Max Shen

Pain researcher and coach with a background in AI research and policy

Max Shen is a computer-science-trained researcher, writer, and pain coach who created Debug Your Pain with Tanner Holman. His work combines chronic-pain inquiry, movement, interviews, models, and individual coaching. Earlier engineering and policy publications use Max Langenkamp; his professional transition announcement connects that record to his present practice. [1]; [2].

MIT engineering, AI-policy research, and biosecurity preceded the present practice. Pain became disabling during graduate study, while his later fellowship biography records pain research and SecureDNA work continuing together. The change was an overlap of occupations before it became a departure. [3].

[External image / 外部圖片] Max Shen, wearing a light-colored shirt, photographed against green foliage.
Max Shen in the photograph on his personal website. Max Shen / official personal website; photographer not named · Rights retained by the source and respective rights holders; no open license stated. · Image source

Learning to build, with other people

As a teenager in New Zealand, he recited Wilfred Owen’s Dulce et Decorum Est to an audience he remembers as five thousand. His August 2026 letter also recalls burying Val, the German Shepherd who had accompanied nine childhood years. These memories make poetry, performance, affection, and loss part of his early life without supplying a complete childhood biography. [4].

The earliest detailed engineering accounts show a student learning by attempting projects whose attraction precedes a mature sense of their difficulty. Rapspeare, which he calls his first real programming project, attempted to generate a Shakespearean rapper. His account, first written in August 2017 and published on Medium in May 2018, credits Andrej Karpathy’s recurrent-network demonstrations and encouragement from Soren Bouma. He combined Shakespeare from Project Gutenberg with rap lyrics, experimented with alternating their lines, and adapted existing Lasagne/Theano code. Saving trained weights let friends generate text without paying for a GPU. He reduced training time when results stopped improving enough to justify the cost of an AWS machine; he also admits leaving an unwieldy main function intact because he feared breaking working code. This is an unusually useful account of technical formation: dependence on others’ code, financial constraints, incomplete understanding, small improvements, and the pleasure of showing

friends something that runs. [5].

Another 2018 project was Sleepsense, built by a team in MIT's introductory interconnected embedded-systems course, 6.08. An ESP32, motion sensor, button, and rechargeable battery recorded movement on a mattress; server-side Python processed the readings, and a website displayed a night's history. The work involved mundane but consequential engineering: distinguishing acceleration from background noise, preventing a battery-estimation loop from hanging, reducing Wi-Fi use, and making the recording light tolerable at night. Attempts to produce a meaningful sleep score with more elaborate models failed, and the team settled on a heuristic based on movement density. Langenkamp acknowledges that some nights received poor scores despite feeling reasonably restful and that physiological measurements would be needed for stronger scientific validation. This was a functioning student device with a provisional interpretation of its measurements, rather than a demonstrated instrument for measuring sleep stages. [6].

In January 2019 he and his friend Tuomas Oikarinen spent a month as visiting researchers at the University of Sannio in Benevento, Italy, under Silvia Ullo's guidance and alongside graduate researchers Maria Pia Del Rosso, Alejandro Sebastianelli, and Federico Piccirillo. They used satellite imagery and a three-dimensional convolutional network to investigate landslide detection. Obtaining usable training examples proved harder than designing the network. Original Sentinel-2 downloads were enormous for their laptops, and manual cropping was slow; Google Earth Engine enabled a more workable download pipeline. They assembled about 120 image pairs and tried augmentation, discovering that added noise could make training slower and worse. The reported accuracy, approximately 0.65, remained modest. His account ends with a paper submitted to IGARSS and a visit to the European Space Agency. The episode records collaboration, travel, data preparation, and an imperfect result—not an isolated algorithmic breakthrough. [7].

Later coursework retained this practical character. A 2021 report by Alice Zhang, Maximillian Langenkamp, and Sule Kahraman describes building a USB-security demonstration with a Raspberry Pi Zero W after discovering that other boards lacked the needed capabilities. Their obstacles included incompatible devices, software libraries, and differences between operating systems. The project belongs to MIT's 6.858 computer-systems-security course. It helps explain the eventual credibility of a policy researcher discussing hardware: he had encountered how a specification meets the recalcitrance of actual components. [8].

