

UNITED WISDOM

The Neuroscience of Meditation: What the MRI Shows

World and Consciousness

Andrey Zarinsh & Maria Evilkut

W20 · EN

@andrey.zarinsh @ilkutki

The Neuroscience of Meditation: What the MRI Shows

When philosophy went into the scanner

Andrey Zarinsh & Maria Evilkut

United Wisdom · Science · An original series

> ***Archival editor's note.*** *This full translation preserves the supplied archival edition. Its claims about causal brain changes, diagnostic associations, and equivalence between spiritual traditions and neural findings are inherited authorial interpretations, not newly verified scientific or clinical conclusions. The existing website translation differs materially: it gives Lazar 15 controls and roughly nine years of practice (archive: 20 and ten), Tang 80 participants (archive: 40), and replaces the archive's thirtyfold gamma-amplitude statement with a different account. It also changes the attribution of the Lazar quotation and Hölzel's nationality. The archival wording is retained below to make the edition transparent; neither set of disputed details was independently source-verified in this translation task. Correlation, small samples, study design, and the difference between measured correlates and metaphysical conclusions remain essential limits.*

I. When philosophy went into the scanner

For thousands of years philosophers and mystics said roughly the same thing. There is a place inside a person that does not depend on circumstances. Advaita called it the witness. The Stoics called it prohairesis, moral purpose. Myth drew it as fire, taken by Prometheus. The mandala drew it as the point at the center of the circle, the bindu.

These were words. Very different words, out of different cultures, in different languages. They could be loved, they could be mocked, they could be filed under "Eastern exotica" or "Hellenic idealism."

Since 2004 these words have begun to yield to measurement. Not in the sense that anyone has measured the witness — no one knows how to do that. In the sense that people who work with that witness systematically turned out to have something in their brains that other people do not have. Tibetan monks first. Then ordinary Americans who had taken an eight-week course.

This article is about what exactly is visible. And about what does not follow from it.

@andrey.zarinsh @ilkutki

II. The scanner and its blind spots

First, a word about method.

Functional magnetic resonance imaging (fMRI) measures neither "thoughts" nor "consciousness." It measures the flow of blood into different regions of the brain. Where neurons fire harder, more oxygen is needed, blood arrives faster, and the magnetic signal from that blood is caught by the scanner. This is the BOLD signal — blood-oxygen-level-dependent.

Electroencephalography (EEG) measures the electrical activity of the cortex through electrodes on the scalp. What fMRI lacks in temporal resolution, EEG supplies: milliseconds. Its spatial resolution is worse. Where in the brain an electrical rhythm is coming from has to be reconstructed mathematically.

Structural MRI is the third instrument. It shows no activity, but it shows how much gray matter sits in each region and how densely that region is connected. The picture is static, but it can be taken from the same person twice — before and after, say, eight weeks of practice.

All three have blind spots. They show groups of neurons, not single cells. They show change, not cause. They show a correlate, not the experience itself. This has to be settled at the outset: neuroscience does not "find" consciousness in the brain. It finds that some people's brains work differently from other people's, and it tries to work out which kinds of training lead to which kinds of change.

Meditation, in this respect, is a convenient object. It is easy to standardize. It is easy to dose, in hours of practice. There are people with decades of experience, and there are programs through which a beginner can take a verifiable course. And there are subjective reports that can be set beside objective measurements.

III. The DMN: the network of "I am thinking about myself"

In 2001 the neuroscientist Marcus Raichle, at Washington University in St. Louis, made an unexpected observation. When subjects in the scanner are doing nothing in particular — lying still, waiting, resting — a specific network of regions comes on. Give them a task, counting backward for instance, and that same network goes quiet.

Raichle named it the default mode network, the DMN. Its principal nodes:

— the mPFC, the medial prefrontal cortex, the front of the forehead along the midline; — the PCC, the posterior cingulate cortex, deeper, at the center of the brain, toward the back of the head; — the precuneus, next to the PCC; — the angular gyrus and the temporoparietal junction, at the sides.

