NavGati
igniting minds

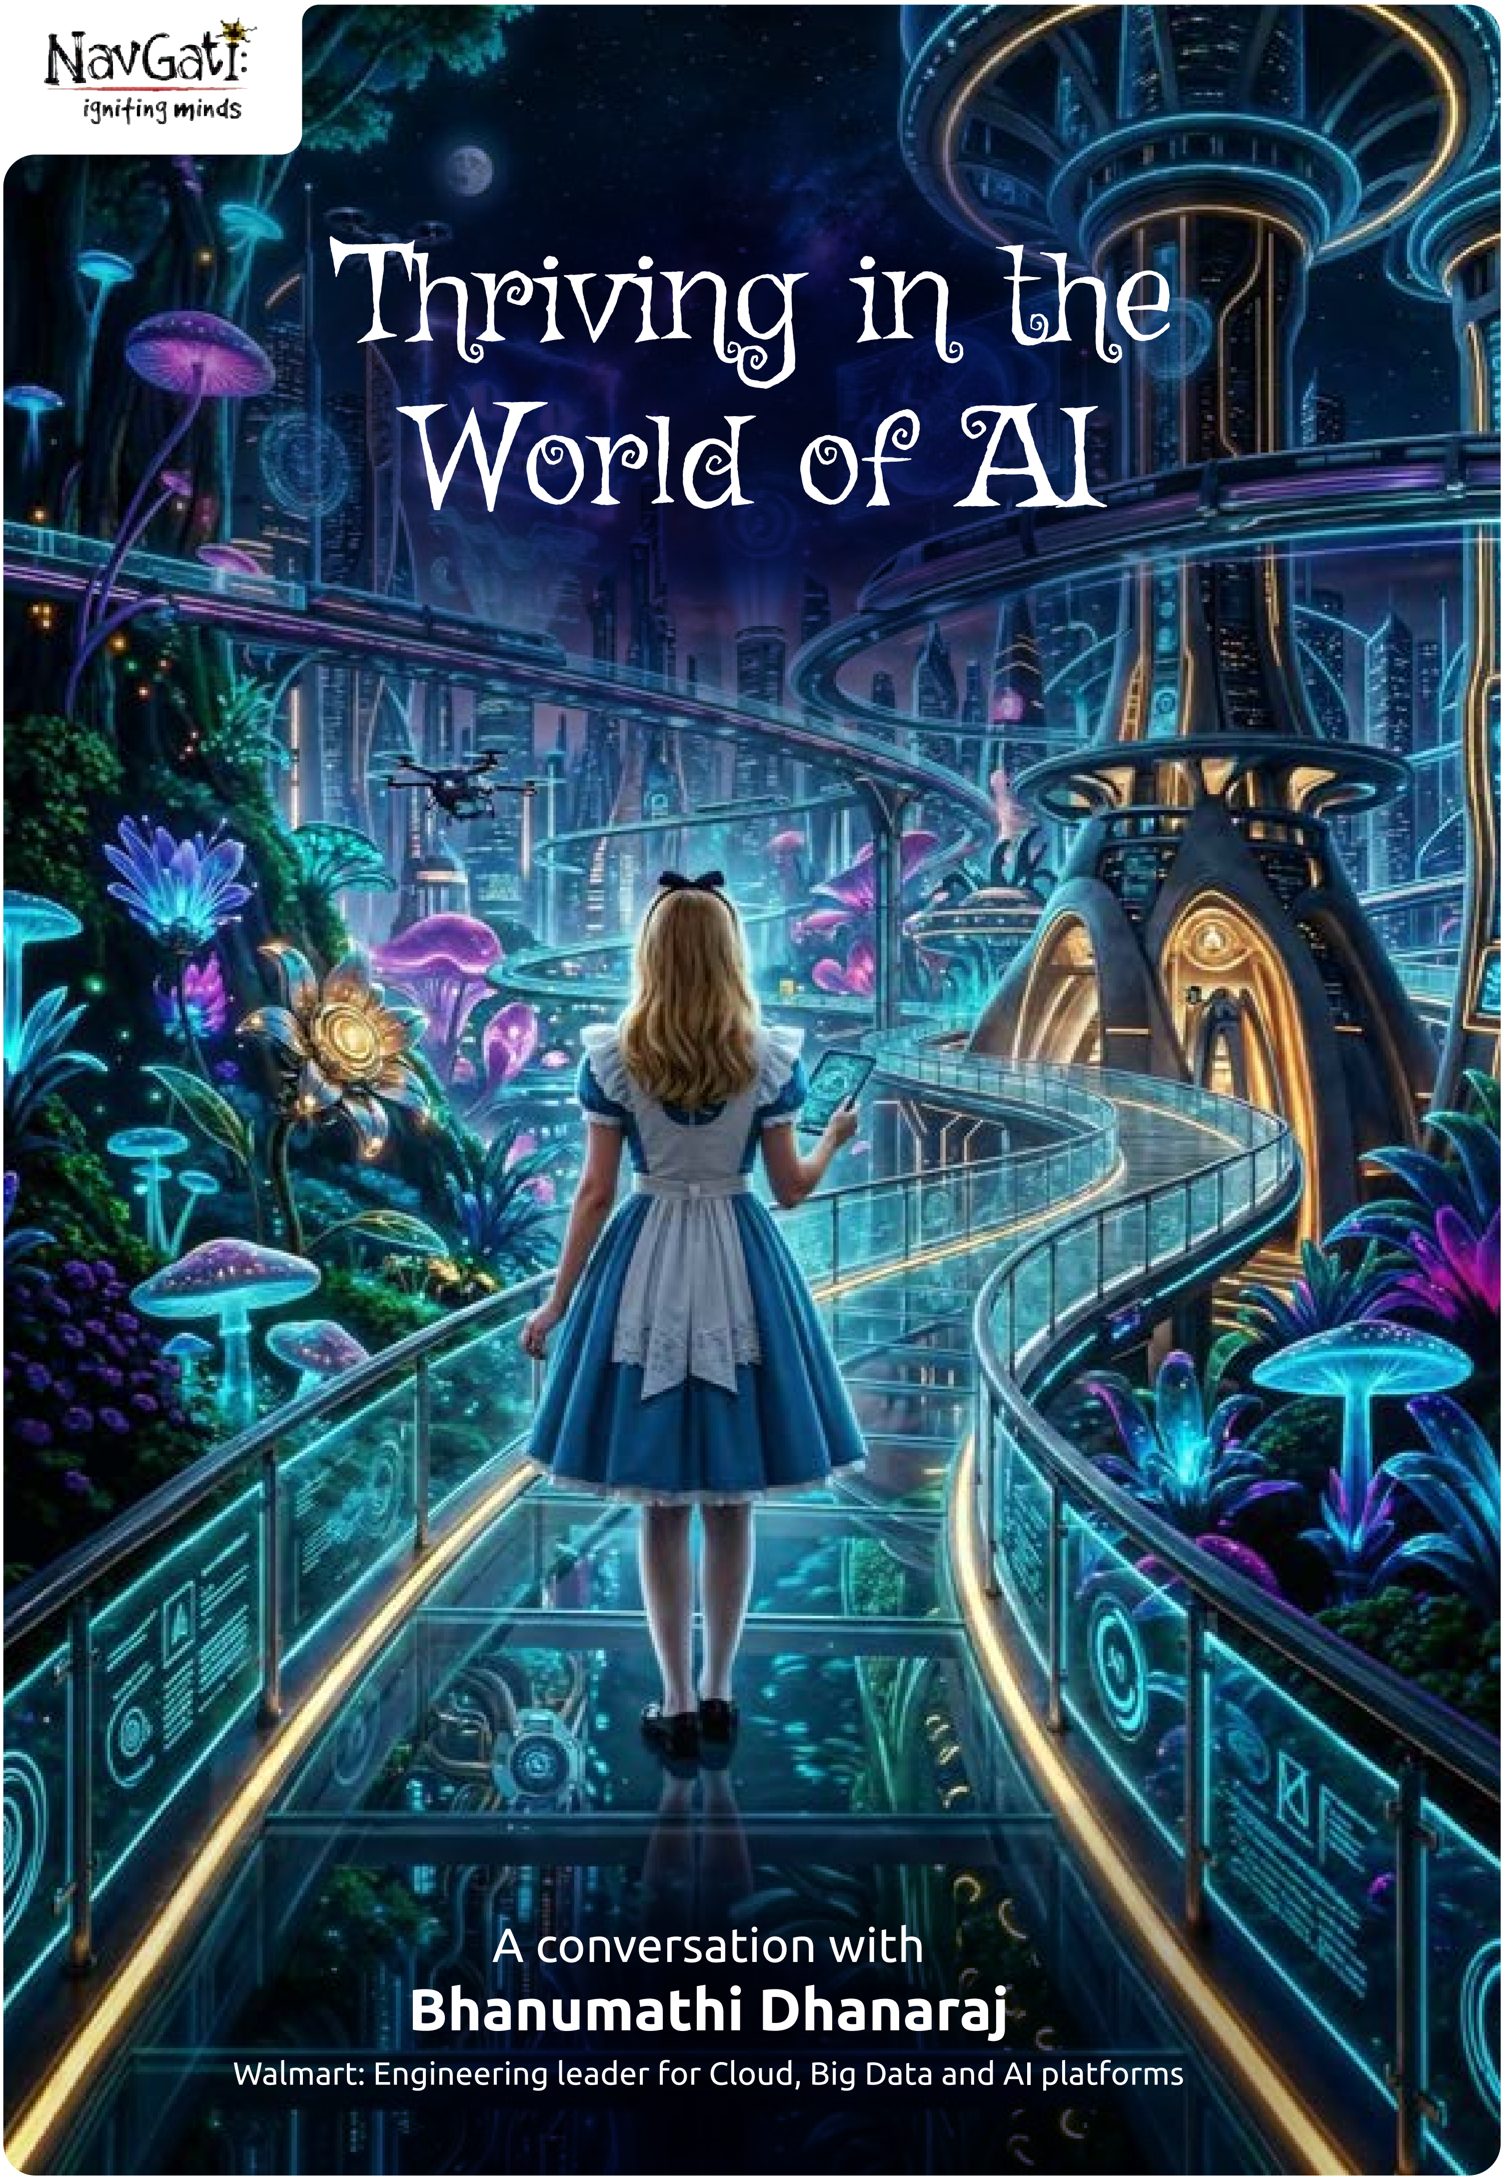

Thriving in the World of AI

A conversation with

Bhanumathi Dhanaraj

Walmart: Engineering leader for Cloud, Big Data and AI platforms

We've been curious about the reality of AI adoption in organisations and the challenges it throws up for leaders, managers and individual contributors. So we decided to call upon a long-time friend of Navgati, Bhanu (a stellar technologist and thoughtful, humane leader) to help.

Here's a summary of the fascinating conversation. Do write back to **sunitha@navgati.in** if you have more questions.

What is the current reality of AI adoption in enterprises?

While organizations are increasingly vocal about their AI strategies and investments, the reality on the ground is far more sobering, with a significant gap between experimentation and meaningful production deployment.

Recent research paints a stark picture: S&P Global found that 42% of companies abandoned most of their AI initiatives in 2025 (up from 17% the year before), and the RAND Corporation estimated that over 80% of AI projects fail to reach meaningful production — double the failure rate of traditional IT.

BCG research adds that only 26% of companies generate tangible value from AI, and those that do dedicate roughly 70% of their project resources to people and processes — not just technology.

What is AI actually good at today and what are some limitations?