His selected-projects page names several institutional settings without furnishing a complete employment calendar: working with TensorFlow at Cerebras; reinforcement-learning models of human behavior in MIT's Computational Cognitive Science Lab; data analysis and policy at Georgetown's Center for Security and Emerging Technology; a role at In-Q-Tel whose substance he marks as redacted; and work on 5G and AI policy at the White House Office of Science and Technology Policy. The list also mentions explorations in economics and philosophy. It establishes a range of actual technical and policy encounters, but does not justify assigning a precise duration, seniority, or set of

responsibilities to every affiliation. [9].

Engineering enters policy and institutional research

The April 2020 preprint Hiring Fairly in the Age of Algorithms, with Allan Costa and Chris Cheung, addresses automated hiring through technical and legal questions together. It examines datasets, models, definitions of fairness, commercial applications, employment law, and the difficulty of scrutinizing systems protected as trade secrets. The authors propose algorithmic transparency reports for employers using such software. The object of concern is not merely whether an algorithm performs well on a benchmark: it is whether people affected by a consequential decision can understand the system and its risks. This was a student-authored research proposal, not enacted regulation. Its combination of technical inspection and institutional design would recur in his subsequent work. [10].

At CSET, Langenkamp and Melissa Flagg coauthored AI Hubs: Europe and CANZUK, published in April 2021. The report maps AI-related talent and investment across Europe, Canada, Australia, New Zealand, and the United Kingdom, presenting these regions as resources for international cooperation in a strategic environment shaped by competition with China. Its intended readers are policymakers making choices about partnerships and capacity. The report shows him translating dispersed technical activity into a geographic and institutional picture; it also locates him within Washington’s emerging AI-policy research world. That framing is a policy argument, rather than a neutral ranking of countries’ intelligence or worth. [11].

His MIT master’s thesis, How Open Source Machine Learning Software Shapes AI, records a BS in computer science and engineering in 2021 and an MEng in electrical engineering and computer science in May 2022. Dylan Hadfield-Menell supervised it. The acknowledgments show how the work became possible: conversations with Tim Hwang and Teddy Collins while Langenkamp worked remotely for CSET prompted an application for funding; Open Philanthropy supplied support. Hadfield-Menell helped narrow an initially sprawling inquiry. Daniel Yue contributed equally to the research for the first half, combining technical persistence with an openness to unconventional ideas. For the interview component, Langenkamp conducted eight formal interviews of about fifty minutes and fifteen informal conversations, preparing notes promptly and using iterative thematic analysis. These details identify the thesis as an apprenticeship in interviewing and institutional explanation as well as computational analysis. [12].

The related paper, How Open Source Machine Learning Software Shapes AI, appeared at the 2022 AAAI/ACM Conference on AI, Ethics, and Society, with Daniel N. Yue of Harvard Business School as coauthor; it explicitly states equal contributions. Examining 139 repositories alongside practitioner interviews, the authors ask who pays for an infrastructure that enables much more visible AI applications. They describe software’s roles in standardization, experimentation, and community formation, and compare estimated production costs with potential economic benefits. Their striking estimate that

benefits might exceed costs by about a hundredfold depends on broad economic projections and assumptions, not measured profits returned to these particular repositories. The productive tension is between software's usefulness and the limited incentives to maintain it. The contribution is shared with Yue; neither the thesis nor the conference artifact warrants treating their work as a solo accomplishment. [13]; [14].

A subsequent essay with Hadfield-Menell translates the project for a wider audience. It identifies Nadia Asparouhova's Working in Public as an inspiration and emphasizes the mismatch between widespread dependence on machine-learning software and rewards for maintenance. The essay's method is to make an enabling layer visible: what appears to be the achievement of an AI system also rests on tools, standards, volunteer labor, organizational priorities, and funders. The thesis had given Langenkamp a way to follow those connections through quantitative traces and conversations with practitioners. That experience matters later when qualitative interviews return in his pain research, although the questions and evidentiary problems change. [15].

