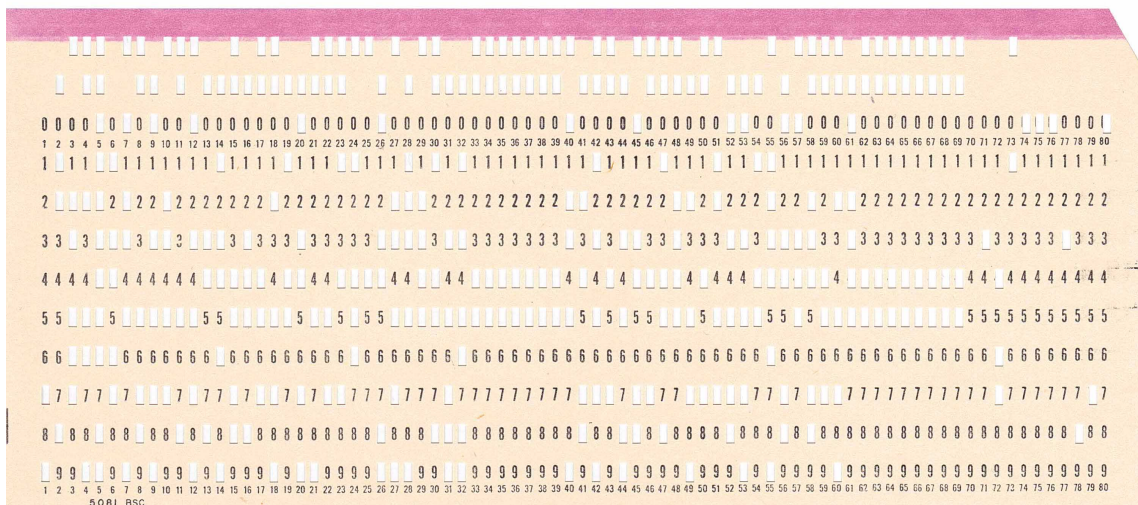


Technology And Analytics

Getting AI to Scale

by Tim Fountaine, Brian McCarthy, and Tamim Saleh

From the Magazine (May–June 2021)



CentreCentre/Patrick Fry Studio

Summary. Most companies are struggling to realize artificial intelligence’s potential to completely transform the way they do business. The problem is, they typically apply AI in a long list of discrete uses, an approach that doesn’t produce consequential change. Yet trying... [more](#)

Most CEOs recognize that artificial intelligence has the potential to completely change how organizations work. They can envision a future in which, for example, retailers deliver individualized products before customers even request them—perhaps on the very same day those products are made. That scenario may sound like science fiction, but the AI that makes it possible already exists.



Listen to this article

00:00 / 24:49

Listen to more HBR, on the [Noa app](#).

What's getting in the way of that future is that companies haven't figured out how to change themselves to meet it. To be fair, most have been working hard to incorporate digital technologies, in some instances genuinely transforming the way they serve their customers and manufacture their offerings.

To capture the full promise of AI, however, companies must reimagine their business models and the way work gets done. They can't just plug AI into an existing process to automate it or add insights. And while AI can be employed locally across functions in a laundry list of specific applications (known as *use cases*), that approach won't drive consequential change in a company's operations or bottom line. It also makes it much harder and more costly to get AI to scale, because each far-flung team must reinvent the wheel with respect to stakeholder buy-in, training, change management, data, technology, and more.

But that doesn't mean companies should try to overhaul the whole organization with AI all at once. That would almost certainly end in failure. A complete makeover is an enormously complicated process involving too many moving parts, stakeholders, and projects to achieve meaningful impact quickly.

The right approach, we've found, is to identify a crucial slice of the business and rethink it completely. Introducing changes throughout an entire core process, journey, or function—what we call a *domain*—will lead to a major improvement in performance that isolated local applications simply cannot match. It also will enable each AI initiative to build off the previous one by, for example, reusing data or advancing capabilities for a common set of stakeholders. We've seen this approach trigger an organic cycle of change within domains and, ultimately, build momentum for the use of AI throughout the larger organization as business leaders and employees see it work. Moreover, this approach

promotes a mindset of continuous improvement in the workforce, which is crucial because AI technology is advancing rapidly, requiring organizations to think of AI transformations as ongoing rather than one-time efforts.

