

Why your AI proof of concept isn't enough – and how to fix it



► Introduction

Today's business landscape evolves at an increasing and unprecedented speed, driven largely by one factor, Artificial Intelligence (AI). Organizations have a clear need to harness this technology for tangible business value, but this demand becomes more pressing with each passing day, as Generative AI (GenAI) and other AI technologies continue to advance at an astonishing pace.

New tools and frameworks appear regularly. Since the beginning of 2025 alone, Chat-GPT and Gemini have each released three significant updates. DeepSeek disrupted the space with a model its developers claim to have built at a fraction of the cost of competitors, and Grok-3 debuted with differentiated capabilities in terms of processing speed and access to real-time information.

In parallel, other AI technologies have exploded. As organizations increasingly turn to agentic AI to handle complex workflows, this new technology's fully autonomous systems leave yesterday's chatbots in the dust. Advancements in processing power and algorithms have increased the feasibility and prominence of edge AI – where local devices analyze and action data, instead of pushing it to the cloud – while also enabling new use cases in hyper-personalization and boosting the precision of predictive analytics to unprecedented levels of accuracy.

This fast-moving pattern presents an obvious challenge for organizations: simply keeping up. They must stay agile and update their strategy to take advantage of new developments but also need to scale solutions to avoid getting caught in an endless testing loop, where developers trial every new technology and projects never move beyond Proof-of-Concept (POC).

Indeed, many enterprises succeed in creating AI POCs, but then struggle to scale these initiatives into robust, enterprise-level solutions. Accurate, high-performing algorithms are part of the answer. Depending on the nature of the use case, POCs will leverage several forms of underlying AI, but no matter which technologies teams select or how stringent the performance criteria, accurate algorithms are not enough for successful scaling.

Therefore, scaling demands more than just great technology; it requires robust process and governance, strong underlying data, and effective change management to drive tool adoption. Given that businesses want to scale AI to avoid ceding advantage to competitors, the question quickly becomes: "How can we scale AI POCs successfully?"

With this challenge in mind, we spoke with Luke Gerdes, a Partner at Uvance Wayfinders, to explore how organizations can successfully scale AI projects, whilst delivering against key metrics.


► Laying the foundation for scalable AI solutions

Organizations that successfully scale AI consider impact before they commit the significant cost, time, and energy required to realize an enterprise solution. Project teams need to dedicate as much attention to benchmarking solutions' cost savings and revenue gains as they do to measuring technical accuracy. POCs that cannot offer a compelling and measurable Return on Investment (ROI) in a reasonable timeline should not move ahead.

“ Finding the perfect balance is essential, and somewhat of a Goldilocks situation. If the scope of AI is too narrow, significant value may be overlooked. Conversely, if it's too broad, the implementation can become cumbersome, overly complex, and costly. Defining the problem's boundaries with precision ensures that the level of automation is 'just right'.

– Luke Gerdes

”



Modular design is another critical element to implement from the very beginning, especially given the aforementioned pace of change. Think of it akin to Lego bricks: teams can swap out a single block as technologies evolve, without rebuilding everything. Ultimately, well-modularized POCs drive impact sooner, because developers can better prioritize and focus on path-critical components that play the greatest role in restricting scale, instead of working to re-write mangled code from scratch.

Business leaders sometimes view modularity as an arcane technical consideration, but it has immediate connections to strategy and use case selection. As modular design allows for re-use, where some blocks support more than one offering, this approach encourages growth.

Organizations don't have to start with monolithic 'do-it-all' systems that can require massive investment. Instead, they can right-size their entry points to AI by building the highest value tool first, then bolting on new modules to handle downstream use cases. The first module predicts sales, another optimizes inventory, and a third generates reports—all linked but independently tweakable. Ultimately, modular AI is about building smarter, not bigger, and is best suited for complex real-world applications where off-the-shelf tooling just won't cut it.

► People and process, not just technology

Assuming that POCs are well designed and demonstrate enough ROI to scale, organizations need to focus on process to scale AI successfully. Here, Luke highlights the “critical role of data governance, which involves deciding how AI solutions use, and perhaps even change, business-critical data”. Formalizing this process is especially important in heavily regulated industries, where data privacy and regulatory restrictions abound. Fortunately, most organizations have years of experience defining access rights for employees and for the technologies that preceded AI, so defining effective data governance for AI typically involves adapting existing practices, rather than creating new approaches from scratch.

Data stewardship meanwhile can be far more challenging, often needing to involve the entire organization. Effective stewardship entails making sure that data has high construct validity, providing a faithful representation of processes and events that happened in the real world. If front line factory workers fail to calibrate sensors, or if sales personnel keep their ‘real’ leads in a spreadsheet instead of the company's CRM, data stewardship will suffer because systems will record inaccurate information.


Fixing this problem almost always requires a focus on people. Process changes that incentivize employees to prioritize data quality can ensure AI has the data it needs to run long term, but this requires organizations to make significant commitments to change. AI needs in-house business data to drive the most impact for an organization, and it's often infeasible to fix widespread data stewardship issues while still doing business the way it's always been done.

“

If data is the new oil, then poor governance and inadequate stewardship create uncapped wells that waste business-critical information that could be used to support AI.

– Luke Gerdes

”





► Adapting to a constant state of flux

When AI is involved, everything changes constantly. The data shifts as operational and customer contexts evolve. Models, fine-tunings, and other customizations are retrained in response. As every component of an AI solution requires continuous monitoring and updating, effective version control becomes path critical.

