90% of AI POCs never scale. Will yours?

Blend™

THE CRITICAL 7

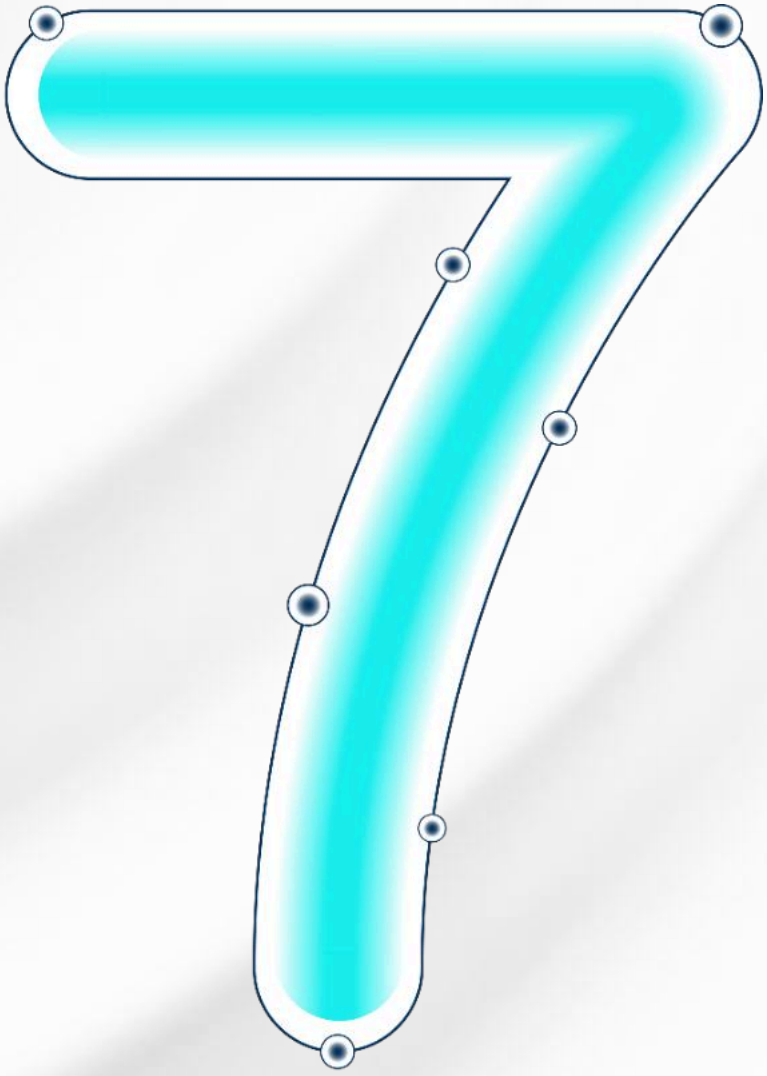

PROVEN STRATEGIES FOR SCALING AI

Contents

Foreword

Oz Dogan

President of Solutions and Service Lines, Blend

The pace of change we've witnessed since GenAI burst on the scene in late 2022 has been breathtaking. Gartner **expects** 80% of independent software vendors to incorporate generative AI capabilities into enterprise applications by 2026. Nearly every enterprise has AI pilots underway. AI agents, which can take actions and interact with each other to orchestrate sophisticated tasks, are just around the corner. Humanoid robots are already transforming warehouses and assembly lines. And this is just the beginning.

At Blend, we think the time is right to offer our framework for Scaling AI successfully. The Critical 7 is a collection of insights drawn from our client engagements that offer a roadmap for organizations looking to integrate AI in Business Operations. I am convinced that businesses that bind AI to business strategy, build essential technical foundations, continuously innovate around AI's unique strengths, and infuse AI into their culture will quickly outperform their competitors.

Blend's AI framework delivers proven value. Companies that use it are four times more likely to get their AI initiatives off the ground. Once they're up and running, success builds on itself to accelerate AI-driven transformation.

AI is the biggest technical advancement of our lifetime, and we are just getting started. I'm incredibly excited about the potential of these transformative technologies. As you read the following pages, I hope you'll come to share my enthusiasm. This is a once in a lifetime moment. Let's get to work building the future we envision.

Oz Dogan

Meet Our Experts

Ozgur Dogan is President of Services and Solution Lines at Blend, where he leads the delivery of AI, data, and technology solutions that drive real business impact. With over 20 years of experience, he brings deep expertise in scaling transformation for enterprise clients.

Ozgur Dogan
President, Solutions & Service Lines

Rob Fuller is Chief Solutions Officer at Blend, where he leads the development of AI and data solutions that align strategy with execution. He brings deep experience in enterprise transformation and helps clients turn innovation into measurable impact.

Rob Fuller
SVP, Chief Solutions Officer

Mike Mischel is Senior Vice President of AI Transformation at Blend, helping clients scale AI from vision to value. He specializes in turning complex challenges into actionable strategies that drive adoption, impact, and long-term success.

