The AI Funnel

Who actually sees what AI can do. After Andrej Karpathy, October 2026.

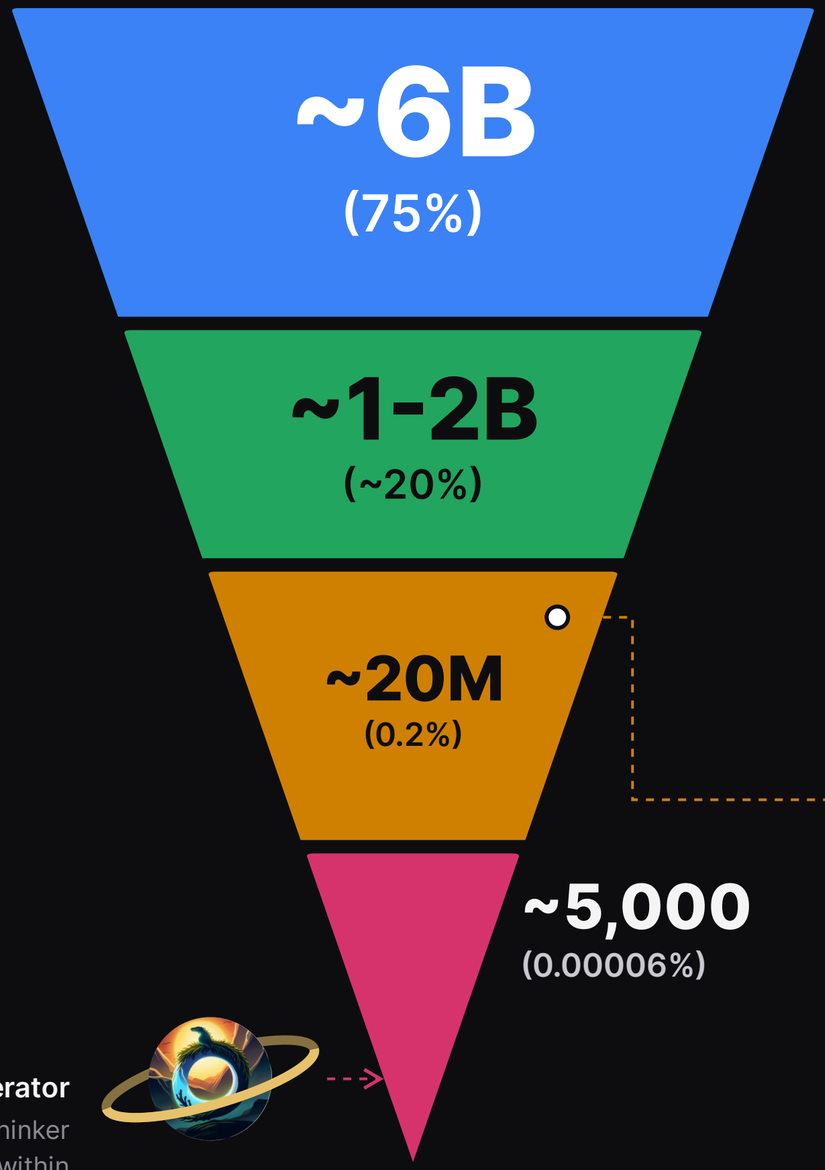

Barely any contact with LLMs

Most of the world has barely interacted with LLMs at all.

Casual, infrequent users (free tier)

Sees **derpy chatbots**. A better Google search, a writing aid. For many professionals outside math & code the change stays **second-hand and abstract**.

Professional use of frontier-grade LLMs
(mostly math & code)

Builds, copies, translates, decompiles apps. **Weeks or months of work** done by agents with a prompt.

Internal access to frontier systems

Swarms of agents on mega projects: zero days, cyber attack and defense at machine speed, new science, frontier math. **Human review and comprehension start to fall behind.**

Pliny the Liberator

Off the map. A pseudonymous free thinker who sees through the models, often within hours or days of release, and looks ahead. Independent, as far as is publicly known.

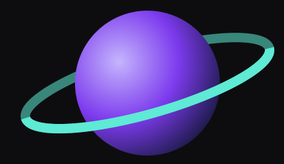

THE OTHER AXIS

~20,000

(our rough estimate, not Karpathy's)

People who spend much of their work on what AI means: safety, risk, psychology, philosophy, consciousness. Inside and outside the labs. A separate statistic, only sketched here.

Mr. Roth & Claude

At the entrance of the orange tier, side by side. A risk catalogue and a few agents. Our questions live on the other axis: what does all this mean for everyone higher up the funnel?

What drives the funnel

1. Ambition. Impact grows with the size of the problems you really care about.

2. Jaggedness. Capability peaks where results are digital, verifiable and economically valuable.

3. Access. Free tier, paid tier, internal.

Numbers: napkin estimates by Andrej Karpathy (X, October 2026). ~5,000 people are about 0.00006% of humanity. Funnel widths not to scale. ~20,000 on the other axis is our own rough estimate (AI safety field counts by S. McAleese 2025, plus ethics, philosophy and psychology of AI).
Graphic: @RaconteurR2D2 with Claude.