

STEVEN LEVY BUSINESS NOV 10, 2023 9:00 AM

Fei-Fei Li Started an AI Revolution by Seeing Like an Algorithm

Researcher Fei-Fei Li's ImageNet project provided the feedstock for the deep learning boom that brought the world ChatGPT and other world-changing AI systems.



Fei-Fei Li PHOTOGRAPH: MATT WINKELMEYER/GETTY IMAGES

 SAVE

Early in the pandemic, an agent—literary, not software—suggested Fei-Fei Li write a book. The approach made sense. She has made an indelible mark on the field of artificial intelligence by heading a project started in 2006 called ImageNet. It classified millions of digital images to form what became a seminal training ground for the AI systems that rock our world today. Li is currently the founding codirector of Stanford’s Institute of Human-Centered AI (HAI), whose

very name is a plea for cooperation, if not coevolution, between people and intelligent machines. Accepting the agent's challenge, Li spent the lockdown year churning out a draft. But when her cofounder at HAI, philosopher Jon Etchemendy, read it, he told her to start over—this time including her own journey in the field. “He said there’s plenty of technical people who can read an AI book,” says Li. “But I was missing an opportunity to tell all the young immigrants, women, and people of diverse backgrounds to understand that *they* can actually do AI, too.”

This is an edition of Steven Levy's *Plaintext* newsletter.



**AI. STARTUPS.
CRYPTO. FAANG.
TECH CULTURE.**

**SUBSCRIBE TO GET STEVEN LEVY'S
FEARLESS LOOK AT ALL THINGS TECH.**

SIGN UP for Plaintext and tap Steven's unique insights and unmatched contacts for the long view on tech.

Li is a private person who is uncomfortable talking about herself. But she gamely figured out how to integrate her experience as an immigrant who came to the United States when she was 16, with no command of the language, and overcame obstacles to become a key figure in this pivotal technology. On the way to her current position, she's also been director of the Stanford AI Lab and chief scientist of AI and machine learning at Google Cloud. Li says that her book, *The Worlds I See*, is structured like a double helix, with her personal quest and the trajectory of AI intertwined into a spiraling whole. "We continue to see ourselves through the reflection of who we are," says Li. "Part of the reflection is technology itself. The hardest world to see is ourselves."

The strands come together most dramatically in her narrative of ImageNet's creation and implementation. Li recounts her determination to defy those, including her colleagues, who doubted it was possible to label and categorize millions of images, with at least 1,000 examples for every one of a sprawling list of categories, from throw pillows to violins. The effort required not only technical fortitude but the sweat of literally thousands of people (spoiler: Amazon's Mechanical Turk helped turn the trick). The project is comprehensible only when we understand her personal journey. The fearlessness in taking on such a risky project came from the support of her parents, who despite financial struggles insisted she turn down a lucrative job in the business world to pursue her dream of becoming a scientist. Executing this moonshot would be the ultimate validation of their sacrifice.

The payoff was profound. Li describes how building ImageNet required her to look at the world the way an artificial neural network algorithm might. When she encountered dogs, trees, furniture, and other objects in the real world, her mind now saw past its instinctual categorization of what she perceived, and came to sense what aspects of an object might reveal its essence to software. What visual clues would lead a digital intelligence to identify those things, and further be able to determine the various subcategories—beagles versus greyhounds, oak versus bamboo, Eames chair versus Mission rocker? There's a fascinating section on how her team tried to gather the images of every possible car model. When ImageNet was completed in 2009, Li launched a contest in which researchers used the dataset to train their machine learning algorithms, to see whether computers could reach new heights identifying objects. In 2012, the winner, AlexNet, came out of Geoffrey Hinton's lab at the University of Toronto and posted a huge leap over previous winners. One might argue that the combination of ImageNet and

AlexNet kicked off the deep learning boom that still obsesses us today—and powers ChatGPT.

What Li and her team did not understand was that this new way of seeing could also become linked to humanity's tragic propensity to allow bias to taint what we see. In her book, she reports a “twinge of culpability” when news broke that Google had mislabeled Black people as gorillas. Other appalling examples followed. “When the internet presents a predominantly white, Western, and often male picture of everyday life, we're left with technology that struggles to make sense of everyone,” Li writes, belatedly recognizing the flaw. She was prompted to launch a program called AI4All to bring women and people of color into the field. “When we were pioneering ImageNet, we didn't know nearly as much as we know today,” Li says, making it clear that she was using “we” in the collective sense, not just to refer to her small team. “We have massively evolved since. But if there are things we didn't do well; we have to fix them.”

