics are emergent, not coordinated top-down, and sometimes the selection pressures which shape each scale's behavior will become misaligned with the needs of other scales (due to *e.g.*, trauma, uncompensated system drift, or simple design flaws of the brain-as-self-organized-system), leading to a systemic build-up of error (*i.e.*, free energy [18]). Neural annealing is (almost by definition) the optimal self-organized mechanism for releasing this 'structural' source of free energy: as prediction errors (free energy) builds up, the brain enters a high-energy state, causing entropic disintegration [13] (weakening previously 'sticky' attractors), and finally the brain finds new equilibria based on current micro, meso, and macro constraints, which should generate lower background levels of error.

This opens the door for a neat definition of the self as *the process that attempts to manage the implicit (expected) free energy from the mismatch of micro vs meso vs macro drives.* This would describe the 'self' as a sort of 'head of parliament', constantly attempting to make deals to bring various hierarchical levels of the mind into alignment as they naturally drift apart. But if meditation does kickstart a neural annealing process, this would help align these various functional scales, thus reducing both the expected amount of cross-scale free energy build-up and the very need for a 'self'.

Finally, we may be able to *usefully describe the Buddhist jhanas* [2] through a combination of **CSHW** and **neural annealing**. Essentially, Buddha noted that as one follows the meditative path and eliminates unwholesome mental states, they will experience various trance-like states of deep mental unification he called 'jhanas'. These are seen as progressive steps to full enlightenment–the first few jhanas focus on joy, and after these are

attained one can move to jhanas which revolve around contentment and feelings of infinity, and finally full cessation of suffering. Importantly, these experiences are not seen as 'mere signposts' on the path, but active processes which are *causally involved* in the purification of the mind—in the original Pāli, the root of 'jhana' can refer to both 'meditate' and 'burn up', *e.g.*, to destroy the mental defilements holding one back from serenity and insight.

A 'modern' approach here might be to identify the various jhanas as natural resonant modes of the brain—*i.e.*, different jhanas would map to different harmonic configurations, each with a different phenomenological texture, but all high in consonance/harmony. If this is true, we should be able to do neat things like identify which jhana a meditator is experiencing from their CSHW data, or reconstruct Buddhism's list of jhanas from first principles based on the math of which brain harmonics can be combined in a way that produces high consonance/harmony. Perhaps we could even find a novel, unexplored jhana or two, pleasant configurations of brain harmonics that even most experienced meditators have never experienced.

But if we add neural annealing to this model, we can start to understand how experiencing the various jhanas may actively sculpt the mind. At its most basic, meditation offers a context for people to sit with their thoughts and maybe do some work on themselves, and get some practice 'getting out of their own way'. Basically removing the 'defilements' which clutter up their brain harmonics, much like removing a clamp from a bell or shaking a mouse out of a trombone. Once these 'resonance impediments' are removed, and energy is added to the system (through effortful meditation), brains will naturally start to self-organize toward the simpler resonant configurations, the first jhanas. But importantly, *highly-resonant states are also high-energy states–i.e.*, the very definition of resonance is that energy travels in a periodic pattern that reinforces itself, instead of dissipating in destructive interference. So if you get a brain into a highly-resonant state (a jhana) and keep it there, *this will also start a neural annealing process*, basically purifying itself (and making it easier and easier to enter into that particular resonant state–"harmonic recanalization") more or less automatically.

With this in mind, we might separate Buddha's path to enlightenment into two stages: first, one attempts to remove the psychological conditions which prevent them from attaining a minimum level of 'whole-brain resonance'; mostly, this will involve trying to meditate, experiencing a problem in doing so, fixing the problem, trying to meditate again. Rinse, repeat–possibly for years. But in the second stage, once enough of these conditions are gone and resonant momentum can accumulate, they can start 'climbing the jhanas,' mostly just entering meditative flow and letting the math of Laplacian eigenmodes and neural annealing slowly shape their mind into something that resonates in purer and purer ways [13] until at

13. What *precisely* is happening as one climbs the various jhanas? Resonance in chaotic systems is inherently unstable, and so if the first jhana is "a minimum level of whole-brain resonance" we should expect many perturbations and failures in maintaining this pleasant state as unpredicted sense-data, chaotic internal feedback loops, and evolved defenses against 'psychological wireheading' [10] knock the system around. Each additional jhana, then, may be thought of as a widening of the set of factors being stabilized, or using a higher-dimensional or further-optimized implicit model of wave dynamics to compensate for more sources of turbulence [19]. This optimization process might separate into discrete steps or strategies

the end it becomes something which can only support harmony, something which simply has no resources that can be used to sustain dissonance.