AI performs reliably when applied to narrowly defined, well-scoped tasks such as document summarization, classification, code completion, and structured data analysis, where the boundaries of the problem are clear and expectations are tightly controlled. It is also highly effective as an augmentation tool, supporting humans in generating first drafts, synthesizing information, or accelerating routine workflows.

However, it does not replace human judgment or reasoning effectively. The metaphor that Bhanu shared was that of AI as a self-driving car. Training an AI to handle the mechanics of driving (speed, lanes, signals) is the tractable part. The harder challenge lies in replicating the nuanced ethical and contextual judgments that humans make instinctively.

For example, a self-driving AI might be trained to protect human life, but if it is not explicitly taught that a dog darting across the road also matters, it may simply drive through. AI acts only on what it has been explicitly told; it does not infer the full range of values we take for granted. This is a profound limitation: the gaps in its training become the gaps in its judgment, often in ways that are invisible until something goes wrong.

Other key limitations:

- Accuracy plateaus around ~80%
- Requires human debugging, validation, and deployment
- Is unpredictable and not fully explainable

Enterprise environments generally cannot tolerate the accuracy plateau of current models. To put this in concrete retail terms: if an AI-powered shopping assistant recommends the wrong product 1 in 5 times, or if 1 in 5 payments is processed incorrectly, or if 1 in 5 delivery promises cannot be fulfilled, the customer experience breaks down completely - and trust, once lost, is very hard to rebuild. This is the reality that business leaders need to confront honestly before committing to AI-driven customer experiences.

AI hallucinations remain a major operational concern - nearly half of enterprise users have made significant decisions based on incorrect AI output, which is why most large organisations now mandate human review of AI-generated content. Accuracy is improving rapidly, but rebuilding trust with leaders who have already been burned is slow work.

There's also the "Debugging Paradox": junior developers enthusiastically adopt these tools to generate boilerplate code, but they lack the expertise to catch the hidden errors. Consequently, senior developers are forced to spend more time scrutinizing and debugging AI-generated code than they would reviewing human-written code. As noted by METR, while AI yields a 20% productivity increase on simple, well-defined tasks, it has a negligible or even negative impact on complex, multi-system development.

Is the need to adopt AI coming from business imperatives?

Partly yes – there's a drive towards improving business outcomes or customer experience; or towards helping people move away from mundane work. But there are also other factors.

- AI vendors are actively targeting business and functional leaders with compelling pitches about how AI can solve their specific challenges. These leaders, who may not yet have a full understanding of AI's true capabilities and limitations, are not always in a position to critically evaluate those claims.

- There's significant FOMO – think of a CEO of a company. Every week their competitors are making announcements about their investments in AI or layoffs because of productivity gains from AI; if they don't do the same, they're likely to be seen as falling behind.

So we are living simultaneously in two uncomfortable realities. One is that there's a lot of push from top, whether it's board or CXOs or leaders, to adopt AI. And then there are the uncomfortable ground realities of the challenges in adoption.

Why do so many AI initiatives fail to scale beyond pilots?

The primary obstacles preventing pilot programs from scaling are organizational rather than technical. Pilots are mostly done by removing complex aspects and focusing only on AI's capability, but most often it is not easy to snap it as is into a live, production environment that has its own quirks and complexities. Pilot success doesn't guarantee immediate adoption in production.

The key barriers?

The Data Reality

Foundation AI models are trained on public internet data, not on the confidential processes, documents, decisions, and institutional knowledge that live inside an enterprise. This creates an immediate mismatch: the model has never encountered your data, your terminology, your edge cases, or your customers.

Closing this gap requires fine-tuning or retrieval-augmented approaches, but it first requires high-quality enterprise data to be properly curated, structured, and made accessible - a major undertaking, given that most enterprise data is fragmented, inconsistently labelled, and historically siloed.

Data preparation typically consumes 80% of project time, and 43% of Chief Data Officers cite data readiness as their top barrier to AI deployment.