Losing ordinary abilities during graduate study

Shen locates his disabling pain episode in graduate school, without a sufficiently secure onset date to assign it to a particular semester. In his first-person account, several intense days chasing a programming bug were followed by painful wrists and forearms, stiff fingers, and rapidly deteriorating typing. He had already invested in an ergonomic keyboard and careful posture. A doctor suggested tendonitis or repetitive strain injury; physiotherapy brought exercises. Four weeks later he opened doors with his forearms and ate frozen meals because he could not chop vegetables. Fear and anger accompanied the possibility of living with dictation. He ordered a dozen popular books on pain, read research, and experimented with awareness practices. He reports substantial improvement, followed by the return of typing, pull-ups, guitar, and cooking; years later, one mild flare meant a week's pause from piano. This is his account of illness and recovery, not an independently reconstructed medical record. Its most concrete stakes are the things that stopped being possible and became possible again. [16].

A later professional transition post adds that he actually used voice commands and feared having to leave his graduate program. He describes well-intentioned practitioners whose explanations confused him and advice that sometimes made things worse. Reading neuroscience, self-experimentation, and conversations with Professor Fan Wang helped redirect his attention. His retrospective time estimates vary: this post describes becoming functional in roughly two weeks and substantially better within two months, while the longer recovery account places pain-free typing and pull-ups around three months. These are approximate recollections of stages of recovery, not a clinical timeline to force into exact agreement. The post presents the return of handstands as a practical marker of the distance traveled. [17].

The interruption changed the stakes of his inquiry: an explanation now had to make

ordinary activity possible again. Regained pleasures also mattered without proving a theory.

His own interpretation emphasizes fear, expectations, attention, and the nervous system’s protective activity. An unsuccessful structural explanation became part of his reason for searching elsewhere. That experience should not be generalized into a portrait of all clinical pain science: the International Association for the Study of Pain’s 2020 definition explicitly recognizes biological, psychological, and social influences and distinguishes pain from nociception. The distance between Shen’s encounters and those principles is a real part of his story, but does not show that the whole field lacks them. [18].

Biosecurity continues alongside the new inquiry

The new interest did not immediately replace his earlier occupation. His older homepage simultaneously lists SecureDNA work on safeguards for DNA printers, chronic-pain modeling, writing, and movement exploration. The Foresight Institute’s 2025 fellowship biography describes him as leading SecureDNA’s policy and hardware initiative concerning security standards for benchtop synthesis and as an MIT research affiliate investigating chronic pain’s onset and resolution. These are records of a period of overlap, read alongside his later departure rather than as guarantees of current job titles. [19]; [20].

The 2023 working paper Clarifying the Problem of DNA Screening belongs to this period of defining what a security intervention must achieve. In the same year, Richard Moulange, Langenkamp, Tessa Alexanian, Samuel Curtis, and Morgan Livingston published Towards Responsible Governance of Biological Design Tools. That preprint argues that tools for designing biological systems require a governance analysis of their own, rather than importing every assumption from language-model debates. It considers development practices, evaluation, access, transparency, cybersecurity, and resilience. Its public presentation was at the NeurIPS workshop on Regulatable ML, a distinct venue from the main conference. The work documents a collaborative policy agenda; it does not settle the capabilities or risks of every subsequent tool. [21]; [22].

The published Applied Biosafety article “Practical Questions for Securing Nucleic Acid Synthesis” provides a particularly clear view of the collaborative environment. Sophie Rose, Tessa Alexanian, Langenkamp, Helena Cozzarini, and James Diggans organized questions arising from guideline review, policy research, and expert conversations. They consider customer and sequence screening, domestic and international coordination, commercial pressures, and benchtop devices. Disclosure statements place the authors across nonprofit security initiatives and commercial synthesis: Langenkamp at SecureDNA, Alexanian at IBBIS and the Nuclear Threat Initiative, Diggans at Twist Bioscience and the industry consortium. These affiliations help explain why implementation questions matter. A recommendation must function amid providers’ different obligations and incentives. First published online in March 2024, this article represents a team’s synthesis of practical problems, rather than an individual invention of screening policy. [23].