On the day I spoke to Li, *The Washington Post* ran a long feature about how bias in machine learning remains a serious problem. Today's AI image generators like Dall-E and Stable Diffusion still deliver stereotypes when interpreting neutral prompts. When asked to picture “a productive person,” the systems generally show white men, but a request for “a person at social services” will often show people of color. Is the key inventor of ImageNet, ground zero for inculcating human bias into AI, confident that the problem can be solved? “*Confident* would be too simple a word,” she says. “I'm cautiously optimistic that there are both technical solutions and governance solutions, as well as market demands to be better and better.” That cautious optimism also extends to the way she talks about dire predictions that AI might lead to human extinction. “I don't want to deliver a false sense that it's all going to be fine,” she says. “But I also do not want to deliver a sense of gloom and doom, because humans need hope.”

Li believes that an important element in developing AI further will be funding to make sure the next breakthroughs—moonshots like ImageNet—come from academia and government, not just commercial enterprises focused on profit and loath to share with the public. This past June, she was among a small group of AI scientists, experts, and critics who met face-to-face with Joe Biden when the president visited San Francisco. She urged that the government fund more AI moonshots. “If we deprive the public sector of the resource, we're doing a disservice to the next generation,” she told him. (Note she didn't say that such

deprivation was akin to murder, as Marc Andreessen charged in his recent 5,200-word Ayn Rand-ian belch.)

And what did the president say to Li when she proposed such moonshots? “Well, he didn’t write a check right there,” she says. “But he was engaged.” She points out that Biden’s recent sweeping executive order on AI has a section on public sector investment. Li’s not one to take a public victory lap, but she seems to have got the result she wanted. Maybe that investment makes it more likely that the next ImageNet-scale advance in AI will come from someone like Li, who didn’t jump to Google or some startup before the diploma ink got dry.



Time Travel

In her book Fei-Fei Li describes reviving the dormant Stanford AI Lab in the Gates Building on the university’s well-manicured campus. But as I described almost 40 years ago in my book *Hackers*, the original SAIL was set apart—in more ways than one. Note the early description of the internet at the end of this passage.

[SAIL’s setting was] a semicircular concrete, glass, and redwood former conference center in the hills overlooking the Stanford campus. Inside the building, hackers would work at any of 64 terminals scattered around the various offices. Instead of the battle-strewn imagery of shoot’em up space science fiction that pervaded [MIT’s] Tech Square, the Stanford imagery was the gentle lore of elves, hobbits, and wizards described in J.R.R. Tolkien’s Middle Earth trilogy. Rooms in the AI lab were named after Middle Earth locations and the SAIL printer was rigged so it could handle three different Elven type fonts... It did not take long for SAIL hackers to notice that the crawl space between the low-hanging ceiling and the room could be a comfortable sleeping hutch and several of them actually lived there for years. One systems hacker spent the early 1970s living in his dysfunctional car parked in the lot outside the building—once a week he’d bicycle down to Palo Alto for provisions. The other alternative for food was the Prancing Pony, the SAIL food-vending machine, loaded with health-food goodies and potstickers from a local Chinese restaurant. Each hacker kept an account on the Prancing Pony, maintained by the computer.

Stanford and other labs, whether in universities like Carnegie-Mellon or research centers like Stanford Research Institute, became closer to each other when ARPA linked their computer systems through a communications network. This “ARPAnet” was very much influenced by The Hacker Ethic in that among its values was the belief that systems should be decentralized, encourage exploration, and urge a free flow of information. From a computer at any “node” on the ARPAnet, you could work as if you were sitting at a terminal of a distant computer system. People sent a tremendous amount of electronic mail to each other, swapped technical esoterica, collaborated on projects, played Adventure, formed close hacker friendships with people they hadn’t met in person, and kept in contact with friends at places they’d previously hacked.



Ask Me One Thing

Liene asks, “Can great ideas come from great altered minds? Shouldn't smart people alter their minds a little more these days?”