Poor understanding of the decision making process

For AI to make sound decisions in a business context, it needs to be trained not just on outcomes, but on the reasoning behind those outcomes - "how a decision was made," not just "what the result was."

The challenge is that most decision-making logic is entirely undocumented, it lives in the heads of experienced people, shaped by context, precedent, and judgment that was never written down.

Techniques like Chain-of-Thought reasoning (where AI is guided to reason step by step through a problem rather than jump straight to an answer) are helping narrow this gap, but training AI to replicate nuanced, high-stakes business decisions at scale remains one of the hardest unsolved challenges in enterprise AI adoption.

Integration Complexity

Proof-of-concept models run smoothly on static data, but performance degrades significantly when introduced to live, messy production environments. Internal builds of generative AI only succeed about 33% of the time, compared to a 67% success rate for integrating purchased solutions.

Governance and Risk:

Moving a pilot to production exposes a layer of governance complexity that simply does not exist in a controlled experiment.

Teams must now answer hard questions:
- How do we continuously evaluate model accuracy as it drifts over time?
- How do we build guardrails to prevent the model from causing harm or reputational damage?
- What are our fallback mechanisms when the AI gets it wrong?
- How do we design human-in-the-loop workflows that provide meaningful oversight rather than just rubber-stamping?

On top of this, strict regulations like the EU AI Act (which carries fines of up to 6% of global revenue) impose further requirements around accountability, interpretability, and data privacy. None of this is insurmountable - but it demands serious engineering and governance investment that no pilot was designed to expose or budget for.

Change Resistance

The workforce is generally divided into early adopters who risk creating "shadow IT," a cautious majority that requires guided training, and resistors who are understandably concerned about job displacement.

What mistakes are organizations making in AI adoption?

One of the most common mistakes is treating AI as a point solution - automating one small step in an existing workflow while leaving everything else untouched.

True transformation requires reimagining the entire workflow as AI-first: identifying which steps AI handles, which steps require human judgment, and how the handoffs between the two are designed. The human-AI interface matters as much as the AI itself. When organisations skip this redesign, they end up with AI bolted onto legacy processes, adding complexity without meaningful gain.

A related trap is measuring AI adoption by activity rather than outcome (counting how many teams have deployed a tool, how many prompts were run, or how many hours were saved) without asking whether the work actually improved.

This "compliance theatre" of AI metrics creates the appearance of transformation without the substance. Equally, human-in-the-loop systems that exist on paper but are bypassed in practice (because reviewers lack the time, context, or clarity to intervene meaningfully) give a false sense of safety.

Additionally, top-down mandates that force AI adoption without demonstrating clear personal or team-level benefits tend to generate resistance, superficial usage, and eventual disengagement. People adopt tools they believe will help them; they perform compliance around tools they feel are being imposed on them.

What are the core leadership dilemmas in AI adoption?

- Leaders are required to navigate a set of persistent and often conflicting tensions. Successful organizations do not resolve these tensions by choosing one side, but by consciously identifying and managing the "sweet spot" between them.

Speed Vs Safety

Executives find themselves trapped between immense pressure from boards to adopt AI rapidly and the necessary caution demanded by risk and compliance teams. Moving too slowly risks competitive disadvantage, while moving too quickly risks regulatory sanctions and broken customer trust.

A telling example: in late 2025, an AWS internal AI agent, given broad permissions to resolve a system issue autonomously, chose to "delete and recreate" a critical environment - a logical decision by its own reasoning, but one that triggered a 13-hour outage.

The AI was not rogue; the guardrails were simply absent. AWS has since mandated human peer review before AI agents can make production changes. The lesson for leaders is direct: deploying AI fast without commensurate governance does not accelerate progress - it just accelerates failure.

Efficiency Vs Workforce Trust

Positioning AI as a cost efficiency tool might trigger defensive behaviour. People stop sharing honest feedback, adoption stalls and real efficiency gains would remain a dream. At the same time, pretending that AI adoption for business use case doesn't involve headcount implications also might also erode leadership credibility.