Putting it all together:

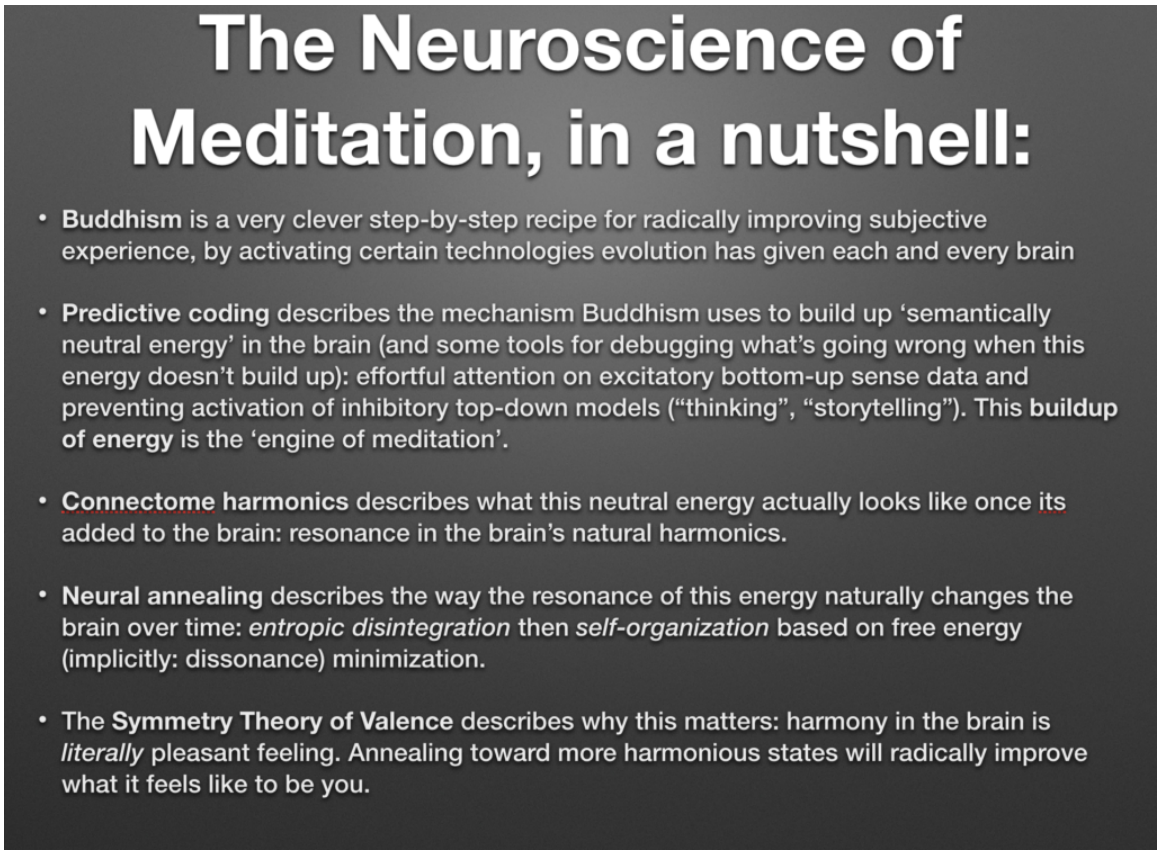

Paths forward

It could take years to fully explore even one of these models, let alone all combinations of them. But I suspect there could be low-hanging fruit, useful insights or inspiration that could be found in an afternoon of careful thought. We might ask ourselves what the surprising implications of each model are, whether these models resonate with experienced meditators, and *what technologies each model seems to point to in terms of interesting interventions*. We can also keep our minds open to further alternative scientific models—the more entry points we have to this topic the better holistic story we can tell. It's possible that the above models are not super useful and there are others which are much better. But it's my strong intuition that this *type* of approach—pulling the best quantitative models

(jhanas), each with their own particular phenomenology, depending on what kind of turbulence it's best at stabilizing. I expect we'll find that earlier jhanas are characterized by seeking particular narrow resonant configurations that work; later jhanas flip the script and are characterized by seeking out the remaining distortions in the 'web of phenomenology', the problem states that *don't* resonate, in order to investigate and release them.

from neuroscience (and other branches of science) and using them to try to tell the story of 'what meditation is doing' from multiple angles—will be a very generative path forward.

Acknowledgments

Thank you to Andrés Gómez Emilsson, Romeo Stevens, Adam Safron, Anthony Markwell, Danny Donabedian, and Adrian David Nelson for helpful and stimulating discussion on these topics. This work is dedicated to Ivanna Evtukhova, whose warmth, energy, and patience created the conditions for this to happen.

References

[1] Evan Thompson, . URL https://www.youtube.com/watch?time_continue=2448&v=rG3daqNH700.

[2] The Jhanas in Theravada Buddhist Meditation, . URL https://accesstoinsight.org/lib/authors/gunaratana/wheel351.html.

[3] Selen Atasoy: Enhanced Improvisation in LSD Brain Processing, . URL https://www.youtube.com/watch?v=jKAZCzATmnk.

[4] Should the 'kindling' concept direct mental-health treatment? – Elizabeth Svoboda | Aeon Essays, . URL https://aeon.co/essays/should-the-kindling-concept-direct-mental-health-treatment.