AI Operations (AIOps), which formalizes this complex loop of versioning and testing to ensure that AI systems maintain their relevance and effectiveness, is a critical capability for scaling AI successfully. It requires the careful design of decision frameworks that delineate the responsibilities for updating models and validating changes, all while evaluating cost/performance trade-offs. For example, a smaller and lighter model that is less accurate than heavier alternatives is a better business choice for some use cases, because its reduced size makes the model simpler to maintain and requires less compute to run in production.

The complexity of AIOps makes it a specialized skillset that successful organizations hire against. But the need for technical expertise is broader. Organizations that don't invest in personnel with contemporary skillsets will struggle to scale AI effectively.

Unless organizations hire highly skilled personnel and invest to maintain their skills as the field rapidly advances, it's unlikely in-house talent will be enough. Enterprises do not always require a large and expensive permanent team, and some tools allow small teams to build impressive solutions, but any organization without a dedicated strategy for technology talent is likely to stumble in their efforts to drive long-term value from AI.

Furthermore, while tech talent is critical to success, effective enterprise AI requires participation from Subject Matter Experts (SMEs) who know the details of the use case an AI solution seeks to address. As Luke puts it, "We know how to measure drift in a supervised learning algorithm—by looking to see if the error increases between predicted and actual values—but model monitoring needs different sorts of business-specific tests for GenAI because those technologies always involve a stochastic element that introduces some randomness into answers." Given that SMEs know "what right looks like", best practices to maintain complex AI solutions often include periodic reviews with domain experts to help quantify performance over time, leading to more desirable results.

Some use cases also require SMEs to serve critical safety functions, such as in heavy industries. While AI is well-positioned to increase factory productivity by recommending set points for equipment, full automation could amplify risk. Human-in-the-loop approaches are required when bad recommendations jeopardize human safety. Balancing workloads between people and machines allows organizations to comply with regulatory standards while maximizing the potential of AI.

Case Study: Process optimization for a large production plant

Challenge

Optimizing set points for industrial processing equipment to increase production efficiency. Here, the client was facing the classic manufacturing problem whereby higher quality outputs led to lower volume, and vice versa.

AI solution


AI recommended set points that transcended the traditional trade-off curve, achieving both higher quality and volume simultaneously. Human subject matter experts reviewed and approved AI recommendations to ensure safety and reliability.

Key result

The AI-enabled optimization process broke the old production paradigm by simultaneously increasing product volume and product quality—all without any significant capital investments in new equipment.

Key takeaway

Designing AI solutions requires a balanced approach between automation and human oversight, resulting in customized solutions that fit specific use cases.



► The importance of change management

No matter how performant solutions may be, they're without value if unused. Resistance to adopting new processes and an overreliance on legacy workflows are some of the common challenges that undermine AI adoption.

Wayfinders can help with this transformation. Organizations can receive support to find direction, design strategic paths, and gain assurances during implementation. Here, Luke recommends keeping business strategy and end users top of mind throughout the scaling journey. Feedback from the former ensures solutions remain anchored in business value and facilitate employee behaviors that match strategic priorities. Input from the latter means that tools fit the reality of day-to-day workflows, even when those workflows must adapt to meet new strategic requirements or technological shifts.

While many clients may have internal subject matter experts, they often lack the cross-functional knowledge necessary to scale AI solutions effectively. Here, Wayfinders acts as a strategic partner, designing and implementing AI solutions that are not only technically sound but aligned with business goals.

An iterative test-and-learn approach can also help to build buy-in. Because they're rapidly built and limited in scope, POCs always generate a lot of feedback that can improve a solution's first enterprise-grade minimum viable product. But it's often advantageous to repeat the pattern by rolling-out the first enterprise release to a sub-section of the target user group to gather feedback, refining the solution before its release to the wider community.

“ The goal for organizations when it comes to AI adoption should be to avoid simply implementing technology for technology's sake and instead focus on practical use cases that can be scaled over time. As soon as you lead with the technology and try to find a problem, you're a hammer looking for a nail.

– Luke Gerdes

”

► Conclusion

Unlocking the full potential of AI

Put simply, AI has the power to transform business operations for the better, but this requires a strategic and comprehensive approach. Scaling is ultimately a complex journey that extends far beyond just available technology. Business alignment is pivotal, as it ensures that AI initiatives are not just innovative but directly contribute to organizational goals.

However, with careful use case selection, appropriate project design, robust governance, diligent data stewardship, and the integration of human expertise, organizations can effectively scale AI solutions, realize significant benefits and improve overall ROI.

Because Wayfinders supports clients from opportunity to outcome, our strategic consulting services bridge this gap and help drive these necessary transformations. As an AI-native and impact obsessed organization, we break down traditional industry siloes, helping organizations overcome resistance to change, while fostering a culture of continuous innovation.



Want to scale your AI solutions beyond the POC stage?

Reach out to our Wayfinders team to discuss your next steps with our experts.



[Contact us to find out more](#)

► Contributors



Luke Gerdes

Partner, Fujitsu Uvance Wayfinders

Luke Gerdes is a partner with Fujitsu Uvance Wayfinders. With a career spanning over 15 years in analytics and research, he works at the intersection of management consulting, artificial intelligence, and machine learning to lead large, use case-backed transformations that take ideas from proof of concept to scaled production. He delivers more than just technology and focuses on end-to-end impact by surrounding innovative solutions with the process, governance, and teams required for these technologies to thrive. Prior to management consulting, Luke was an academic and researcher. He was a civilian professor at West Point, where he leveraged network science, machine learning, and related technologies to model terrorist organizations and other adversarial groups. Luke describes himself as a “computational social scientist who’s learned to redeploy his skills to drive business value from end to end.”



luke.gerdes@fujitsu.com



[Luke Gerdes](#)