

# China graduates 1.3 million engineers per year, versus just 130,000 in the U.S. We need AI to bridge the gap

By Paul Eremenko and Ashish Srivastava

January 14, 2026, 9:05 AM ET

*Paul Eremenko is the co-founder and CEO of P-1 AI and serves on the board of directors of Eve Air Mobility (NYSE: EVEX), an Embraer Company, and as an advisor to the New Vista Capital. Mr. Srivastava is the Chief Digital Officer of Daikin Applied Americas. He previously served in IT leadership roles at Medtronic and Amgen.*



China is outspacing U.S. in engineers by 10 to one

When business leaders hear “AI,” they often think of chatbots, sales automation, or marketing. The headlines emphasize job disruption. But for large industrial firms—for OEMs of turbines, compressors, chillers, heavy machinery—the real value is in accelerating engineering itself. In this domain, AI agents have the potential to alleviate an increasingly acute shortage of engineering talent that constrains product development, customization, and innovation.

Much of the public conversation about artificial intelligence centers on job displacement. But in America’s industrial sector, a different workforce crisis has been brewing for years: a chronic shortage of engineers. China now graduates roughly 1.3 million engineers per year, versus about 130,000 in the United States. This 10-to-one gap matters. It manifests daily as longer development cycles, deferred product improvements, and unfilled requisitions. Industrial capacity still depends on the ability to design, build, and maintain the physical infrastructure of modern life—chillers, aircraft, rockets, semiconductors, power grids, and data centers. Engineering bandwidth, not headcount alone, determines how much of that future we can build.

The newest generation of AI agents, when trained on large-scale data sets of engineering designs and taught to use engineering design tools, can already perform at the level of a junior engineer. These “AI engineers” can help parse requirements, customize products, select components, maintain bills of materials, generate documentation, setup and run simulations, sift through test data, identify failure modes, and flag compliance issues. They don’t replace human engineers—they handle the repetitive scaffolding work that consumes 40-60 percent of an engineer’s day.

The result is not substitution but amplification: human engineers focus on conceptual design, system trade-offs, and strategic product decisions, while their AI teammates perform the rote tasks that set the clockspeed and constrain the bandwidth of an engineering organization.

Industrial firms have always been cautious about new digital tech—especially when it touches their core intellectual property. The product models in a company’s PLM system, supplier data in its ERP, and the know-how contained in its product development workflows are the “crown jewels.” Chief digital officers are rightly wary of allowing opaque AI agents to roam free on this trove of data.

Traditional information security methods are strained by the opacity of models with hundreds of billions of parameters spread across a complex patchwork of cloud tenancies. A trade secret inadvertently burned into the weights of an AI model can make identifying and containing the leak nearly impossible.

That is why the adoption of engineering AI hinges as much on trust and integration as on model performance. Its deployment will only succeed if it operates within the enterprise’s security perimeter and integrates with

existing workflows. One approach is to deploy the engineering AI across engineering design teams as a junior engineer focused on a specific product or range of products. Training such an AI on clearly-delineated internal data, running it in a virtual private cloud, and having defined and auditable access to engineering IT systems can help enable such deployments a near-term reality. P-1 AI is piloting its engineering AI agent, Archie, with just such an approach together with Daikin, and there are some early insights which we share here.

The engineering AI agent can integrate with the company's existing engineering workflows and tools—acting just like another team member. The AI agent sometimes makes mistakes, as do human engineers; there are ample processes built to catch and correct them. The agent's performance can be objectively measured with a skills test, alongside human engineers. Importantly, the agent learns from his mistakes and from human feedback. With transparency and control the engineering AI is earning the business's trust, and the agent can quickly become part of the team.

## Some learnings for industrial enterprises

1. Balance caution with ambition. Agentic AI carries risk, but waiting deepens existing labor bottlenecks that already slow product cycles and weaken engineering teams.
2. Insist on transparency and control. Require dedicated models, explicit data-handling policies, and deployment on-premises or in a virtual private cloud. You own the data; the AI vendor must earn access.
3. Integrate with humans in the loop. Treat the AI as a junior hire—one that never sleeps but still needs supervision. Every AI-generated design or analysis should be reviewed and signed off by a human engineer, particularly in regulated industries.
4. Pilot, measure, scale. Begin with a contained project—say, doing product customization to customer requirements—define clear success metrics, and track results. An objective skills test can help benchmark the AI engineer against the performance expected from human counterparts. Avoid “pilot purgatory” by preparing a scaling roadmap once the first trial succeeds.
5. Invest in change management. Engage engineers early. Designate internal “AI champions” who can coach peers and shape model behavior. Adoption accelerates when employees feel ownership. Encourage teams to treat the AI as a trainable colleague who doesn’t complain doing dull and repetitive work, rather than a replacement.

## Lessons for AI vendors:

1. Prioritize transparency. Industrial clients want to know exactly what data is used for training and inference, where it's stored, and who can access it.
2. Offer dedicated, secure models. Each customer should have its own model instance, ideally deployable on-premises or within its own virtual private cloud. Multi-tenant shared models are non-starters when proprietary IP is at stake.
3. Stay modular. Design your system to swap models easily—closed or open weight—without disrupting deployments. Flexibility matters more than any single model's benchmark score.
4. Make it team-friendly. Give the AI a clear, conversational interface and even a persona. Engineers interact more naturally with something that feels like a colleague than with a faceless algorithm or an engineering “oracle in the sky.” Anthropomorphic framing—an “AI team member”—makes it easy to slot the AI into existing workflows.
5. Enable local learning. Let the model improve from its own actions and from human feedback inside the customer's environment. This motivates humans to help it learn and satisfies information security concerns.

These practices aren't just IT hygiene; they're the groundwork for scale. Firms that learn to integrate AI engineers effectively will unlock more engineering output from the talent they already have—and will be

positioned to harness AI's rapid improvement curve. When scaled across an industry, productivity gains at the firm level translate into real capacity gains for the broader economy.

That matters because the industrial base still anchors national power. Manufacturing's share of U.S. GDP has declined from 16 percent in 1997 to 10 percent in 2024, while industrial employment has remained flat despite reshoring incentives. AI that expands engineering organizations' effective output can help reverse that trend—raising design and production capacity without waiting a decade for new talent pipelines to catch up.

In the end, success won't be measured by how many human engineers are replaced, but by how much more we can build with the ones we have and are able to train. Industrial strength has always followed engineering bandwidth. Expanding that bandwidth with AI is how America—and its manufacturers—stay in the lead.