Ultimately, the companies that can't take full advantage of AI will be sidelined by those that can—as we already see happening in several industries, like auto manufacturing and financial services. The good news is that over the past year many companies (even firms with limited analytics capabilities) have begun developing the skills required to capture AI opportunities, as the Covid-19 crisis forced them to alter the way they did business almost overnight. Now the challenge will be applying those skills to pull off larger initiatives. (To find out if you haven't scoped your initiatives correctly, see the following sidebar.)

Signs You're Thinking About AI Too Broadly or Too Narrowly

TOO BROADLY	TOO NARROWLY
The work identified in one domain can't be completed within three or four waves of work over 12 to 15 months.	You're solving a niche challenge while leaving the root causes of problems untouched or not taking into account interrelated processes.
There are more than a dozen leaders with different goals who get to say what should happen next and there's no clear business owner with accountability.	The business leader in the target area doesn't feel ownership because the project won't move the needle, and you haven't involved leaders from across a specific value chain.
You need to redesign the whole data and tech architecture of the company to get any value.	You've created a solution that doesn't integrate with other upstream and downstream processes.

© HBR.org

In this article, we'll draw on our experience working with hundreds of clients, including some of the world's largest organizations, to describe what companies need to do to get AI to scale.


Step 1: Set the Strategy

It can be challenging to get the scope of AI initiatives just right. We advise CEOs to target areas of the business where AI will make a big difference in a reasonable period of time; it's relatively easy to find a sponsor, get stakeholders to buy in, and put together a team; and there are multiple interconnected activities and opportunities to reuse data and technology assets.

Potential impact. The chosen domains should be large enough to significantly improve either the company's bottom line or customer or employee experiences. One airline we advise determined that it had 10 main business domains fitting that description: cargo, crew, revenue management, e-commerce, customer service, airports, maintenance, network planning, operations, and talent. But it started with cargo, where it had identified a portfolio of AI initiatives that could be completed in about 18 weeks. The first would deliver some \$30 million in additional profit by enabling more accurate forecasting of cargo volumes and weight and increasing the use of shipping capacity.

In another case a telecom provider chose to redesign its process for managing customer value (which spans all the ways a company interacts with its customers), using AI to understand and address each customer's unique needs. That work quickly reduced the time it took to execute marketing campaigns by 75% and enabled the company to lower customer churn by three percentage points. The company expects those improvements to add \$70 million to its bottom line by the end of 2021.

Interconnected activities. Promising domains encompass a clear-cut set of business activities whose recalibration can solve systemic problems like chronic inefficiencies (such as lengthy loan approval times), high variability (rapidly fluctuating consumer demand), and routinely missed opportunities (difficulties getting products to customers). In many cases AI solutions may address the root causes of these problems, partly through the insights delivered and partly through organizational improvements.



CentreCentre/Patrick Fry Studio

The airline identified six closely intertwined cargo activities: negotiating rates, allocating space, booking reservations, documenting shipments, managing ground operations and delivery, and billing. Customer satisfaction and pricing were both dependent on factors such as the availability of space on short notice, the ability to track shipments in real time, and the speed of delivery. When the six activities were reconfigured so that they could feed data into an AI-supported platform, the company was able to significantly reduce systemic waste while greatly improving the customer experience—bolstering its margins and its reputation at the same time.

Sponsor and team. In a promising domain you can readily identify the following:

- an internal business champion responsible for the entire value chain involved (at the airline, it was the vice president of cargo)
- dedicated senior business staff (at the airline this included the senior director of cargo and two of his direct reports) who can fill the roles of “product owner” (the person responsible for solution delivery), translator (who bridges the analytics and business realms), and change lead (responsible for change management efforts)
- a team of AI practitioners, such as data science and engineering experts, designers, business analysts, and a scrum master (these practitioners may also be drawn from a central team in the organization)
- a cluster of frontline users or knowledge workers responsible for day-to-day activities (at the airline, they included 250 sales

and reservation agents across the Americas, Asia Pacific, and Europe)

Drafting employees from across the domain life cycle (regardless of where they formerly sat within the organization) and giving them accountability for the work builds engagement with an initiative and creates excitement and momentum. Those factors are crucial to getting employees to think beyond business as usual in devising solutions and help the project clear inevitable unexpected hurdles.