Mike Mischel
SVP, AI Transformation

Adam Mincham is Chief Marketing Officer at Blend, where he leads brand, growth, and go-to-market strategy. With deep experience in customer experience and digital transformation, he helps position Blend at the forefront of AI and data innovation.

Adam Mincham
Chief Marketing Officer

Introduction

It took ChatGPT **just two months** to reach 100 million monthly active users, making it the fastest-growing consumer application in history. It's estimated that **more than 700 million** people will regularly use AI tools by 2030.

Spending on AI is expected to grow 30% annually and top $630 billion by 2028, with generative AI comprising 32% of all AI investments, **according to IDC.** The speed with which AI has raced to the forefront of corporate strategy has been astonishing.

Yet the vast majority of AI projects so far have failed to deliver tangible benefits or scale beyond the proof-of-concept stage. At Blend, we have observed that these failures have several traits in common. From ensuring strategic alignment and strengthening data infrastructure to overcoming technical barriers and managing workforce transformation, the seven challenges outlined in this eBook are critical to realizing AI's full potential.

AI has gone mainstream so quickly that many business leaders have been caught off guard. Wary of risks and fearful of failure, they focus their organization's efforts on incremental improvements rather than transformative change. The failure of these cautious initiatives to yield dramatic results breeds skepticism within the organization and solidifies resistance to change. Lacking a strong data foundation, AI models struggle to deliver reliable insights. High infrastructure costs and the need for radical new development techniques amplify perceived disruption and the image of AI as a money pit.

In contrast, organizations that forge ahead with AI projects aligned with business strategy, deliver demonstrable ROI, and are championed by senior executives who see workforce **productivity gains of more than 40%**, lower costs, and happier customers. These companies have taken the Critical 7 to heart as they future-proof their businesses against the inherent uncertainty of AI's next evolution.

The insights in the following pages were developed by Blend's experts leading AI projects at some of the world's largest companies. Let them guide your strategy as you prepare for the age of AI.

The Critical 7
Scaling to Production AI

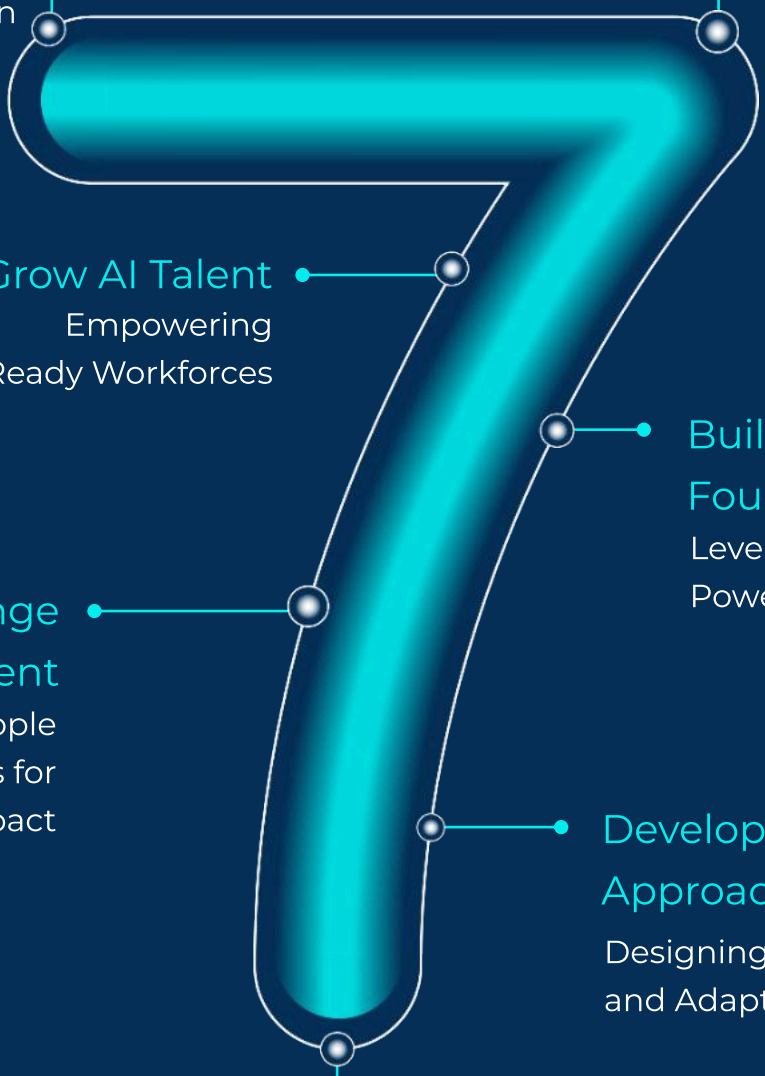

Integrate Business Strategies
Maximizing Value from AI Initiatives

Create Trust
Setting Expectations, Leading Ethics, and Driving Adoption

Grow AI Talent
Empowering AI-Ready Workforces

Build Strong Data Foundations
Leveraging AI to Power AI

Enable Change Management
Aligning People and Process for Meaningful Impact

Develop a Technical Approach
Designing for Performance and Adaptability

Accelerate Innovation
Building for Continuous AI Evolution

Chapter 1

Integrate Business Strategies:
Maximizing Value from AI Initiatives

Integrate Business Strategies

AI success depends on strategic alignment with business objectives, not just technological ambition. While most enterprises have AI initiatives in place, our experience is that 80% of AI projects never achieve broad production use.