His December 2024 Institute for Progress essay, “Securing Benchtop DNA Synthesizers,” extends the work into hardware policy. He describes consulting more than a dozen industry professionals, reviewing technical literature, analyzing system architecture and security threats, and working with professional forecasters through Metaculus. The recommendations include standards and certification capable of addressing devices and their screening arrangements. The essay illustrates the actual labor behind a policy publication: asking people how systems operate, translating their answers into a threat model, and considering measures that could be audited or enforced. Forecasts in that process remain estimates of possible futures, and the publication itself is not evidence that its proposed standards were adopted. [24].

This phase has its own institutional world: synthesis providers, nonprofits, security researchers, policy institutes, and workshop audiences. It also supplies a practical continuity with the open-source thesis, since an intervention depends on infrastructure and people who must actually carry it out. The continuity does not make pain another policy problem. In coaching, the person experiencing the phenomenon can change it while describing it; interpretation and interaction become part of the object of inquiry.

Writing, attention, and life beyond the research agenda

Shen’s public writing was already broader than either biosecurity or recovery. His older homepage recommends the soundtrack to Orfeu Negro, Philip Glass’s Words Without Music, and Yeats’s poetry, plays, and philosophy. Alongside blindfolded dancing, it includes scallion pancake with egg and riding a bicycle gently downhill with hands off the handlebars. He says he likes meeting people. Such a list gives a less instrumental view of his interests than a publication bibliography can: music, food, balance, reading, and company appear as pleasures, rather than materials gathered to demonstrate a single intellectual position. [25].

His February 2024 essay “In the World, Like Lovers” distinguishes merely occupying a place from experiencing familiarity and participation within it. A home differs from a domicile; absorbing work and close relationships change how the surrounding world becomes available. The essay borrows intimacy as a way to explain meaning, rather than presenting a documented romantic biography. It concerns the ordinary experience of being involved in something before stepping back to explain it. [26].

The following essay, “We Who Glimpse Icebergs,” asks readers to notice the usually unnoticed work of bodily experience—seeing one’s nose in the visual field, swallowing, and sensing oneself in a familiar setting. A breakdown makes background processes conspicuous. Langenkamp uses such observations to question accounts of people as detached thinkers making explicit choices from complete representations. These essays reveal an intellectual interest in experience that is not confined to the clinic, and a writing method that begins with something a reader can notice rather than with a technical definition. [27].

In an undated essay on technological inevitability, he argues with technically minded friends who treat brain-computer interfaces or advanced AI as unavoidable because they are physically possible and commercially attractive. His objection is that those conditions do not determine what societies will build; political action, institutional choices, and competing purposes also matter. This is his argument about technological futures, not a prediction established by the historical examples he invokes. It situates his policy interests within a wider concern about people surrendering judgment to narratives of inevitability. [28].

A June 2025 essay remembers a dirt-bike crash at seventeen: his helmet struck the windscreen, and his body moved automatically before he found himself unhurt. He uses the memory to reflect on vulnerability and care as different ways of inhabiting the same world. The crash is a publicly narrated event, not evidence for a retrospective medical diagnosis. The writing makes room for mortality and gratitude alongside professional questions. [29].

His April 2026 essay about feeling stuck in a simulation recalls an episode after graduation, on a flight to Brazil, when intense concentration on breathing made his surroundings feel flattened and his body strangely remote. A phone conversation with his friend Nick helped him feel connected again. He contrasts compulsive scrutiny of whether reality is real with participating in it, including Afro-Brazilian dance. This is a first-person account of an unsettling experience and a friendship that mattered at that moment; it does not establish a clinical condition. It also shows that his interest in embodiment includes travel, conversation, and dance rather than only exercises prescribed for pain. [30].