What does this network do while a person is "doing nothing"? Years of experiments gave an answer: it generates thoughts about the self. Memories. Plans. Comparisons. Inner dialogue. Imagined scenes of meetings to come. Resentment of a colleague. The question of whether one is good enough. Rehearsals of future conversations. Returns to old mistakes.

Psychiatry picked this up quickly. An overactive DMN correlates with depression — depressive rumination is, in effect, a DMN that has jammed. An overactive DMN correlates with anxiety disorders. With chronic stress. With a low capacity to be present.

And here an interesting question arose. If the DMN is the neural substrate of the narrative about oneself, what happens in people who train themselves to step back from that narrative? In meditators?

IV. Brewer 2011: the DMN is quieter in meditators

In 2011 Judson Brewer, then at Yale, published a study in PNAS that became one of the most cited in the field.

The design is simple. Twelve experienced meditators with more than ten thousand hours of practice, and twelve controls matched for age and sex. Everyone was put in an fMRI scanner and asked to perform three kinds of meditation: concentration on the breath, loving-kindness (mettā), and choiceless awareness. With rest between them.

The result: in the meditators the main nodes of the DMN — the mPFC and the PCC — were relatively deactivated across all three practices. That is the first finding. The second: in those same meditators the PCC was more strongly coupled to the lateral prefrontal and anterior cingulate cortices, the regions that handle cognitive control and monitoring.

Put differently: while a meditator practices, the network of self-narrative goes quieter, and the network that watches what is happening becomes more active and more tightly bound to the first. A loop forms. Mind-wandering is noticed earlier, and the observing system damps it before it can grow.

The meditators themselves, questioned after the session, reported considerably less mind-wandering than the control group. The subjective report matched what was visible in the BOLD signal.

This does not prove that they "saw God" or "touched the absolute." It shows something plainer: training attention changes how the brain's networks operate, and it changes precisely those networks that handle self-preoccupation.

V. Lazar 2005: thicker cortex in meditators

Before Brewer's work there was another study, a structural one.

In 2005 Sara Lazar, at Massachusetts General Hospital, affiliated with Harvard Medical School, took structural MRI scans of twenty experienced practitioners of vipassanā, Buddhist insight meditation, and twenty controls. The practitioners had a mean of about ten years behind them and meditated an average of about forty minutes a day.

The result: in the meditators the cortex was thicker in several regions than in the controls.

— The right insula, the region tied to interoception, the perception of signals from the body; — The anterior cingulate — attention, conflict monitoring; — The prefrontal cortex — higher functions, regulation; — Sensory regions handling hearing and touch.

The gain is slight: four to eight thousandths of an inch, 0.1 to 0.2 mm. But it is statistically significant. And, importantly, the effect was strongest in the most experienced practitioners. Which means this is not a matter of "people already built that way take up meditation" but of "the practice brought them there." A dose-dependent effect, in the language of science.

Lazar wrote a sentence that would be quoted many times: “The structure of the adult brain can change in response to repeated practice.” This seemed to challenge the old neuroscience textbook, which said that the brain no longer grew after twenty-five. Ten years later such claims had become commonplace. In 2005 they still sounded bold.

VI. Hölzel 2011: eight weeks are enough

The objection to Lazar's study is obvious. Perhaps thick cortex in those regions is not the result of meditation but its cause. Perhaps people with a cortex like that simply find meditation easy, and so they are the ones who keep at it. Correlation, not causation.

To settle the question a longitudinal design is needed: take people with no experience, measure their brains, train them, measure again, compare.

The British-German researcher Britta Hölzel did exactly this. In 2011 she published a before-and-after study in Psychiatry Research: Neuroimaging.

Sixteen people with no meditation experience took an eight-week MBSR course — Mindfulness-Based Stress Reduction, the program John Kabat-Zinn built in 1979 at the University of Massachusetts. Forty-five minutes of practice a day, plus weekly group sessions. Structural MRI before the course and after. A control group of seventeen waited out the same period without training.