A primary reason: lack of a clear connection to financial or operational goals.

We see the same pattern in client interactions. Executives say they want to apply AI for a competitive advantage, but the projects they fund deliver only incremental improvements to existing processes.

AI projects must be chosen based on value that sets your business apart from your competitors. That begins with understanding how your business derives value. Projects that yield only incremental benefits can often be dispatched more quickly by purchasing a solution from a vendor.

AI shouldn't be simply an enabler of business strategy but a partner in business transformation. It is a chance to move beyond top-down, point-in-time strategies, to adaptive strategies that change in real time.

Take the example of an apparel manufacturer. Upon detecting a change in consumer preferences for a style or color, the company would traditionally assemble a team to analyze the data, agree on product changes, model the impact on manufacturing and the supply chain, and cascade any planned changes through the organization. The process can take months.

Key Takeaways

Drive the Business Case
AI projects must be grounded in financial or operational value. Ensure you have built a sound business case that accurately projects ROI (return on investment) prior to beginning work on a POC.

Align to Business Strategy
First understand where the most value is for your organization. Generative models can help shape adaptive, real-time strategies by surfacing insights from across the organization — making strategy everyone's business.

Human First Design
Reimagine workflows by designing with AI agents and humans working in parallel. Appoint AI champions with executive support to drive transformation and adoption across the business.

An AI agent can now be deployed to monitor consumer sentiment in social media, communicate changes to management, and simultaneously calculate the operational impact, costs, and revenue. Upon approval, another agent can draft communication of the changes appropriate to each part of the organization, while a third prepares purchase orders, and another creates manufacturing change plans. All of these activities can proceed in parallel, compressing the time to implementation to hours.

"With advances in technology, a 20% increase in EBITDA (earnings before interest, tax, depreciation, and amortization), can now become a reality," says Mike Mischel, Senior Vice President of AI Transformation at Blend.

When AI and strategy are woven together, the model can evolve the strategy continually instead of being revisited once a year. Ideas from deep in the organization that would never catch the attention of top executives can be surfaced automatically. Strategy becomes everyone's business.

Don't try to boil the ocean. If AI can save $3 million on a $150 million engineering project in year one, that's a start. Choose high-impact applications and expand AI adoption incrementally. Even small changes can have bottom-line value when carefully targeted. For example, if your goal is to speed up HR onboarding, focus on the sales organization where faster time-to-value has tangible business benefits.

Look for projects that reinvent process and experience with AI at the center. This often means discarding linear thinking and reimagining processes around AI agents working in parallel with humans.

Choose AI **"champions"** to shepherd projects through to completion. They need executive support and the freedom to tackle tasks that carry acceptable levels of risk.

> Companies focus on headcount reduction, which usually affects already marginalized groups like call centers and HR."

Mike Mischel
SVP, AI Transformation

Chapter 2

Build Strong Data Foundations:

Leveraging AI to Power AI

Build Strong Data Foundations

As enterprises grow more sophisticated in their use of AI, their dependence on data increases - whether in building the right retrieval-augmented generation (RAG) for unstructured data, creating embeddings for data representation in AI use, or choosing the right data and metrics to answer questions. The union of good data with good data documentation creates transparency and trust.

The scope of the challenge is initially limited to the scope of the AI POC but will expand over time as more capabilities are needed. AI can help solve this problem and cope with disparate data and systems, but success is always dependent on using the right data.

AI models require data and lots of it. That's a problem for many large enterprises, whose information is fragmented across multiple databases, SaaS applications and departmental silos. Unifying terabytes of disaggregated data is such a daunting task that it can stop AI initiatives in their tracks. **Over 40% of data leaders** face challenges due to these fragmented models, which hinder AI adoption at scale.

Total data unification should be the goal long term, but we can't let it prevent us from adopting AI. A better approach is to bridge silos dynamically, enlisting AI as an aid. AI bots can scour data stores, looking for and labeling common data types. They can highlight duplicate or conflicting information and identify ROT (redundant, obsolete, and trivial) data that shouldn't be part of a model. Reducing bottlenecks caused by inconsistent or obsolete data paves the way for faster decision-making in real time.