Building a public practice with Tanner Holman

The Debug Your Pain authorship page explains the site's origin as a response to the hundreds of hours Shen says he spent finding useful information. He wanted to make available the resource he had needed at the beginning. It credits his friend and colleague Tanner Holman's knowledge of physiology and movement as complementing his technical perspective. It also identifies Shen's training in Pain Reprocessing Therapy and earlier work in Harvard's Moral Psychology Research Lab and MIT's Computational Cognitive Science Lab. Research experience, a particular training program, and movement knowledge are distinct preparations; the page does not establish a medical license. Nor does it identify which collaborator wrote every lesson or made every coaching decision. [31].

Their May 5, 2025 announcement, "Pay Us to Resolve Your Chronic Pain," turns that educational project into an explicit service. They said they had been coaching for six months and offered an initial conversation to assess fit. Holman's account includes a period in a back brace and reconsideration of his career. The offer targets persistent pain, including cases where stress affects symptoms or previous medical approaches have not helped, while explicitly stating that the coaches are neither doctors nor physiotherapists. Testimonials include a client reporting an immediate shoulder-pain change during a guided

awareness exercise and a later follow-up. Such stories explain what the coaches and clients find persuasive; without a comparison group or systematic follow-up they do not establish the service’s effectiveness. The announcement discloses a commercial offer but no price schedule, revenue, or basis for estimating the founders’ income. [32].

The site’s Pain Debugging Protocol shows how the general explanation becomes work for a reader. It combines understanding, processing, and integration, acknowledges assessment for structural problems, and encourages keeping a document of observations, hypotheses, and experiments. Explaining one’s model aloud resembles the programmer’s practice of talking through a bug. Other elements concern finding a calmer setting, developing regular awareness practice, and returning gradually to movements and activities that have become threatening. Free guided material and a contact route for individual coaching connect the writing to ongoing support. The labor lies partly in helping someone notice what happens during a particular action, rather than merely accept a proposition about the nervous system. This assembled protocol has not been validated by the trial evidence for another named therapy simply because it shares some of its methods. [33].

Shen’s March 2025 essay on John Sarno makes the relationship between practical benefit and explanatory disagreement explicit. He credits Healing Back Pain with helping his recovery, but challenges the completeness of Sarno’s account of repressed emotion, reduced blood flow, and pain. Phantom-limb pain is one of his tests of its reach. He prefers a predictive account in which expectations, context, and protective responses can help perpetuate symptoms. The essay also argues for reaching different publics through varied forms of pain education. Its confident estimates of the number helped by Sarno rely heavily on testimonials and extrapolation; they should remain Shen’s claims. The useful distinction in his own practice is that he can value an intervention without accepting every mechanism attached to it. [34].

The nearest major controlled study in his cited repertoire is the randomized Pain Reprocessing Therapy trial by Yoni Ashar and colleagues, published online in September 2021. Among 151 adults selected for primary chronic back pain, 66 percent assigned to PRT were pain-free or nearly so after treatment, compared with 20 percent in the placebo arm and 10 percent receiving usual care; benefits were followed for a year. That is meaningful evidence for a particular intervention and population. It is not a trial of Shen and Holman’s coaching, all repetitive-strain complaints, or chronic illness generally. Keeping that limit close to the service prevents the profile from borrowing clinical authority the practitioners have not themselves demonstrated. [35].

There is also an institutional question here. Research papers, a free website, a newsletter, guided practices, and paid individual time reach people in different ways. The writing can circulate without the author being present; coaching depends on attendance, trust, attention, and adaptation to a person. The public record makes this combination visible without establishing how large the practice is or how it supports the broader research. Its small scale and uncertain economics should not be filled in with either a heroic story of

independence or an allegation about commercial motives.

Movement, models, and public teaching

Shen identifies Fan Wang, the MIT neuroscientist, as a mentor and influence on his inquiry. His current site places active inference, phenomenology, and Bayesian statistics alongside Pain Reprocessing Therapy, Judith Blackstone's Realization Process, and Moshe Feldenkrais. These are different intellectual and practical traditions, not interchangeable credentials. Their coexistence helps locate an inquiry moving between formal descriptions of learning and first-person bodily experience. The site establishes his publicly named repertoire; it does not show that Wang's laboratory endorses his coaching or every claim on his website. [36]; [37].