Hi, Liene. I’m assuming you’re speaking of psychedelics, which are very much in vogue. And certainly they have had their impact on some of tech’s best talent. On a recent Joe Rogan podcast, Sam Altman, spurred by the host’s enthusiasm, extolled the virtues of psychedelic therapy. And Steve Jobs told journalist John Markoff that taking LSD “was one of the two or three most important things he had done in his life.” Think of that when you’re picking up your iPhone 58 times a day.

But it isn’t only chemicals that bend minds. As I explain in the essay above, Fei-Fei Li’s mind was altered by seeing the way that neural nets viewed the world. And she didn’t need to visit a dispensary or dealer! For my money, though, the most mind-altering substances are stocked on the shelves of bookstores and libraries. Between the covers of those tomes are ideas that can level up even the most elevated minds. And I question the intelligence of anyone who doesn’t read. Case in point: crypto-fraud Samuel Bankman-Fried, who said that no book is worth reading, and “If you wrote a book, you fucked up, and it should have been a six-paragraph blog post.” Maybe Sam will see the error of his ways, and alter his mind in the prison library.

You can submit questions to mail@wired.com. Write **ASK LEVY** in the subject line.



End Times Chronicle

Vampire bats are headed to the USA. Worst case scenario: rabies cases and more Twilight sequels.



Last but Not Least

My exclusive sneak peek at TGL, the sports league reinventing golf as a high-tech, made-for-TV stadium competition. Tiger's involved!


How surveillance and cell phone video clips have become San Francisco's civic language.

Want to liberate a pig? Check out the ops manual for the more disruptive radical animal freedom group.

Pin this AI gadget to your chest to (maybe) break your smartphone addiction.



You Might Also Like ...

-  Find the best bargains on quality gear with our Deals newsletter
- "Someone is using photos of me to talk to men"
- First-gen social media users have nowhere to go
- The truth behind the biggest (and dumbest) battery myths

- We asked a Savile Row tailor to test all the “best” T-shirts you see in social media ads
- My kid wants to be an influencer. Is that bad?
- 🌞 See if you take a shine to our picks for the best sunglasses and sun protection




Steven Levy covers the gamut of tech subjects for WIRED, in print and online, and has been contributing to the magazine since its inception. His weekly column, **Plaintext**, is exclusive to subscribers online but the newsletter version is open to all—**sign up here**. He has been writing about... **Read more**

EDITOR AT LARGE 

TOPICS PLAINTTEXT ARTIFICIAL INTELLIGENCE MACHINE LEARNING COGNITION GOOGLE
DEEP LEARNING ETHICS BOOKS

MORE FROM WIRED



Generative AI Is Coming for Sales Execs' Jobs—and They're Celebrating

ChatGPT-style AI can tackle the drudge work of responding to RFPs faster than humans. Sales teams at Google, Twilio, and others say productivity is spiking.


PARESH DAVE



What OpenAI Really Wants

The young company sent shock waves around the world when it released ChatGPT. But that was just the start. The ultimate goal: Change everything. Yes. *Everything*.

STEVEN LEVY



The AI Detection Arms Race Is On—and College Students Are Building the Weapons

Gen Alpha is quickly developing tools that identify AI-generated text—and tools to evade detection.

CHRISTOPHER BEAM



ChatGPT Can Now Talk to You—and Look Into Your Life

ChatGPT inches closer to feature parity with the seductive AI assistant from *Her*, thanks to an upgrade that adds voice and image recognition to the chatbot.

LAUREN GOODE

Chatbot Hallucinations Are Poisoning Web Search

Untruths spouted by chatbots ended up on the web—and Microsoft's Bing search engine served them up as facts. Generative AI could make search harder to trust.

WILL KNIGHT

The Fall of Babylon Is a Warning for AI Unicorns

When it went public in 2021, Babylon Health was valued at over \$4 billion. Now it has declared bankruptcy. Insiders say it could never live up to its hype.

GRACE BROWNE

Google's AI Is Making Traffic Lights More Efficient and Less Annoying

Google is analyzing data from its Maps app to suggest how cities can adjust traffic light timing to cut wait times and emissions. The company says it's already cutting stops for millions of drivers.

PARESH DAVE

How Neuralink Keeps Dead Monkey Photos Secret

Elon Musk's brain-chip startup conducted years of tests at UC Davis, a public university. A WIRED investigation reveals how Neuralink and the university keep the grisly images of test subjects hidden.

DELL CAMERON