Key Takeaways

Improve Data Quality
Leverage AI to clean, label, and contextualize fragmented or inconsistent data—reducing ROT and improving trust in model outputs.

Reduce Mis-Use of Data
Use AI to surface and resolve conflicting definitions (e.g. "revenue" across departments) and promote standardized understanding organization-wide.

Manage Data Governance
Empower AI to detect anomalies, classify sensitive data, and enforce governance policies at scale—without slowing agility or innovation.

Reduce Misuse of Data

Agents can reduce data misuse by contributing to contextual understanding. For example, accountants may define revenue as money earned during a period that meets accounting standards, while sales teams factor in bookings, even if no money has changed hands. AI can point out such discrepancies and suggest standard definitions. This permits organizations to connect enterprise data across functions to enhance accessibility and security.

> "Ask the model how much data preparation can be automated with what techniques to ensure high quality. Then ask AI for help cleaning that data."

Rob Fuller
SVP, Chief Solutions Officer

Enlist AI for Data Governance

Lack of strong governance will quickly send AI initiatives off the rails. Poor data quality leads to inaccurate predictions and flawed decision-making, undermining trust. Inadequate security measures risk privacy breaches. A lack of standardized architecture can slow AI deployment.

Once again, AI is your friend enhancing governance. AI models can detect anomalies and inconsistencies and suggest corrections based on patterns or policies. It can identify sensitive information and classify many data types automatically. Policy violations or risky behavior can be spotted in real-time, and remedial actions suggested. Governance can also be built into retrieval-augmented generation to ensure that AI-generated responses conform to business policies.

Chapter 3

Develop a Technical Approach:

Designing for Performance and Adaptability

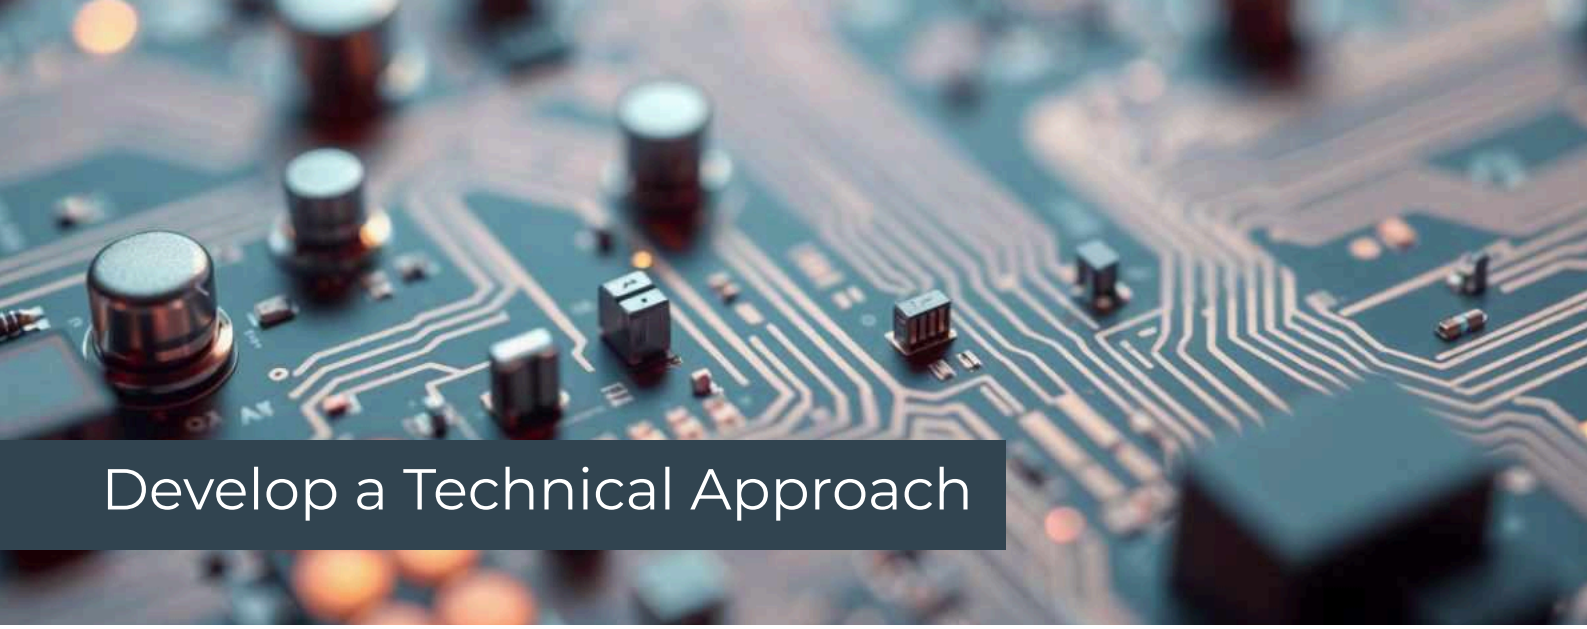

Develop a Technical Approach

Engineering AI: Architectures, Models, and the Science of Scale

Instead of a monolithic AI strategy, enterprises should adopt a flexible framework that allows for continuous improvement and scaling. Bet on models getting continually better and cheaper, but also your needs becoming more specialized as AI use expands. Don't lock into one model or even one model provider for all your needs.

