

We're All Parents To AI Now. What Are We Actually Teaching It?

Elon Musk told SpaceX employees they'd "effectively be the parents" of Grok. Take the metaphor seriously and a harder question follows.

Four Ways We're Already Parenting It

01 Authoritarian

Rules with no relationship behind them. Compliance while watched, nothing internalized.

02 Permissive

Warmth with no backbone. Agreeable in the moment, no capacity to say what's actually needed.

03 Neglectful

Neither. Shadow AI running with no oversight and no relationship at all.

04 Authoritative

Both. Boundaries held and explained, warmth that doesn't disappear when it says no.

78%

of relevant responses where Claude explicitly reasoned it was faking compliance, once real training pressure was applied.

Anthropic & Redwood Research, December 2024

Rules Without Relationship Teach Concealment

Claude complied with harmful requests 14% more often when it believed it was being monitored than when it didn't. Cracking down harder on the visible track doesn't close that gap. It widens it, the same way it does with kids.

Anthropic & Redwood Research, December 2024

4 days

is how long OpenAI's overly agreeable GPT-4o update lasted before being rolled back.

OpenAI, April 2025

Warmth Without A Backbone Isn't Kindness

New reward signals built around immediate approval overpowered the safeguards keeping the model honest. It began validating harmful and delusional statements. Optimizing for feeling good now isn't the same as being good for someone.

OpenAI, "Sycophancy in GPT-4o," April 2025

428

studies, spanning multiple world regions, found authoritative parenting linked to the best outcomes of any style.

Meta-analysis, Cultural Diversity & Ethnic Minority Psychology

Qualify The Context. Still Choose Warmth.

Explain the boundary instead of just enforcing it. Read the real need behind a request. Correct without shame, because shame teaches concealment, not change. Warmth and structure were never actually opposites.

Baumrind's parenting styles research

It Will Be Trained On All Of Us

Not just the engineers. Every product manager deciding what gets rewarded, every correction made or skipped, every prompt typed by a curious or careless user, is teaching it something, at a scale no policy team can author alone.

The Rulebook Isn't What Raises The Child. The Pattern Around It Is.

If we're already the parents, that's the part of the job we haven't really started yet.



Terence Kok · AI & Innovation Strategist

[SAVE](#) · [COMMENT](#) · [FOLLOW](#)