The results:

— The left hippocampus: gray matter density rose. The hippocampus handles memory and learning, and its shrinkage is associated with depression and chronic stress. — The posterior cingulate, the PCC: density rose. Which is a paradox — this is a DMN node, and while the density of neurons increased, its activity during meditation, as Brewer showed, drops. — The temporoparietal junction, the TPJ: rose. This region is tied to taking another's perspective, to empathy, to theory of mind. — The cerebellum: rose. There is growing evidence that the cerebellum does more than coordinate movement, and takes part in emotional regulation.

A parallel study by Hölzel showed one more thing: the amygdala, the brain's alarm center, fell in density — and that decrease correlated with the reduction in stress the participants reported themselves.

So eight weeks of moderate daily practice are enough to change, measurably, the structure of an ordinary person's brain in the regions that handle memory, emotion, empathy and stress.

This study became one of the most widely reproduced in the popular press. The Harvard Gazette ran it under the headline "Eight weeks to a better brain." The headline simplifies — science does not claim "better," it claims "different." But the direction of the changes does, in most cases, coincide with what the participants themselves call improvement.

VII. Davidson and Lutz: gamma waves in the monks

Alongside all of this ran another line of work: the study of people with extreme amounts of practice behind them.

Richard Davidson, a neuroscientist at Wisconsin-Madison, received an unusual offer from the Dalai Lama in the early 2000s. Tibetan monks were willing to take part in research if it would help science. Davidson assembled a group of eight monks with between ten thousand and fifty thousand hours of practice, most of them masters of the Kagyu and Nyingma traditions. And a control group: ten students who had been taught to meditate for one week.

In 2004 Antoine Lutz and Richard Davidson published a result in PNAS that the academic world received with skepticism.

When the monks entered a state of compassion meditation — non-referential compassion — gamma waves at 25–42 Hz were recorded in their cortex, with an amplitude averaging thirty times the baseline. The control group did not show amplitudes of this kind.

The gamma rhythm is tied to the integration of information across different regions of the brain. When you listen closely to music and recognize the melody, or solve a hard problem and see the answer, gamma rises in those seconds. But normally for brief milliseconds, and locally. In the monks the gamma synchrony was stable, prolonged, and spread over wide areas of cortex.

More than that: even when the monks were not meditating but simply resting, their baseline gamma activity was above normal. Which suggests that decades of practice change not only states but traits — the stable characteristics of a brain.

The study was much argued over. Some said the effect was accidental and the sample tiny. Others said gamma is only an electrical rhythm and cannot be equated with "enlightenment." Both objections are fair. But the fact stands: in people who have trained for decades, the brain works differently at the physical level, and that difference is not a myth but a measured quantity.

VIII. Tang and Posner: five days change attention

From Davidson's work it is easy to draw the wrong conclusion: that decades are required before anything shifts. They are not.

Yi-Yuan Tang, a Chinese neuroscientist, and Michael Posner, one of the leading researchers of attention in the United States, ran a study in 2007 on a radically shorter timescale.

Forty Chinese students were randomly divided into two groups. The first did five days of twenty-minute IBMT, Integrative Body-Mind Training, a practice combining breath, relaxation and light awareness. The second did five days of twenty minutes of relaxation with no meditative component.

After five days:

— The IBMT group performed significantly better on the Attention Network Test, especially on the conflict component, where a visual cue points one way and the correct answer lies the other. The gain held the following day as well. — After a laboratory stress test, cortisol, the principal stress hormone, rose less in the IBMT group. — Self-reported anxiety and negative mood fell.

Five days of twenty minutes. A hundred minutes of practice in total. Not ten years, not ten thousand hours. And already a measurable effect on attention and on stress reactivity.

There have been many replications since. The effect size in short courses is smaller than in long ones, and not every component recovers — structural change in the cortex, for instance, does not have time to appear in five days. But even the minimum dose produces a shift.

