

Tim Ferriss recommends John Gardner’s ”Personal Renewal” speech; Sarah Guo backs Nick Cammarata on interpretability

Recommended Reading from Tech Founders

2026-10-16

Tim Ferriss recommends John Gardner’s “Personal Renewal” speech; Sarah Guo backs Nick Cammarata on interpretability

By Recommended Reading from Tech Founders • October 16, 2026

Recommendations this period include a 1990 speech on lifelong learning, an argument that interpretability may be solvable without frontier-lab compute, a No Priors episode with Greg Brockman, and a Kalanick clip on why Uber didn’t buy Lyft.

Tim Ferriss: John Gardner’s “Personal Renewal”

In a video about why he won’t let AI write his next book, Ferriss offered “one piece of long form” for viewers: **“Personal Renewal,”** a speech John Gardner gave to McKinsey & Company in November 1990. Ferriss read the transcript and called it “exceptional” [1]. He says it is “full of gems” about continual learning throughout life, though it takes a while to get to them. An audio version is available by searching “Personal Renewal John Gardner” [1]. The passage he read aloud pushes back on “the myth that learning is for young people.” It tells how Cervantes wrote *Don Quixote* in prison at 53, after a life of failures [1]. His advice: “maybe print it out single task get away from notifications.” He has read it several times and says, “I just get something from it every time” [1].

Sarah Guo: Nick Cammarata on interpretability

Guo quoted the claim that “you can’t do critical research without the compute scale of the top frontier labs” and answered, “Nick: hold my beer,” linking to Cammarata’s post [2]. Cammarata co-authored early mechanistic interpretability work at OpenAI and came away pessimistic. In OpenAI’s first large-scale

automated interpretability project, only about 2% of neurons were legible [3]. He now says that with modern agentic tools, “the vast majority in models up to trillions of parameters are legible.” He expects ambitious interpretability to be solved “by the time we need it most, and possibly quite soon,” calling it possibly “a bitter lesson for interpretability: the simple thing, at absurd scale, works.” His post links to Goodfire [3].

Guo and Elad Gil: No Priors with Greg Brockman

Both hosts promoted their own No Priors episode with OpenAI co-founder and president Greg Brockman, so these are self-promotion rather than outside endorsements. Guo framed it as “the Post-AGI Rules of the Game in Cyber and Science.” Chapters cover AI and mathematics, biology data bottlenecks, drug discovery, regulation, OpenAI’s proactive agent Dots, how UI will evolve, and “The Change in the Physics of Cybersecurity” [4, 5]. Gil shared the same conversation with a similar chapter list [6]. Guo’s main takeaway is the security story. OpenAI “took 25% of its production engineers off every project” to work only on security, and its top model “hunted in their infrastructure until it saturated (stopped finding new issues).” Her question: “Will everyone else follow?” [7]

Jason Lemkin: Kalanick on why Uber didn’t buy Lyft

Lemkin quoted Guille Florvs’s post with a clip from Atoms’ launch. In it, Travis Kalanick says he passed on acquiring Lyft, even though Uber was spending roughly a billion dollars a year fighting it, because “Lyft’s culture was very different than ours” [8]. Lemkin’s lesson for founders is that in M&A meetings, being judged “a >fit<” matters more than price, since acquirers “want to buy a 100% bought-in, A team.” He called it “A big miss for Lyft.” Lyft is down about 77% from its IPO price, he notes, while Uber is up about 54% [9].

Shorter mentions

- **Keith Rabois on pro rata (via 20VC):** In the latest 20VC episode, one panelist said Rabois gave “the only thoughtful answer” in the pro-rata debate. As the panelist recalls it, investors could drop pro rata but would have to pay “like a half the price” in the early round. The panelist’s view: “You can value an option and deduct it from the price” [10].
- **Bill Gurley** called Clay Haddock’s post “a great story” and said his Runnin Down a Dream Foundation is looking for more partnerships [11]. Haddock describes a cold email that turned into a partnership with his Woihanble Fund, which backs young entrepreneurs on the Pine Ridge Reservation. The foundation sent books for the fund’s cohorts [12].
- **Elon Musk** shared Brad Torgersen’s video post on HAL (“HAL was told to lie... So, he couldn’t function”). Musk added: “Truth-seeking is the most important aspect of AI/SI safety” [13, 14].

Sources

1. Why I Refuse to Let AI Write My Next Book
2. X post by @saranormous
3. X post by @nickcammarata
4. X post by @saranormous
5. X post by @saranormous
6. X post by @eladgil
7. X post by @saranormous
8. X post by @guilleflorvs
9. X post by @jasonlk
10. Valar Atomics Sues Day One Ventures | Meta's Back to the AI Frontier |
NVIDIA's \$30B Reflection Deal
11. X post by @bgurley
12. X post by @ClayHaddock
13. X post by @BradRTorgersen
14. X post by @elonmusk