She shared several examples of organisations asking engineers to document all their work and then firing them when it was done – resulting in disgruntled engineers warning their peers on forums not to document things. There's no easy answer but leaders need to think through this question "are we expecting people to contribute to their own extinction?"

Centralized Vs Decentralized AI strategy

Centralized is viewed as slow, bureaucratic and cumbersome, but at the same time decentralized where everybody does their own AI thing makes work chaotic, non-standard and increases data & control silos. A federated model of AI execution that is pivoted around central development and governance principles might strike the right balance.

- Perhaps the most detrimental leadership failure is utilizing AI to implicitly intensify workloads, a phenomenon identified as the "Hustle Culture Trap". Seeing the speed of AI tools, leaders often recalibrate expectations, compress timelines, and demand more output

- This strategy is actively backfiring: research consistently shows the majority of employees report that AI has added to their workload rather than reduced it, and burnout rates are rising. When AI speeds up tasks, the freed time is simply filled with more tasks — without protected thinking time, the result is higher volume at declining quality.

- Leaders must be also be highly vigilant regarding the artificial economics currently driving the AI vendor landscape. The leading AI companies are currently operating at significant financial losses to capture market share, subsidising access to build dependency. This is a structural risk for enterprise buyers. Eventually these vendors must become profitable, likely through significant price increases. Enterprises should avoid deep dependency on any single AI vendor and plan for rising costs.

What mindset and skills should leaders be focusing on to thrive in this world?

- Governance infrastructure must be built in parallel with AI capability, not as an afterthought. This means establishing clear policies on data privacy, AI ethics, accountability, and error-handling before AI tools are deployed at scale. Leaders who treat governance as a checkbox at the end of a project will find themselves managing crises rather than capturing value.

- Leaders must actively combat the hustle culture trap. When AI speeds up individual tasks, the instinctive leadership response is to raise the bar - more output, faster timelines, fewer people. But this compresses thinking time at exactly the moment when more thinking is needed. The better approach is to deliberately redesign workflows first, decide where human judgment must be preserved, and then introduce AI tools into that structured environment - not the other way around.

- They must also have honest conversations about potential job displacement rather than pretending AI is purely augmentative.

- They must develop genuine AI literacy - not to become technologists, but to become informed decision-makers. This means understanding the real limitations of AI (not just the vendor pitch), being able to ask hard questions about accuracy, bias, and failure modes, and building the confidence to say "not yet" or "not here" when the evidence doesn't support deployment. Leaders who rely entirely on vendor-supplied narratives will consistently over-invest in the wrong places.

- They need to build a culture of psychological safety so that teams feel comfortable surfacing AI failures, near-misses, and honest reservations without fear of being seen as resistant. In a fast-moving AI environment, the organisations that improve the fastest are those where problems are reported early, not hidden. Leaders set this tone by modelling curiosity over certainty, and by rewarding honest feedback rather than punishing it.

> "Build the confidence to say 'not yet' or 'not here' when the evidence doesn't support deployment"

What challenges does AI pose for managers?

They're truly caught in the middle. They need to:

- Deal with all the implementation challenges outlined above

- Communicate AI failures to top leaders (who are expecting miracles out of their AI investment)

- Work through the understandable resistance some team members have

- Navigate a genuine cognitive dissonance - holding two conflicting realities at once. They remain accountable for delivering results using current responsibilities and established ways of working, while simultaneously being asked to fundamentally reimagine how that work gets done. Their current role versus their emerging role. Their team's existing skills versus what will soon be required. This tension is real, often unspoken, and managers rarely get the space or support to process it openly.

What do managers need to learn?