Separating the prompt from the RAG and the model creates modular designs that can more easily support future capabilities. Layering in AI-controlled guardrails and dynamic routing can help scale the adoption of new models as they develop. The goal should be to create a modular and adaptable AI architecture that flexibly accommodates new innovations as well as changes in the business.

While experimentation is critical, letting a hundred flowers bloom is a path to waste and confusion. Evaluation of new models and services should take into account their costs, infrastructure impact, and skills demand. Proposed new AI applications should meet defined business value metrics such as targeted productivity improvements, gains in customer satisfaction scores, or access to new markets.

Technical challenges fall into three major areas: development, scaling, and accommodating differing modalities.

Key Takeaways

Adaptable Architectures
Design modular AI systems that decouple prompts, models, and RAG components — enabling rapid iteration, flexible scaling, and integration of future innovations.

Probabilistic vs. Deterministic Comfort
Embrace AI's probabilistic nature by setting fit-for-purpose accuracy thresholds, building human-in-the-loop evaluation systems, and turning uncertainty into smarter decisions.

Domain Models
Fine-tune smaller models or use RAG to meet domain-specific needs—balancing performance, cost, and speed without retraining large foundation models.

Deterministic vs Probabilistic

AI algorithms are probabilistic; they generate different outcomes based on probability distributions. This can be a jarring change, but probabilistic reasoning is useful in uncertain scenarios. They enable more nuanced decision-making and represent the messy reality of the world instead of pretending everything is black and white.

That means software testing needs to be done differently. Models must cycle through the same problems continuously, with results evaluated by human experts or other AI models. AI evaluation frameworks are critical to ensuring businesses adopt the right models for the right issues. Part of the system design is putting in guardrails to ensure the output is always acceptable and to retry if not.

Generative AI models can be improved through fine-tuning and RAG to tune accuracy and performance, but the most important skill to develop is figuring out the right use cases for generative models compared to more traditional ML or deterministic programs.

Because AI model outputs can be unpredictable, it is essential that trust be designed into them. When challenged, models should be prepared to explain their reasoning process and embellish results with citations and footnotes. Human experts should be involved in testing and training to guard against model drift and hallucinations.

> "Turn the probabilistic process into a power."

Rob Fuller
SVP, Chief Solutions Officer

AI Infrastructure at Scale: Designing for Performance and Adaptability

AI models also scale differently from traditional applications. The training process can tie up computing resources for days, but inferencing (or using a trained model to make predictions) uses much less computing power. Organizations should consider using cloud platforms for training while inferencing on their own infrastructure. Hyperscalers can quickly adapt to continued hardware improvements and reduce the need for costly fixed infrastructure.

While prices for training and inference of models are coming down, each use case needs to consider the token costs and performance realities to ensure the use case is viable - or design an architecture to make it feasible.

Put your faith in RAG and fine-tuning. The cost of training foundation models is so high that few but the largest enterprises can consider it. RAG systems use a retrieval database such as a vector store or search index to deliver results faster and more cheaply with lower maintenance overhead. There is no need to retrain the entire model when conditions change, and RAG reduces the risk of errors and hallucinations by grounding answers in well-governed sources.

Fine-tuning adapts a base model to an exact domain, such as legal, medical, technical, or company-specific. The result is higher accuracy, more relevant outputs, and better handling of technical data. It requires more up-front work than RAG, but the payoff is better alignment with your specific needs. Both options scale better than foundation models because updates only require updating the retrieval database rather than retraining the entire model.

Know When "Good Enough' is Good Enough

It can be complex to assess when a model is sufficient to be put into production, and diminishing returns of further training set in. **"It’s the 80-20 rule,"** says Rob Fuller, Senior Vice President of Technology Solutions at Blend. **"You can get to 80% accuracy or more with 20% of the effort, but improving the accuracy of the last 20% requires 80% of the effort."**

Knowing what accuracy level is right depends on the situation. A scenario in which lives are on the line requires greater reliability than one involving classifying documents. The more backups and human checks you have in place, the greater your tolerance for error can be.

Be aware of the limits of probabilistic reasoning. Perfection may not be achievable, but the tradeoff can be more creative out-of-the-box thinking. Make probabilistic reasoning a virtue.

These equations will change over time as model quality improves and prices fall. Using smaller models and applying RAG can significantly reduce infrastructure costs. Organizations should evaluate the full range of available options when choosing models.

