

HOW TO MOVE AI FROM PILOT TO PRODUCTION

BY MARTIN BITTNER



Generative AI is widely considered a decade-defining technology that is creating opportunities across all sectors and industries, for private and professional use. No technology has ever seen an adoption curve as rapid as AI tools—surpassing the internet and social media in terms of

time till mass adoption, according to a September 2024 “On the Economy” blog post from the Federal Reserve Bank of St. Louis (“The Rapid Adoption of Generative AI”). In the same vein, IDC’s FutureScape 2025 report states that investment in enterprise AI adoption will cross the \$150 billion level.



Implementing AI at scale is a transformative process, and it must be seen not as a sprint but a marathon that requires preparation, alignment, and commitment.

McKinsey experts say that about 80% of companies are already using AI in at least one business function ("The State of AI in 2025").

Based on these data, it might appear that AI, and especially enterprise AI, is an almost magical success story rewriting the playbook across all industries. LinkedIn co-founder Reid Hoffman, in AllBusiness.com's "15 Quotes on the Future of AI," said "AI is going to reshape every industry and every job."

However, there is a different side to this story: The MIT NANDA report published in summer 2025 ("The GenAI Divide: State of AI in Business 2025") sent shockwaves across boardrooms when it found that about 95% of AI pilot projects never make it to production and/or fail to generate a return on investment. This finding is not an isolated one but is well in line with conclusions from IDC's recent CIO Playbook ("88% of AI Pilots Fail to Reach Production—but That's Not All on IT"), which shows only about 12% of pilots/proof-of-concepts progressing to full deployment.

Two questions emerge. How do we reconcile the general perception of AI, and the news stories of rapid and successful AI adoption, with the MIT NANDA report? And what can leaders do to ensure that they are part of the 5% that see an actual ROI and enterprise adoption?

I have spent the past 10 years developing automation solutions for complex enterprise workflows and processes. As a two-time founder and CEO of a company that delivers AI workflow automation for the enterprise, I see firsthand every day what it takes to successfully automate the countless processes that drive biopharma, insurance, industrials, and so many other industries behind the scenes, and what leads to actual value creation.

The following learnings summarize some of the key insights I have collected and offer a perspective into how leaders can ensure that their choices will lead to actual AI deployment in production, at scale, and more important, to a clear return on investment.

BEWARE: AI IS NOT SOFTWARE

Artificial intelligence (and more specifically, in the context of this article, generative AI) is inherently different from software or code. AI operates as a probabilistic system that makes predictions based on likelihood and statistics, as opposed to the deterministic nature of software that

performs the exact same operation every time it is called upon. What that means is that adopting AI has to be done in a way that prioritizes dependability, reliability, and verification of the validity of its outputs through technical means as well as human oversight instead of assuming it will work the same way it worked in a demo.

Bringing AI on board is similar to hiring a junior employee. A newly hired employee learns from the company's SOPs, handbooks, prior work outputs, and above all by observing and interacting with colleagues. AI—in contrast to software—can do the same, learning from the data and the human feedback it is presented with. That is one of the greatest promises of AI, especially agentic AI, when deployed correctly: the ability to learn and fully adopt a specific organization's specific way of doing things, thus becoming an invaluable member of your team. Yet, this also means that training and oversight (often called "human-in-the-loop") are essential and have to be implemented from Day One.

Also, applying one generic AI solution will not fit the needs of many companies. We often think about the corporate world classified by sectors and industries, but in reality, even within one space, every organization is different and there is no "average company." For example, Toyota, Rolls-Royce, and Tesla are all car manufacturers, but a generic solution developed for the average of the three would require each of the three companies to change its specific processes to adapt to this new tool and, as a result, incur significant costs. This is unlikely to add much value to any of the three organizations.

AI can overcome this traditional way of thinking about the world, given its above-mentioned ability to learn and thus adapt to specific workflows in a tailored, n of 1 fashion. This means the right enterprise AI framework mirrors and augments your own processes, instead of forcing you to adopt someone else's way of doing things.

INTEGRATION MAKES ALL THE DIFFERENCE

To make it brief: AI in isolation is pointless. The value of bringing AI into the enterprise is almost solely derived from deep integration into your data and across your organization. The power of AI lies in its ability to ingest and integrate large amounts of (structured as well as unstructured) data that form the input layer for meaningful

insights into the inner workings of your company and any output that needs to be produced. Whether it is about monitoring and planning production processes, performing root cause analysis for downtimes and incidents, or generating reports, memos, and regulatory documents, it is the data distributed across your organization that enables AI to generate all of these insights and outputs in a fraction of the time and cost it currently takes.