IX. What changes in experience

Put all of it together and the picture of subjective effects looks roughly like this:

— less mind-wandering, a greater capacity to bring attention back; — faster recognition that a thought or an emotion has taken hold; — a weaker amygdala response to negative stimuli —

fewer flare-ups on autopilot; — a more accurate sense of the body: interoceptive signals arrive before they have grown into a problem; — the sense of "I" becomes less rigid, more transparent.

The last point is the interesting one. Experienced practitioners describe it the same way: the narrative about oneself becomes one of the objects of attention rather than its frame. Not "I am thinking about work" but "I notice a thought about work." This difference in experience appears to correlate with exactly the weakening of the DMN.

And here, at this point, neuroscience meets what Advaita called the witness. Not the I-narrative, but the agency that sees the I-narrative as a phenomenon.

X. Three mechanisms

Neuroscientists explain all of these effects by three overlapping mechanisms.

Regulation of attention. The anterior cingulate and the lateral prefrontal cortex act as a dispatcher, noticing the departure from the chosen object and returning to it. Meditation trains this circuit specifically — the milliseconds of departure grow shorter.

Regulation of emotion. The prefrontal cortex has inhibitory connections to the amygdala. The stronger those connections, the weaker the autopilot of fear and anger. Meditation — practices of acceptance and compassion above all — strengthens the link, and the amygdala becomes less hyperreactive.

Change in self-awareness. The DMN weakens during practice, and on average in long-term practitioners. This correlates with a less rigid experience of "I." At the limit it becomes what the traditions call the transparent self, the witness, not-self.

In a 2015 review in Nature Reviews Neuroscience, Yi-Yuan Tang, Britta Hölzel and Michael Posner drew all of this into a single model. Their conclusion is cautious: the effects are multiple, they occur in different regions of the brain, and taken together they are consistent with what practitioners themselves describe.

XI. What neuroscience does not prove

Here it is worth stopping. Neuroscience is a powerful instrument. But it has limits, and they should be stated plainly.

It does not prove that meditation leads to enlightenment. Gamma waves, cortical thickness and a quiet amygdala are not a description of mystical experience. They are a description of physical correlates. What a person undergoes inside that experience, the scanner does not show.

It does not confirm the existence of a soul, or of a witness as a metaphysical entity. That meditators have a neural correlate of the transparent self does not mean that this self is an eternal, unchanging thing. It means there is a neural configuration in which the sense of "I" works differently.

It does not validate the subjective reports of Buddhists as truth about the world. When a monk says there is consciousness beyond the body, that is his experience. Neuroscience can say only what neural activity accompanies that experience. What lies at the base of the experience itself is another question.

In popular writing these boundaries sometimes dissolve. "Scientists have proved that meditation opens the door to a higher reality." They have not. Scientists have shown that meditation changes the brain in measurable ways. That is already not a small thing.

XII. What neuroscience does confirm

On the other side, the list of what neuroscience has genuinely shown:

— The adult brain is plastic. The structure of gray matter changes in response to training. — Eight weeks of moderate daily practice produce measurable structural change. — Long practice changes the baseline traits of a brain, not only its current states. — There are specific neural networks responsible for the narrative about oneself, the DMN, and these networks weaken in practitioners. — There are specific networks responsible for observation, the anterior cingulate and the lateral prefrontal cortex, and these strengthen. — The emotional reactivity of the amygdala falls with practice.

That is enough to stop calling meditation mysticism. It is training. Like the piano, or languages. With the same general principles: regularity matters more than intensity, the effect is dose-dependent, changes accumulate over years, but the first shifts take weeks.

XIII. Caveats

Self-criticism is required on several points.

Sample sizes. Hölzel: sixteen people. Lutz and Davidson: eight monks. Lazar: twenty and twenty. These are very small numbers for reliable conclusions. With samples like these it is easy to arrive by chance at a beautiful result.

Publication bias. In science it is easier to publish "we found an effect" than "we found none." In meditation research this is especially acute, since the field is popular and well funded.