Chapter 4

Accelerate Innovation:

Building for Continuous AI Evolution

Accelerate Innovation

Fast, Flexible, and Future-Ready: Building for Continuous AI Evolution

AI is evolving at an unprecedented pace, outstripping the ability of many organizations to define internal policies and strategies. However, such guardrails are important for enabling companies to embrace innovation while ensuring AI implementations are structured, scalable, and aligned with business needs.

Speed of innovation is critical for scaling AI. Rapid iteration and experimentation allow continuous prototyping, testing, and refinement. Think differently about problems. We are often trapped by our historical models and constraints. Agents allow a more collaborative process to be employed. Envision how linear processes can be redesigned to work in parallel.

Abandon historical constraints. Automation typically focuses on data routing and rules scenarios, but new capabilities can factor in automated creation, instant analysis and parallel human collaboration. The rules of automation have changed.

Companies that quickly adopt new innovations will have a competitive edge. We recommend creating an **“AI lab”** where technical experts can experiment with the latest technologies and techniques and recommend adoption strategies to AI steering committees. They should continually investigate using smaller, faster AI models, which often outperform large, generalized models for targeted business use cases.

The speed of innovation should be balanced with business impact. A small number of high-impact projects will realize a greater return than many incremental ones.

Key Takeaways

Human in the Loop
Innovation thrives when humans guide, challenge, and refine AI—ensuring outputs are grounded, creative, and trusted

Reinventing the Process
Rethink workflows to enable AI-driven parallel execution, shifting from linear automation to collaborative augmentation.

Modular AI Architectures
Design flexible systems and shared infrastructure that allow rapid experimentation without runaway complexity or cost.

Use the returns from the impactful projects to reinvest and amplify value.

Humans should be kept in the innovation loop. Experts are needed to distinguish between breakthrough ideas and flights of fancy or hallucinations. A collaborative approach also protects against resistance from those who fear AI will take their jobs.

Innovation speed will also follow technology evolution as models get faster, cheaper, and better. The base may seem overwhelming, but keeping with it is often simply a matter of applying more computing power and refining it with the same techniques. Employing modular architectures allows innovation to scale with improvements to models and infrastructure.

From Automation to Augmentation: Mastering AI's Innovation Curve

Grouping investments by business value controls for total costs and allows for refinement and reuse in other areas. For example, a knowledge graph we created for one client was found to have broad enough applicability that it now powers three distinct business applications. Think in terms of how features can have value beyond their immediate use case.

Creating a culture of innovation is about shaping behaviors, mindsets, and systems that encourage creativity, risk-taking, and continuous learning. Employees should be given the time and encouragement to think creatively with rewards based on originality and imagination rather than success. People will embrace AI more readily if the focus is on augmenting human tasks rather than eliminating jobs. Everyone in the organization should be encouraged to innovate, not just a few champions. Leaders should visibly cheerlead innovation efforts.

> "Start with the capabilities of the AI and think about how it can solve a problem in a different way. Think of it as a supporter of the process, not just an app to be built."

Rob Fuller
SVP, Chief Solutions Officer

Chapter 5

Enable Change Management:

Aligning People and Process for Impact

Enable Change Management

Effecting change has traditionally been an arduous and disruptive process. Strategic initiatives conceived at the top of the organization were communicated downward in a cascading effect that, like a game of telephone, invited distortion and misinterpretation. By the time the message reached the front lines of the organization, the strategy had often changed.

AI unlocks a systemic ability to rethink the way organizations manage change. AI models are good at seeing patterns and understanding consequences of change that may not be evident at the top. It can alert senior managers to unintended consequences and suggest alternative strategies. Communication can be customized to individuals and rippled through the organization in parallel.

This new dynamic has important implications for the way organizations handle change. Instead of hard-coding tasks into workflows, organizations should align their people around self-forming teams that dissolve and reassemble as priorities change. Linear work processes should be rethought in parallel, with team members pursuing objectives instead of handoffs. The whole organization becomes flatter and nimbler.

Key Takeaways

Dynamic Enablement
Tailor change strategies to different audiences by using AI to surface risks, identify resistance patterns, and customize support.

Personalized Communication
Reduce fear by transparently addressing the impact of AI and highlighting real, human-centered success stories across the org.

Leverage Behavioral Science
Apply proven behavioral frameworks to foster acceptance—introducing AI in low-stakes scenarios and celebrating bottom-up innovation.

Fostering Acceptance

Many people fear that their jobs will be displaced by AI, and to some degree, those fears are warranted. However, job displacement is a byproduct of many new technologies, and in nearly every case, technology change creates more jobs than it eliminates. The key to encouraging people to embrace AI is to equip them with the skills to use the new tools, reward those who respond positively to change and focus on initiatives that demonstrate clear value.