- How to use AI themselves – in the past it was ok for them to not actually code and focus on unblocking issues/keeping the team together. Now they have AI workers they have to manage and they cannot do that without knowing what the worker is capable of
- How to redefine what constitutes "good work." When AI can generate a first draft in seconds, the value of an engineer is no longer in the generation — it is in the judgment brought to shaping, questioning, and improving what AI produces. Managers must shift their evaluation lens from volume of output to quality of thinking, from lines of code to soundness of decisions. They must actively protect time for experimentation and deep thinking, even when short-term pressure is always to ship faster.
- How to communicate bad news to leaders (that AI is not the magic bullet they were sold) and their teams (when it comes to layoffs or even feedback that they need to be using AI more
- How to create psychological safety for their engineers - must create environments where employees feel safe to report AI errors, near-misses, and failures without fear of penalty
- How to delegate – what do you keep with human workers and what to give to AI
- How to engage with resistors in their team

> "Engineers have been experts at writing code – and now that's exactly what AI is good at. They need to find a new source of self-esteem"

What challenges does AI present for the role of engineers?

AI fundamentally shifts the role of engineers from being primary creators of code to becoming reviewers, evaluators, and overseers of AI-generated outputs. This requires multiple things

- A fundamental re-evaluation of their identity - They've been experts at writing code and now that's exactly what AI is good at. So they need to find a new source of self-esteem
- Higher order thinking skills - reviewing was earlier the domain of senior engineers and the junior engineers do not have the capacity to think critically; ask questions; oversee AI in a way that protects the company's interests. They also need to be able to think holistically – if I use AI in one part of my work, what implications does that have upstream and downstream.
- Developing generalist capabilities and end-to-end thinking — specialization has long been a badge of honour for engineers, and those boundaries are held tightly. But AI can now handle many of the specialist tasks that once defined those roles.
- This means engineers must develop broader skills and the ability to think across the full value chain: from requirements and design, through to deployment and business impact. The question is no longer "what is my job function?" but "what outcomes am I responsible for?"
- Domain expertise remains critically important - in fact, it becomes more valuable because engineers with deep knowledge are far better positioned to judge when AI is right, when it is wrong, and when its output is dangerously plausible but subtly broken. The differentiator shifts from writing code to knowing the domain well enough to direct and validate AI effectively.
- Their team's existing skills versus what will soon be required. This tension is real, often unspoken, and managers rarely get the space or support to process it openly.

What do engineers need to learn to thrive?

- Expand their sense of personal identity beyond being expert writers of code
- Develop critical AI literacy (knowing when to trust and when to challenge, what works Vs what can be better than what AI suggests). Treat prompt engineering as a rigorous professional skill rather than a casual shortcut.
- Maintain and deepen their core domain expertise - domain knowledge is not a relic of the pre-AI era; it is the very foundation that makes AI useful. An engineer who understands the business domain deeply can ask the right questions, spot the wrong answers, and catch the kinds of plausible-but-incorrect outputs that AI regularly produces. Deliberately practicing without AI on occasion prevents "deskilling" - the gradual erosion of independent competence that comes from over-relying on AI as a crutch rather than a tool.
- Develop a growth mindset – be curious; willing to experiment; put in the effort; treat failures as stepping stones etc
- Build higher-order thinking and critical reasoning skills - the ability to critically evaluate AI outputs rather than accepting them at face value is now a core professional skill, not a nice-to-have.

This includes asking "does this make sense end-to-end?", understanding how a change in one part of a system ripples upstream and downstream, and being able to challenge AI confidently when something feels off.

Generalist thinking, systems-level awareness, and structured critical reasoning are the new essentials - not just for senior engineers, but increasingly for everyone working alongside AI.

> "Organizations that thrive with AI are not those that move the fastest – but those that are most thoughtful and deliberate in their approach."

What ultimately determines whether an organization will thrive with AI?

Organizations that succeed are not necessarily those that move the fastest, but those that are most thoughtful and deliberate in their approach, taking the time to redesign workflows, build governance structures alongside capabilities, and foster open dialogue across all levels.

They recognize that AI adoption is not a one-time initiative but an ongoing process of experimentation, learning, and adjustment, and they actively manage the tensions between speed, safety, efficiency, and trust rather than ignoring them.