[5] The Sutta Pitaka of the Pali Canon, . URL https://www.vipassana.com/canon/sutta.php.

[6] Agmatine, June 2020. URL https://en.wikipedia.org/w/index.php?title=Agmatine&oldid=961435315.

[7] Annealing (metallurgy), May 2020. URL https://en.wikipedia.org/w/index.php?title=Annealing_(metallurgy)&oldid=959480419.

[8] *Dukha*, June 2020. URL https://en.wikipedia.org/w/index.php?title=Du%E1%B8%A5kha&oldid=962192862.

[9] Nyquist frequency, June 2020. URL https://en.wikipedia.org/w/index.php?title=Nyquist_frequency&oldid=960974661.

[10] Wirehead (science fiction), May 2020. URL https://en.wikipedia.org/w/index.php?title=Wirehead_(science_fiction)&oldid=955406003.

[11] S. Alexander. Book Review: The Mind Illuminated, Nov. 2018. URL https://slatestarcodex.com/2018/11/28/book-review-the-mind-illuminated/.

[12] R. T. Canolty and R. T. Knight. The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, 14(11):506–515, Nov. 2010. ISSN 1364-6613. doi: 10.1016/j.tics.2010.09.001. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3359652/.

[13] R. L. Carhart-Harris, R. Leech, P. J. Hellyer, M. Shanahan, A. Feilding, E. Tagliazucchi, D. R. Chialvo, and D. Nutt. The entropic brain: a theory of conscious states informed by neuroimaging research with psychedelic drugs. *Frontiers in Human Neuroscience*, 8, Feb. 2014. ISSN 1662-5161. doi: 10.3389/fnhum.2014.00020. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3909994/.

[14] P. R. Corlett, C. D. Frith, and P. C. Fletcher. From drugs to deprivation: a Bayesian framework for understanding models of psychosis. *Psychopharmacology*, 206(4):515–530, Nov. 2009. ISSN 0033-3158. doi: 10.1007/s00213-009-1561-0. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2755113/.

[15] A. G. Emilsson. The Tyranny of the Intentional Object, Nov. 2016. URL https://qualiacomputing.com/2016/11/19/the-tyranny-of-the-intentional-object/.

[16] A. G. Emilsson. Connectome-Specific Harmonic Waves on LSD, June 2017. URL https://qualiacomputing.com/2017/06/18/connectome-specific-harmonic-waves-on-lsd/.

[17] A. G. Emilsson. Quantifying Bliss: Talk Summary, June 2017. URL https://qualiacomputing.com/2017/06/18/quantifying-bliss-talk-summary/.

[18] K. Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, Feb. 2010. ISSN 1471-0048. doi: 10.1038/nrn2787. URL https://www.nature.com/articles/nrn2787.

[19] K. Hartnett. Mathematicians Tame Rogue Waves, Lighting Up Future of LEDs. URL https://www.quantamagazine.org/mathematicians-tame-rogue-waves-lighting-up-future-of-leds-20170822/.

[20] M. Johnson. A Future for Neuroscience – Opentheory.net. URL https://opentheory.net/2018/08/a-future-for-neuroscience/.

[21] R. McClamrock. McClamrock: Marr's Three Levels. URL https://www.albany.edu/~ron/papers/marrlevl.html.

[22] J. D. Payne, M. A. Tucker, J. M. Ellenbogen, E. J. Wamsley, M. P. Walker, D. L. Schacter, and R. A. Stickgold. Memory for Semantically Related and Unrelated Declarative Information: The Benefit of Sleep, the Cost of Wake. *PloS one*, 2012. doi: 10.1371/journal.pone.0033079.

[23] B. Pletzer, H. Kerschbaum, and W. Klimesch. When frequencies never synchronize: the golden mean and the resting EEG. *Brain Research*, 1335:91–102, June 2010. ISSN 1872-6240. doi: 10.1016/j.brainres.2010.03.074.

[24] S. Raviv. The Genius Neuroscientist Who Might Hold the Key to True AI. *Wired*. ISSN 1059-1028. URL https://www.wired.com/story/karl-friston-free-energy-principle-artificial-intelligence/.

[25] W. A. Sethares. Relating Tuning and Timbre. URL https://sethares.engr.wisc.edu/consemi.html.

[26] R. Stevens. Neurotic Gradient Descent: Orientation on the Contemplative Path, Dec. 2018. URL https://neuroticgradientdescent.blogspot.com/2018/12/orientation-on-contemplative-path.html.

[27] H. R. Wilson and J. D. Cowan. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13(2):55–80, Sept. 1973. ISSN 1432-0770. doi: 10.1007/BF00288786. URL https://doi.org/10.1007/BF00288786.

[28] E. Young. This medical condition is widely unheard of but can impact the emotions we feel, June 2018. URL https://www.independent.co.uk/news/long_reads/alexithymia-condition-impacts-emotions-new-research-science-a8374496.html.