Change needs to be orchestrated with the cooperation of the people whose roles must change. While some will continue to prefer to be told what to do, dynamic enablement tailors change management strategies for the audience. Personalized communication and complete transparency mitigate fears. Leaders shouldn't try to sugarcoat the potential for disruption but should stress that AI is a way to support human expertise, not replace it. As the **Harvard Business Review** put it,

"AI won't replace humans, but humans with AI will replace humans without AI."

AI adoption should be orchestrated as a bottom-up and top-down strategy, ensuring innovation at all levels. The more people are involved in redefining roles and choosing the tools to support them, the more enthusiastically they will buy in. It is critical that education be offered to all who want it. Such a policy should align with a message from leadership that the company's future success depends on innovative uses of AI, and all can be part of that vision.

Leveraging behavioral science can help. People fear what they don't understand. Instead of saying, **"AI will eliminate this many jobs,"** say, **"AI takes over repetitive tasks so you can focus on higher-value work."** Celebrate success stories, particularly those that bubble up from below. Introduce AI in small, non-threatening scenarios, such as using note-taking at meetings. Encourage experimentation without consequences. Above all, trust people – and their AI agents - to help guide them. Given access to tools and the freedom to experiment, they will invent new ways to streamline processes and automate routine tasks.

> AI lets you tailor the way you communicate its value — aligning the message with people's personality traits to build trust and drive adoption.
>
> 
>
> **Mike Mischel**
> SVP, AI Transformation

Chapter 6

Grow AI Talent:

Empowering AI-Ready Workforces

Grow AI Talent

The widespread adoption of AI is being hampered by a major skills gap. By **some estimates**, there are enough skilled AI workers to fill only half the available jobs. **Yet Randstad reported** that only 35% of employees at companies seeking AI talent have received relevant training. Just 36% of tech practitioners strongly agree that their organization has the necessary expertise to implement complex AI projects, **according to Capital One**.

The disconnect between demand and the willingness of organizations to invest in developing AI skills is alarming. A **2023 study** by the IBM Institute for Business Value estimated that 40% of the workforce will need to reskill over the next three years due to AI and automation.

General-purpose training in prompt engineering has value but doesn't deliver the accelerated productivity of training customized to individual jobs and functions. Many organizations rely on junior analysts for AI model quality assurance, delegating one of the most important stages of model training to people without the appropriate domain knowledge. AI delivers the greatest benefit when trained by experts.

Collaborating with industry experts can bring specialized insights and the most creative adaptations to training efforts. Implementing hands-on, role-specific education programs can quickly bridge skill gaps and enhance employee confidence in working with AI technologies.

So can slipstreaming AI training into existing tasks. For example, in response to a prompt asking how the organization goes about qualifying sales leads, a generative AI engine can offer to generate a report

Key Takeaways

AI Literacy

Foster a mindset shift by contextualizing AI's role across job functions, helping employees see where and how AI fits into their day-to-day work.

AI Training Moments

Integrate hands-on, role-specific training into everyday tasks to build confidence and normalize AI use across the workforce.

Collaborative Engagement

Combine the strengths of AI natives and subject-matter experts to upskill teams and accelerate learning across functions.

or prepare a PowerPoint presentation to share with colleagues. Knowledge workers should be encouraged to consult generative AI on questions they might previously have asked colleagues. The more they see the positive impact of AI on their role, the more they will incorporate it into their daily work.

Reskilling doesn't have to be an all-or-nothing proposition. Organizations should make function-specific training available to those who show interest and encourage them to pass on their knowledge to peers.

By committing to structured AI training, businesses can not only alleviate the immediate talent crunch but also foster a culture of continuous learning and innovation, which is an essential foundation for long-term success in the AI-powered economy.

AI empowerment should be universal. The risk of selectively granting access to training is that the gap between the knowledge **"haves" and "have nots"** widens. Organizations will accept change more willingly if the tide lifts all boats.

> Embedding AI into an on-the-job training agent gives organizations the ability to constantly watch and insert helpful information into processes."
>
> 
>
> **Mike Mischel**
> SVP, AI Transformation

Chapter 7

Create Trust:

Setting Expectations, Leading Ethics, and Driving Adoption

Create Trust

Widespread reports of **“hallucinations”** in the early days of generative AI created concerns that the models couldn’t be trusted, but those fears have eased as models have improved and people have gained experience using them.

Well-trained and vetted AI models can be trusted with the vast majority of tasks enterprises need them to do. Convincing people of that fact is another matter.

Trust issues extend beyond output quality. Many people worry that AI will negatively impact their jobs and quality of work. The probabilistic nature of AI results also creates unease as inconsistent responses from generative models create suspicion.

There are also legitimate concerns about the potential for data exfiltration or the unintentional leakage of sensitive information and intellectual property into models and bad actors who exploit them.

Trust begins with transparency. Choose models that explain their work to the greatest degree possible. Select commercial generative AI services that clearly state their policies on data harvesting, intellectual property protection, and privacy.