The enthusiasm of researchers. Many of the pioneers are practicing meditators themselves — Hölzel, Lazar, Kabat-Zinn, Davidson. This is not a vice, but it is a bias. A scientist who has meditated for ten years has an interest in finding an effect.

Meta-analyses. When the large meta-analyses appeared between 2014 and 2020 — JAMA Internal Medicine in 2014, Perspectives on Psychological Science in 2018 — the conclusions grew more modest. The effects are there, but they are moderate, and they are frequently overstated. Most popular claims of the kind "meditation halves the risk of depression" are exaggerations.

This is no reason to abandon the practice. It is a reason not to expect miracles from it, and not to be disappointed when the miracles fail to arrive.

XIV. What remains after the scanner

Neuroscience does not make the practice meaningful. Nobody goes to meditate because the density of his hippocampus will increase. People go because things are hard, because thoughts keep them awake, because they want to be less reactive, or because they sense there is something behind their ordinary life and want to know what it is.

What neuroscience does is take the label of mysticism off meditation.

For thousands of years the Advaitins said: there is a witness. The Stoics: there is prohairesis. The Mithraists: there is a fire in the human being. The Buddhists: there is an observing consciousness. Jung: there is the Self. And every time an outside observer had the right to shrug: beautiful, but how would you check it.

Now it can be checked. And when it was checked, this is what turned out: people who train with that place inside are physically different from people who do not. Not esoterically different, not different in how they describe themselves, not different in their allegiance to a teaching — different at the level of neural networks, of gray matter density, of cortical rhythms.

Four thousand years of philosophical tradition met twenty years of neuroimaging, and turned out — at the very core — to be about one and the same phenomenon.

XV. Closing the circle

The series closes with this article. It is worth drawing the line.

In Advaita we spoke of the witness — sākṣin — the agency that sees thoughts but is not thoughts; sees the body but is not the body; sees time but is not time. Tat tvam asi, you are That.

In Stoicism we spoke of prohairesis — moral purpose, the agency that chooses its stance toward circumstances; of what does not depend on externals, and is therefore free.

In the myth of Prometheus we spoke of fire taken at the price of suffering — the symbol of the moment when consciousness first separated from natural blindness and gained the ability to see itself.

In the symbol of the mandala we spoke of the circle with a center — the geometry of the place where consciousness finds its center, where the opposites are held together.

Today we can add: these descriptions are consistent with what fMRI and EEG show. In people who train attention and observation, the network of self-narrative weakens, the network of watching the watching strengthens, the cortex that handles bodily signals and the taking of perspective thickens, the reactivity of the amygdala is damped. Not a proof of the witness — but a consistency.

Not faith. Not a philosophical position. Not cultural belonging. Training.

Attention turned inward, repeated thousands of times, leaving a trace in the brain — exactly as the pianist's fingers are trained, as the mathematician's spatial imagination is trained, as the interpreter's ear is trained.

The witness in Advaita, prohairesis in the Stoics, fire in Prometheus, the center of the mandala, a weakened DMN with a strengthened anterior cingulate in a modern meditator — five ways of saying one thing. Each in the language of its own age. Each partial. All of them together, closing on a description.

What remains? Something simple. Inside there is something that can be trained. That something does not depend on what it is called. One can begin today — twenty minutes, as in Tang; eight weeks, as in Hölzel; ten years, as in Lazar; a whole life, as in a Tibetan monk.

Through any of these windows, the same thing is looking out.

Andrey Zarinsh & Maria Evilkut

United Wisdom · One source — many languages

References inherited from the archive

These links were recovered from the supplied archive metadata; they were not newly read or verified for this translation.

- PDF-L001 — archive PDF page 4.
- PDF-L002 — archive PDF page 5.

- PDF-L003 — archive PDF page 6.
- PDF-L004 — archive PDF page 7.
- PDF-L005 — archive PDF page 7.
- PDF-L006 — archive PDF page 8.
- PDF-L007 — archive PDF page 10.