Reusable technology and data. It's also important to select domains where the data and technology components necessary to run the AI models can overlap. It's much easier when teams don't have to start from scratch every time and can reuse data or snippets of code that have already been prepared for AI. There will likely be a start-up investment for the first model or two created within a domain, but over time new projects can build off past ones, dramatically reducing development time and cost. The resources we're referring to here often include, on the data side, common libraries and metadata definitions, and on the technology side, machine learning scripts, application programming interfaces (APIs) that extract data from legacy systems, and data visualization capabilities.

Executive teams typically will identify about eight to 10 domains where AI can transform their business. Once they do, we recommend that they winnow the list down to one or two on the basis of feasibility and business value.

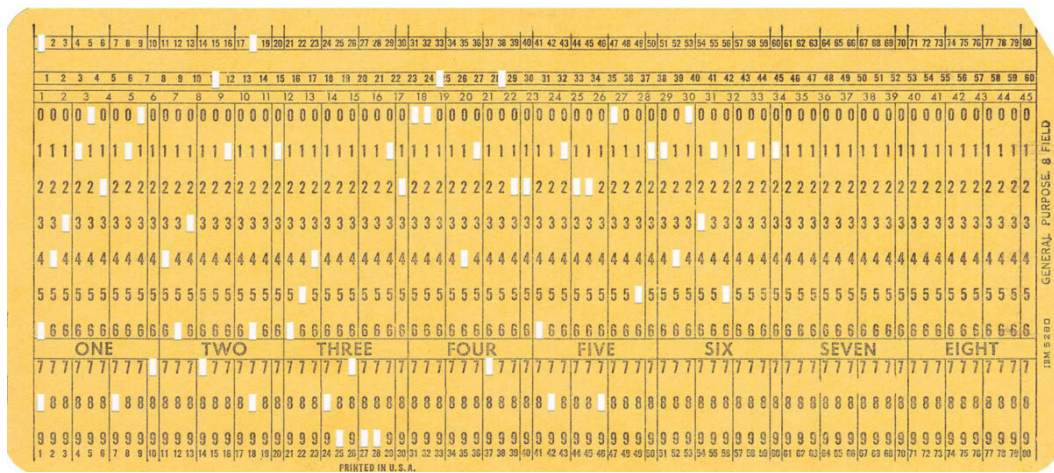
There will be a start-up investment for the first model or two created, but over time new projects can build off past ones, dramatically reducing development time and cost.

At the airline the CEO and his direct reports had held a series of strategy sessions over 12 weeks. They discussed how companies across different industries were innovating with AI, developed a vision for using AI to achieve a double-digit increase in operating profit within 15 months, prioritized which domains to start with, and committed the resources required to move forward. The executives each asked experts within their individual domains to identify what their areas could do differently to reach the profit goal and to assess the potential value and feasibility of their recommendations. In the cargo domain three senior business leaders, along with IT and finance staff, sketched out the opportunity to better fill available cargo space on planes, the expected returns for doing so, and the practicality of accomplishing this in terms of data availability, technology, talent, and so on.

Step 2: Structure the Team

The team responsible for AI initiatives within each domain should contain all the people necessary—from business, digital, analytics, and IT functions—to design, build, and support the new ways of working. To a great extent, once domain teams know their objective and are resourced, they will organize their work on their own, using agile practices. The role of management, beyond creating the teams, will be to ensure that any employees moved onto them from other parts of the business are fully integrated and to remove any organizational barriers that might impede teams' success.

In many cases we've studied, most of the team members needed were already working in the target domain, and leaders simply had to shift them onto the project and then bring in the necessary technical talent from other areas of the company. At the airline, sales, customer service, operations, and finance employees all were involved in the cargo domain transformation, and most of them had reported to the business function from the outset. AI experts, such as data scientists and data engineers, were assigned to the team from the company's AI center of excellence for the duration of the work and reported directly to the senior director in the cargo division, who was the product owner for the new AI.



CentreCentre/Patrick Fry Studio

In some cases companies will have to explicitly reassign people in other, nontechnical roles from various parts of the organization to the team. Consider an energy utility retailer that also sought to use AI to revamp customer value management, including which customers were targeted, which offers were sent to them and through which channels, and how new ideas were tested. The company had to formally move previously siloed marketing campaign experts from across channels and teams under one umbrella. Trying to coordinate their work across separate silos would have created delays and disconnects as requests for input and approvals moved from one department to another. It would also have forced the individuals involved to juggle two sets of obligations.