Transparency includes setting realistic expectations about outcomes. Employees need to know not to accept every response at face value. They can triangulate by asking a question multiple times or submitting it on two or three different models. Solid governance principles and human-in-the-loop validation also help put fears to rest.

Key Takeaways

Set Expectations Correctly

Help employees understand the probabilistic nature of AI by clearly communicating what AI can—and can’t—do, and how it complements human judgment.

Ethical AI Practices

Choose models with transparent data usage policies and embed ethical guardrails into workflows to build trust and prevent misuse.

Risk Mitigation

Use governance frameworks, human-in-the-loop validation, and open communication around data security and job impact to proactively address concerns.

They also need clarity on what AI can and can't do. Probabilistic models aren't well-suited to problems requiring black-and-white answers. They excel in situations where the data or outcomes are uncertain and capturing that uncertainty can lead to better decisions. They help decision-makers evaluate risk factors and benefits, enabling more informed choices in complex scenarios. Probabilistic AI methods can often capture real-world complexity more accurately than purely deterministic approaches. Other machine learning models are more appropriate if the problem demands a hard yes-or-no decision.

"We expect AI to be accurate, but we don't expect that of humans," Fuller says. **"If we are going to take advantage of the humanistic nature of AI models, we need to evaluate them against the criteria we apply to humans."**

Concerns about job loss, data integrity and privacy protection require different approaches. Leaders need to loudly and frequently emphasize AI's role in augmenting human work, not replacing it. The first uses of AI should be on routine tasks that few people want to do. The goal is to show how AI helps humans do more, not make them obsolete. Show relatable examples of how AI made someone's job easier or more impactful, such as freeing customer service agents to focus on complex issues instead of mundane inquiries. As trust builds, so can AI's scope.

Privacy concerns are best addressed by communicating that all possible steps have been taken to ensure that accurate, timely, and unbiased data has been used to train models. Employees should feel free to try to trick AI models into divulging sensitive data. When they're successful, demonstrate that appropriate remediation steps are quickly taken.

> "LLMs are good at ingesting policies. You can use that fact to reliably prevent AI from crossing the line."
>
> 
>
> **Mike Mischel**
> SVP, AI Transformation

Conclusion

AI is Here to Stay

AI is Here to Stay

The rapid pace of change over the last two years has dramatized the futility of making long-term predictions about AI, but we are confident that the following trends will dominate the coming years.

AI becomes a strategic imperative. In the same way that the internet gave birth to a new cadre of high-growth businesses, AI competence will separate leaders from laggards. New companies will emerge that use AI to redefine financial planning, medical care, e-commerce and many other sectors. Incumbent leaders who fail to innovate rapidly with AI will quickly become also-rans.

AI agents will transform the workplace. We believe agents have significant potential to redefine workflows and offload routine tasks from humans. **"The marginal cost of creating new business processes once agents are defined will be radically different,"** says Rob Fuller, Senior Vice President of Technology Solutions at Blend.

However, it will be years before they can be trusted to book full vacation itineraries or orchestrate complex supply chains. Much work still needs to be done to define and implement standards for the APIs that will enable agents to work across multiple commercial services.

The action moves to small models. Training a large model from scratch is impractical for most organizations. Fortunately, most problem domains don't require those investments. Existing models, ideally small ones, can be highly effective and cost-efficient. Small models are cheaper and faster to fine-tune and will quickly come to dominate enterprise inferencing. All edge use cases will utilize small models. The largest drawback of small models is also a strength: they have less knowledge embedded in them, so they are typically just as good at language as large models but are easier to govern.

Focus shifts to responsible AI. Trust continues to be the Achilles' heel of generative AI. Commercial providers will be pressured to make their models more transparent. Enterprises need to put more ambitious evaluation frameworks in place to minimize errors and hallucinations. Organizations must also educate their workforce and customers to understand what AI can and can't do.

Prepare for growing threats. AI will raise the stakes in cybersecurity. AI-generated content will become increasingly indistinguishable from reality, raising concerns about misinformation, fraud, and security breaches. New technologies will need to be embedded in productivity tools to protect against these threats.

Organizations flatten. Knowledge is power. AI democratization will make expertise available to everyone, challenging traditional hierarchies. Organizations that embrace this dynamic will thrive as good ideas bubble up everywhere. **"Ideas on strategy will be based on merit, not just the highest-paid voice in the room,"** says Mike Mischel, Senior Vice President of AI Transformation at Blend. **"Brilliant ideas will come from voices that were marginalized in the past."**

Blend is a premier AI services provider, committed to co-creating meaningful impact for its clients through the power of data science, AI, technology, and people. With a mission to fuel bold visions, Blend tackles significant challenges by seamlessly aligning human expertise with artificial intelligence.

The company is dedicated to unlocking value and fostering innovation for its clients by harnessing world-class people and data-driven strategy. We believe that the power of people and AI can have a meaningful impact on your world. For more information, visit **www.blend360.com**