For an organization to integrate AI in the best ways possible, the right language must be chosen. No, we are not talking about English, French, or Mandarin, but about the programming languages underpinning AI and most enterprise applications. While AI firmly lives in the world of Python—which is ideally suited for experimentation, rapid development and prototyping, and generally data science work—most enterprise applications, the glue that keeps companies running, reside in Java, ensuring stability, scalability, dependability, and security for the applications that power our workflows.

This also means that to deploy AI in the enterprise—not just in isolated, sandboxed pilots but at scale and in production—we need to bridge that divide and leverage Java frameworks that can natively embed AI. In this way, companies get the best of both worlds, with the ability to seamlessly combine the power of AI with the strength of their existing enterprise tech-stack.

IT'S ALL ABOUT THE PEOPLE

Understanding, selecting, developing, and implementing the right technology is only one half of the equation; people are the other. Change management is a critical factor for the success or failure of any AI initiative, and its importance cannot be overstated.

Therefore, it is key to understand incentive structures and perceptions among company leadership and the wider stakeholders of a given process; to identify, coach, and support champions who can galvanize excitement and buy-in around new initiatives; to agree on success criteria and monitor them from the start; and to address concerns and fears openly and constructively. Implementing AI at scale is a transformative process, and it must be seen not as a sprint but a marathon that requires preparation, alignment, and commitment.

AI is undoubtedly one of the most defining trends, if not the single most defining trend, of this decade, and it will permanently change the way companies operate. Those that successfully adapt to this reality will undoubtedly take market share from those that don't, some of which will become the next Nokia or Kodak. But as leaders navigate this transformation, it is worth bearing in mind and communicating clearly to boards and teams alike that change is a natural part of the order of things—

customer preferences change, products become obsolete, new competitors emerge, and global market conditions shift.

To survive and thrive in the market, companies have always had to constantly adapt and evolve, and this time is no different. The only question is whether companies manage to keep up with the world changing around them by making smart choices, or get sidelined by others who are more attentive in their observations and more assertive in their actions.

BRINGING IT BACK TO THE BEGINNING

At the beginning of this article, we asked two questions that we can now answer: There are widespread misconceptions regarding what AI is, how it works, and how to best judge in the procurement phase whether it solves an actual business need in the way we need it to. The old playbooks for judging software demos and pilots do not apply to AI solutions. The challenges, risks, and pitfalls are different ones, and they need to be evaluated accordingly. Companies that embark on a pilot and publicly announce a successful AI deployment often realize shortly thereafter that actual production use of AI is a very different matter altogether. In times of extreme competitive and board pressure, it can all too often be tempting to signal to the outside world that one isn't obsolete but is instead at the forefront of leveraging AI. Confusing early sandboxed and often poorly planned pilots with actual production use of AI is a costly but sadly common mistake.

But what can be done? How can leaders ensure that their companies implement AI successfully at scale and in production? The key to that success is to act on a solid basis of understanding the basic workings of AI, while drawing on one's core strength: No one knows your own company better than you do, and no one has the same depth of insight into your use cases.

Aware of that superpower, leaders can then follow the steps and learnings outlined above to help them move from ideation and use case definition to designing meaningful pilots and managing a successful enterprise rollout. By doing so, they can position themselves in the best possible way against their competitors. In Allbusiness.com's "15 Quotes on the Future of AI," Ginni Rometty, former CEO of IBM, stated, "AI will not replace humans, but those who use AI will replace those who don't." [AQ](#)

Martin Bittner, MD, DPhil, is a two-time founder in workflow automation. He currently serves as co-founder & CEO of Redouble AI, a tech company backed by Y Combinator that leverages GenAI to automate enterprise workflows. Before that, he was the CEO of Arctoris, a successful pharma workflow automation company headquartered in Oxford which he co-founded.

MCE is your reliable partner for continuous success with agile people development solutions.



10,000,000

participants on AMA & MCE programmes in the last 10 years



92%

of Fortune 1,000 companies are our business partners



96%

of participants report they are using what they have learnt at AMA & MCE



1,000+

expert facilitators globally



100+

Open Training Programmes running throughout EMEA



98

year's experience working with our clients around the globe

For more information please contact:



+32 2 543 21 20

@ info@mce.eu

Visit **www.mce.eu**