Often AI project teams can simply be single squads, in which the whole team carries out all the work by itself. But when the tasks are relatively broad in scope, requiring the work of more than a dozen people, a single team will be too unwieldy. In those situations it will make sense to divide the team into several squads, with one squad providing shared capabilities. The telecom company divided its new customer value team into four business squads—one focused on prepaid customers, one on postpaid customers, one on customer acquisition, and one on customer retention. It gave each a mission of either reducing churn or improving cross-sell by 20% by the end of the year. A fifth squad, data utility, with data engineers and developers, was created to support the other four by building technology and assets that could be reused by each one and by developing new AI-enabled analytics models.

Step 3: Reimagine Business as Usual

As we noted earlier, getting the most from AI requires reinventing business models, roles and responsibilities, and operational processes, using new ways of thinking and working. Typically, we find that companies are best served by applying first principles or design-thinking techniques and working backward from a key goal or challenge. For example, firms might envision what a five-star customer experience would look like and then explore in granular detail how they could achieve it.

At the airline the cargo team began by interviewing sales and reservation agents about how they allocated space on passenger planes and decided whether to accept or reject shipment requests. How did agents check on cargo space availability? What other information did they rely on, and how did they weigh the different pieces of information they collected? What concerns did they have when making decisions?

Teams should consider the potential impact that AI initiatives will have on upstream and downstream processes and implement measures to address it.

The team found that the legacy approach was plagued by inaccurate forecasts and guesswork by agents trying to estimate potential cancellations. (With cargo bookings, unlike passenger reservations, there's no penalty for canceling, so it's not unusual for a plane to look fully booked but leave with empty cargo space because of a no-show.) Cargo booking agents were also apprehensive about the impact on customer satisfaction if space was overbooked. To avoid conflicts, agents often waited until the day of the flight to book cargo space for their customers, resulting in suboptimal use of capacity and missed opportunities.

Having identified and understood the issues with the existing processes, the team then mapped out what an ideal process might look like, including the information that agents would need to

determine whether to book, how much they could safely overbook and how far in advance, and how roles would be different. It then spent a few weeks developing a prototype of an AI-enabled dashboard that would provide the necessary information to agents, working in iterative sprints with them to incorporate input from the forecasting models, which were being developed in parallel. The team tested the dashboard with agents for 12 routes representative of the company's global network of 1,500. It compared how cargo utilization and profits differed on routes for agents who followed system recommendations and for a control group who used traditional processes. To build trust in the new system, executives eliminated any repercussions agents might normally face if a flight couldn't accommodate a reservation.

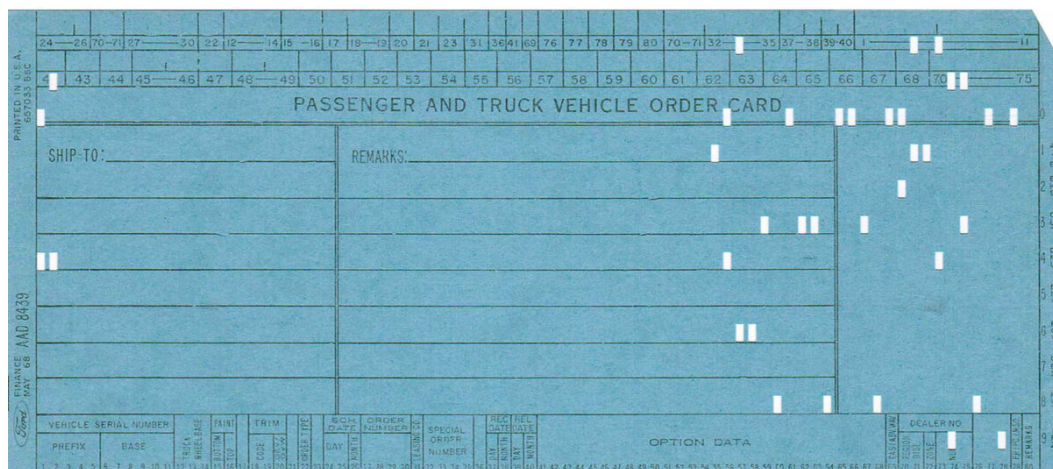
All agents now have access to intuitive dashboards that visually illustrate which flights are underutilizing space. They can view at a glance data on how cargo shipments for recent flights produced revenue. Integrated feedback loops enable the AI systems to continually learn from the agents as they decide whether to accept a cargo request, drawing on their expertise on shipment size and weight balance issues and their knowledge of changes in customers' supply chains, trade routes, and other factors. These new tools provide agents with information that gives them the confidence to sell cargo space well ahead of departure dates.