Enactive approaches to pain predate his contribution. Peter Stilwell and Katherine Harman's 2019 paper treats pain as an embodied process of sense-making in a person's world; Sabrina Coninx and Stilwell's 2021 article considers changes in the field of possible actions. These works offer a way to describe how pain reorganizes what a person can do and how they encounter their surroundings, beyond a model of an isolated signal. They are intellectual context for Shen's work, not clinical validation of his particular service. [38]; [39].

In the 2025 manuscript Beyond Biomechanics, Shen writes with Jozef Frucek, cofounder with Linda Kapetanea of Fighting Monkey. They propose combining active inference and enactive cognition to understand health and movement as adaptive engagement across several scales. The practice examples are concrete: exploratory "Zero Forms" and work with a ball suspended on an elastic line ask participants to respond to changing movement and relationships, rather than increasing load alone. The paper uses photographs from a May 2025 Manhattan workshop. Its disclosures say Shen participated in Fighting Monkey workshops without payment, sponsorship, or commissioning for the manuscript, while Frucek earns income from its workshops and courses. These details locate a relationship between a participant-researcher and a practitioner who also operates a business. The text is an exploratory conceptual manuscript, not a controlled demonstration of Fighting Monkey's health effects. [40].

The computational side remains inspectable through the active-inference-pain repository, hosted under Daniel Gretton's GitHub account and linked by Shen. Its simplified agent learns about danger through signals in an environment, with parameters and beliefs available for examination. Such a model can make a proposed feedback mechanism easier to think about and criticize. It does not reproduce the complexity of a patient's history or identify the cause of a particular person's pain. The repository also prevents the computational work from becoming a claim resting entirely on evocative language; readers can inspect an actual artifact. Its public hosting and collaboration do not warrant assigning Shen sole authorship. [41].

His January 2025 guest lecture, "What Might Our Minds Be, if Not Computers?", belonged

to MIT’ s 6.S094 Computational Psychology course. The official program, held January 27–31, lists him as a research affiliate and speaker among people considering reinforcement learning, Bayesian modeling, and different perspectives on psychology. This was a teaching appearance within a computationally oriented program, not a faculty appointment. A public conversation with Adam Safron and Michael Levin and an August 2025 Joint Dynamics interview with Andrew Cox place the work before adjacent cognitive-science and movement audiences. The latter is an approximately ninety-five-minute conversation hosted in the orbit of a Hong Kong physiotherapy, massage, and training practice. These appearances document circulation across communities; their existence alone should not be treated as endorsement or as evidence of clinical outcomes. [42]; [43]; [44]; [45].

His December 2025 presentation on chronic illness makes a broader ambition visible in the original slides. It moves from biochemical pathways to diagrams of possible stable states, identifies a proposed correspondence among priors, attractors, and set points, and uses bodily demonstrations to introduce regulation. A hand-drawn landscape illustrates an illness state becoming easier to leave during a proposed “reset.” The final questions ask which bodily variables define the states, how attention relates to plasticity, and what makes a useful reset. These are programmatic questions, not resolved mechanisms across the conditions grouped in the slides. The associated published essay explicitly identifies itself as generated by Claude from the talk transcript; the original slide deck is the firmer basis for describing his presentation. Its analogies suggest a research direction but cannot substitute for condition-specific evidence. [46]; [47].

Leaving Cambridge and beginning another phase

In October 2025 he described leaving Cambridge, his home of eight years, for an open-ended journey. He planned music and fruit, followed by offerings of memories, songs, or objects. Rehearsing with eight friends in mind showed that his room lacked enough seating, so he planned for everyone to leave and reenter the living room holding a flower. His account of the gathering remembers tears and relief. The farewell gives Cambridge a social geography of friends and shared rooms alongside its institutional addresses. [48].

