Invisible

TRENDS & INSIGHTS:

AI Training 2025

Matt Fitzpatrick
CEO, Invisible Technologies

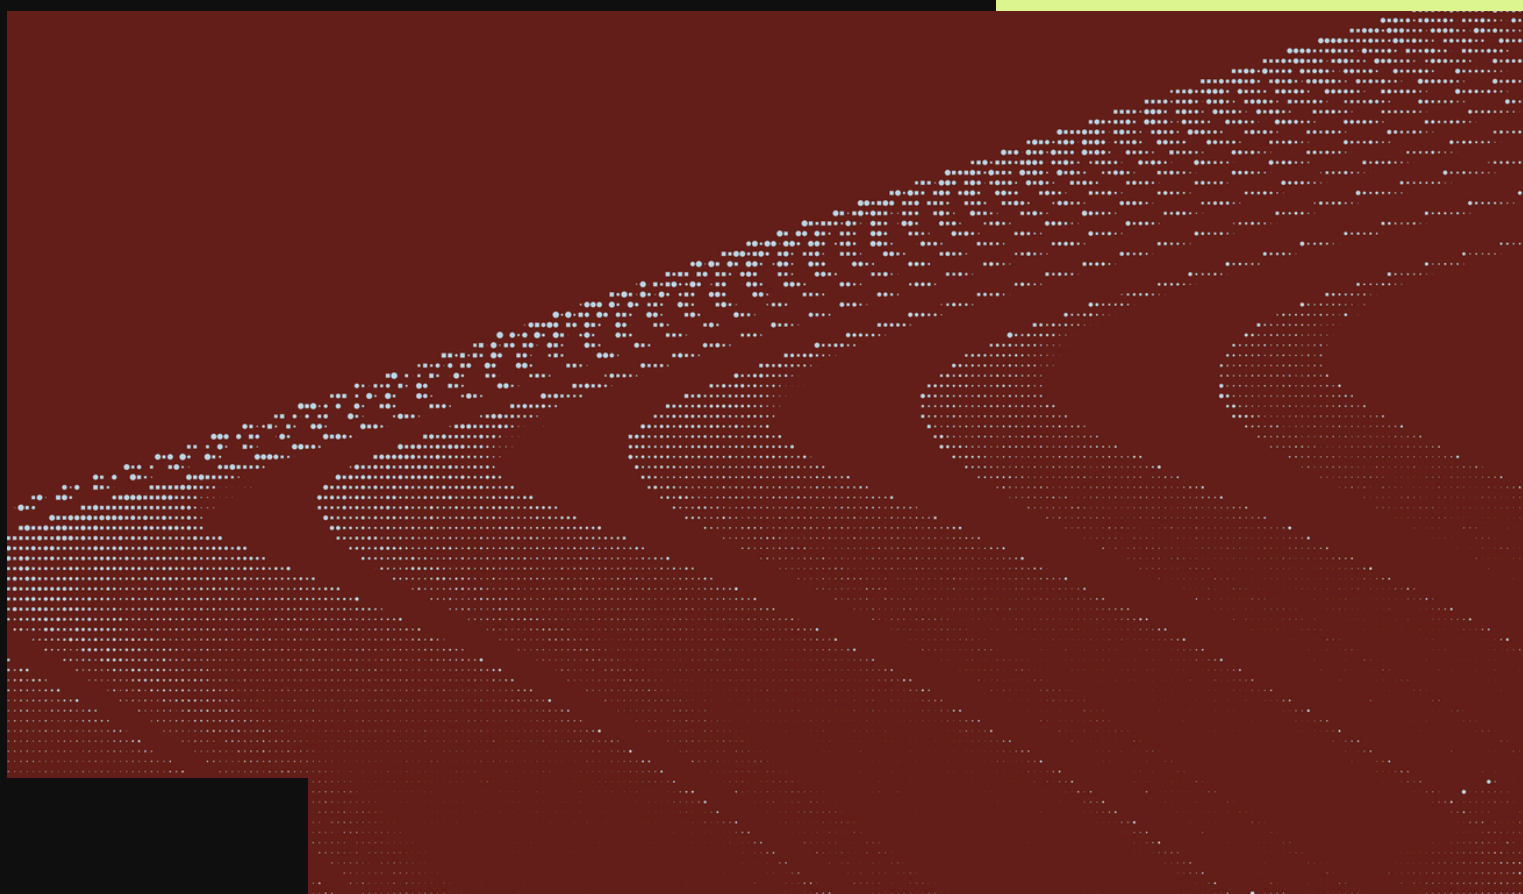

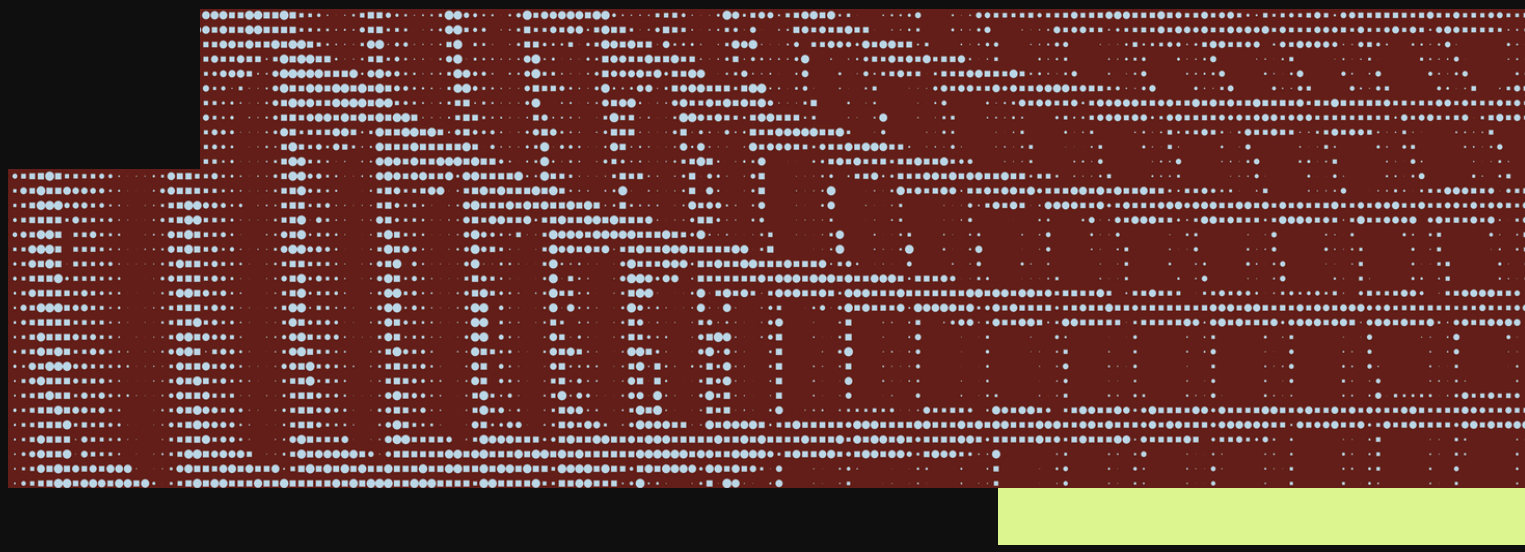

Invisible Technologies has trained nearly every major foundation model, including AWS, Microsoft, and Cohere. And our technology adapts to any sector, function, or system to make AI work in the real world.

As we head into H2, we're sharing the key trends shaping the future of AI training. **H1 marked a clear inflection point:** Chain-of-Thought and agentic reasoning moved from theory to deployment, MCP launched and scaled rapidly, and efficiency became the new north star for model output.

From our vantage point, we're seeing critical shifts in how models are built, trained, and brought to life.