Step 4: Adapt for Organizational and Technological Change

In most cases significant organizational changes, such as adopting interdisciplinary collaboration and agile mindsets, will be required to support the new AI-based processes and models. In fact, our research shows that the companies getting the highest returns on AI are more likely to enact effective change management practices, such as having leaders model desired behaviors, and that such efforts work best when facilitated by CEOs and top executives.

Take the energy utility retailer again. It invested in reskilling employees so that they could effectively work together in the new context and take on new leadership responsibilities; realigned AI

project team members' goals and incentives with their new responsibilities; and backfilled responsibilities in the departments the team members had to leave.



Patrick Fry's series *Print Punch* explores the aesthetics of the punch card and reflects on a time when data was physically tangible. CentreCentre/Patrick Fry Studio

While companies will need to update their tech to support AI, they won't need to do major surgery on their IT infrastructure or data architecture before they begin. Rather, we advise companies to focus on technology that will enable and accelerate AI development and then triage additional investments according to teams' priorities. Cloud-based data platforms and the use of APIs, microservices, and other modern dev-ops practices, for example, can help companies develop new business capabilities two to three times faster.

The telecom provider established a cloud-based platform for raw data from existing transaction and customer service systems so that it could be used more easily by data engineers and data scientists than data from the old warehouse system could. The company also implemented a new analytics workbench, which helped the data scientists train and deploy new models faster, and tools that streamlined data collection, analysis, and model building for its AI-driven customer-value-management system. Those moves allowed it to begin using unstructured data, apply more complex approaches, and work more efficiently.

When prioritizing additional technology investments, teams should map out the capabilities, data, and resources (such as robotics, biometrics, and sensors and connectivity platforms)

they will require and when, and then chip away at each piece as needed. In designing its customer-value- management system, the telecom provider's team realized it would need new technology that automated outbound direct messaging and gave salespeople real-time guidance about the next conversation to have with customers.

Teams should also consider the potential impact that AI initiatives will have on upstream and downstream processes and implement measures to address it. For example, at the airline the AI team developed a reporting tool for managers overseeing the loading and unloading of cargo so that they could effectively support the higher volumes produced by the new sales and reservation process.

A Domino Effect

Once AI development matures within an initial domain and organizations have gotten into a rhythm for reimagining parts of the business, they're ready to expand. The tech foundation they've built and the skills they've learned—for example, how to successfully break down silos, make decisions that used to take weeks in hours, and create more data-driven teams—will help accelerate their efforts in new domains.

At this point companies can pursue multiple domains in parallel. Again, the idea is to build off past work. This might lead companies to prioritize domains that have data and skills in common, such as supply chain and logistics. Or they might pursue the same domain in other business units. The energy utility retailer estimates that nearly 80% of the work done on improving customer value management in one product division (which led to record growth in just a few months, including a 12% increase in customer profit and a 20% increase in customer retention) can be reused in several other business units and accelerate their growth as well.

The companies profiled in this article are all still in the earlier stages of their full AI transformations, but they're on the threshold of a new era. They've gained a taste of what's possible,

and their bold choices have yielded significant returns within the domains they've targeted and new capabilities that discrete use cases couldn't deliver. These companies have created a playbook of methodologies and protocols they can turn to again. As they move on to other domains, their pace will quicken, their AI capabilities will rapidly compound, and they'll find that the future they imagined is actually closer than it once appeared.

A version of this article appeared in the May–June 2021 issue of *Harvard Business Review*.

TF

Tim Fontaine is a senior partner in McKinsey's Sydney office and leads QuantumBlack, an advanced analytics firm owned by McKinsey, in Australia.

Brian McCarthy is a partner in McKinsey's Atlanta office and co-leads the analytics transformation and knowledge development agendas for McKinsey Digital.

TS

Tamim Saleh is a senior partner in McKinsey's London office and is a global analytics leader for McKinsey Digital.