His February 19, 2026 retrospective on the preceding year explicitly pairs leaving biosecurity with moving from Cambridge to San Francisco. It also names journeys to Athens, Ann Arbor, Woodstock, Valencia, and Los Angeles to meet people he calls sages, without providing enough information to reconstruct their identities or formal roles. The addendum describes a current project of qualitative interviews with people who attribute recovery from pain to reading Sarno, alongside coaching and work on scientific accounts of chronic illness. That is a stated research agenda, not a published result. Selecting people who already report recovery can investigate their descriptions and possible processes; it cannot establish how frequently readers recover or compare the book with other interventions. [49].

By August 2026 he was writing from San Francisco, listening to Donny Hathaway and

thanking friends. [50]. His recent reading included Frances Yates on Giordano Bruno and the Hermetic tradition; an earlier poem places dancing, friendship, and helping others among his reasons to practice. [51]; [52]. These interests do not establish a religion or romantic biography.

The most defensible account of Shen’s career is therefore a conjunction of particular skills, institutions, and relationships. Student engineering taught him to work with borrowed code, imperfect measurements, collaborators, and limited resources. Policy and thesis research taught him to ask how systems function through institutions and maintenance. Pain interrupted ordinary competence and gave him a reason to seek explanations that could change lived activity. Holman and Frucek brought different practical repertoires; Wang and other interlocutors offered scientific conversation. Writing, models, teaching, and coaching now provide different ways to pursue the inquiry.

These continuities do not make every event inevitable or prove that one explanation governs them all. His current direction is consequential precisely because it remains a practice in formation: a technically educated person learning to work with subjective accounts, bodily activity, and other people’s need for relief. Its future contribution will depend on the quality of that work and on whether its public claims stay proportionate to what the research and practice actually establish.

Research checked October 3, 2026. Public first-person accounts are attributed as such; source links identify the underlying papers, practice materials, institutional records, and authored essays.

---

1. Current research and practice site↩
2. professional transition↩
3. 2025 Foresight fellowship biography↩
4. “I Am Blessed by You,” August 20, 2026↩
5. “How Would Shakespeare Have Rapped?” May 25, 2018↩
6. “The Difficulty with Tracking Sleep,” May 30, 2018↩
7. “Deep Learning with Satellite Data,” February 2, 2019↩
8. Team report, August 1, 2021↩
9. Selected jobs and projects↩
10. Preprint, April 15, 2020↩
11. CSET report and full PDF↩
12. MIT thesis, title pages, acknowledgments, and methods, May 2022↩
13. AIES paper↩

14. publication DOI↩
15. “How Open Source Machine Learning Software Shapes AI,” with Dylan Hadfield-Menell↩
16. “Beyond the Wrist: Debugging RSI”↩
17. Professional transition account↩
18. IASP revised definition and notes, July 16, 2020↩
19. Earlier homepage↩
20. 2025 fellowship biography↩
21. DNA-screening working paper↩
22. biological-design-tools preprint, November 2023↩
23. Article, affiliations, and disclosures↩
24. Institute for Progress, December 10, 2024↩
25. Personal homepage↩
26. Essay, February 26, 2024↩
27. Essay, March 4, 2024↩
28. “Technology Is Not Inevitable”↩
29. “Is the World a Safe Place?”, June 5, 2025↩
30. Essay, April 10, 2026↩
31. Authorship and background↩
32. Coaching announcement↩
33. Current practice page↩
34. “Useful Stories and Pain Relief,” March 8, 2025↩
35. Original trial, JAMA Psychiatry↩
36. Current site↩
37. Wang’s MIT research profile↩
38. Stilwell and Harman↩
39. Coninx and Stilwell↩
40. Original manuscript, examples, photographs, and disclosures↩
41. Code and documentation↩
42. MIT course program↩

43. lecture recording↩
44. Safron–Shen–Levin recording↩
45. Joint Dynamics episode 133↩
46. Original slides↩
47. talk publication and recording, December 16, 2025↩
48. “How to Run a Goodbye Ritual,” October 14, 2025↩
49. “2025 in 25 Minutes”↩
50. San Francisco letter↩
51. “Hermes’ Lament,” August 7, 2026↩
52. “Why Do I Practice?”, July 24, 2025↩