Here's what's changing and where we're noticing shifting trends in the market right now:

TREND 01

Raising the bar for quality pre-training data

Firstly, the bar for quality pre-training data is rising.

In previous years, chat data full of typos and stereotyping was tolerated, because it helped the model understand how to interact generally. As the models become more sophisticated, each incremental improvement demands more refined datasets, with the complexity of optimization only increasing. Now things like accuracy and completion matter in every single response.

This is the core challenge of pushing model performance further — finding those last bits of high-fidelity data that meaningfully inform the next round of fine-tuning.

TREND 02

Shorter, more frequent training sprints

A second trend we are seeing is an increased demand for shorter, more frequent training sprints.

Previously, models were trained on hundreds of thousands, or millions, of rows, and a training sprint would take months. A major shift in the industry is the move from long training cycles to shorter, more iterative processes. Many foundation models have also learned to separate experimentation from large scale training, allowing them to pinpoint and solve specific issues more effectively.

The focus has shifted from quantity to quality-first data strategies, ensuring that better curated datasets lead to more efficient and capable models.

TREND 03

More specialized & diverse data

Additionally, we are noticing an increasing need for more specialized and diversified data.

Now that readily available data (i.e. scraping the internet) has largely been exhausted, model builders are focusing on capturing deeper nuance and specificity to push AI performance further. Rather than training on massive general datasets, researchers are looking at the long tail of human experience — fine-tuning models with highly specialized and diverse data. For example, two years ago, broad requests for 100,000 rows of general 'chit-chat' data were common. But as models have become proficient in everyday conversation, the focus has shifted. Now, we see requests for dozens of smaller datasets covering distinct topics like fly fishing, recovering from niche foot surgery, or discussing 13th century Japanese architecture.

This targeted approach helps models generate more context-rich, domain-specific interactions, allowing them to push past previous performance plateaus.

TREND 04

Chain-of-thought reasoning

A key evolution of chain-of-thought (CoT) reasoning has been increased visibility into how models reason step by step.

As enterprise and consumer AI roadmaps get more sophisticated, this transparency is crucial. With a clearer view of where CoT succeeds or stumbles, we can train models to reason more concisely and efficiently — especially important as we move from basic instruction-following to solving high-stakes, multi-layered challenges. It also boosts performance without needing more data, works well with prompt engineering, and lays the groundwork for advanced methods like Self-Consistency (where multiple reasoning paths are explored) and Tree of Thought (where models evaluate multiple possible reasoning paths before settling on an answer).

TREND 05

Agentic products

Lastly, we are seeing agentic products emerging as a major trend in AI, shifting the focus from simple one-shot responses to multi-step, orchestrated workflows where models can call other models and tools to complete complex tasks.

This approach enables AI to act more like an assistant that executes real-world tasks — such as generating spreadsheets, handling legal workflows, or recommending financial trades — rather than just responding to queries. Building these systems requires new training strategies, evaluation methods, and a focus on trustworthiness, including citation transparency. Crucially, the effectiveness of agentic systems depends on optimizing for specific outcomes using high-quality data, reinforcement learning, and fine-tuning. As demand grows, especially across enterprise and service sectors, agentic AI is proving to be more scalable and capable than traditional generative models.

Invisible

If you'd like to hear more about the trends we're seeing in this space, don't hesitate to reach out.

Let's connect and explore these insights with you further.

Matt Fitzpatrick
CEO, Invisible Technologies

BOOK A DEMO

About us

Invisible Technologies is the AI operating system for the enterprise. Our end-to-end AI Process Platform structures messy data, builds digital workflows, deploys agentic solutions, evaluates/measures impact, and mobilizes relevant human experts. Invisible has trained foundation models for more than 80% of the world's leading AI model providers, including Cohere, Microsoft, and AWS, and we have the expertise to customize AI for any industry, function, or use case.

Invisible makes AI work in the real world. In 2024, we reached $134M in revenue and were named the #2 fastest growing AI company on the Inc. 5000.