

Operations Practice

Toward smart production: Machine intelligence in business operations

A detailed study of machine intelligence in industrial and manufacturing operations reveals the surprisingly different paths companies can take. But a group of leaders shares similar characteristics.

This article is a collaborative effort by Duane S. Boning, Vijay D'Silva, Pete Kimball, Bruce Lawler, Retsef Levi, and Ingrid Millan, representing views from McKinsey's Operations Practice and the Massachusetts Institute of Technology's Machine Intelligence for Manufacturing and Operations program.



As the superintendent of Vistra Corp's Luminant Martin Lake Power Plant, Wayne Brown is an expert in power generation. Vistra is the largest competitive power producer in the United States, operating power plants in 12 states and producing more than 39 gigawatts of electricity—enough to power nearly 20 million homes. The company has been on a journey to drive operational excellence across its generation portfolio. Launched in 2016, its Operational Performance Initiative has driven a step-change improvement in the efficiency of its assets, generating hundreds of millions in incremental EBITDA along the way.

To maintain and improve its position, Vistra is continually looking for the tools, technologies, and approaches that will help it achieve the next level of performance. Most recently, the company has turned to digital and analytics, including machine intelligence (MI). At the Martin Lake plant, Brown and his team use heat rate as their primary efficiency metric. It measures how much electricity is generated for each ton of fuel consumed by the plant. Optimizing the heat rate is extremely demanding, however. Historically, operators had to continuously monitor and tune hundreds of set points, including steam temperatures and pressures, oxygen levels, and pump and fan speeds. It's a lot for any operator to get right 100 percent of the time,

especially when ambient conditions and the availability of different pieces of equipment are constantly changing.

Now those operators have a smart new ally in the quest for efficiency. Brown and a group of data scientists have built a heat-rate optimizer for Martin Lake. It automatically analyzes hundreds of inputs in real-time using AI. Every 30 minutes, the system generates recommendations for tweaks and adjustments that operators can apply to maximize the performance of the plant. The heat-rate optimizer has been running for roughly nine months, and the plant has been consistently running 1 percent more efficiently. That might not seem like much, but it's worth millions of dollars per year, drives down greenhouse-gas emissions, and frees worker time to focus on other tasks. Scaled across Vistra's large and diverse fleet of power plants and applied to dozens of other operational opportunities, the potential of MI in the business is huge.

The emergence of machine intelligence

Vistra's experience in applying advanced production technologies is playing out at companies across sectors. New research by McKinsey and the Massachusetts Institute of Technology's (MIT's) Machine Intelligence for Manufacturing and

Leading companies are using MI technologies to move the needle on a broad set of performance indicators, achieving three or four times the impact of average players.

Operations (MIMO) program shows that leading companies are using MI technologies to move the needle on a broad set of performance indicators, achieving three or four times the impact of average players.

The work, which was conducted in the first half of 2021, focused on where, when, and how these technologies are being deployed in organizations and about their impact on operational performance (see sidebar “About the research”).

We analyzed our findings to identify the relationships between actions taken and investments made, and tangible and sustainable outcomes. For instance,

the leaders in our survey claimed more than twice the impact in half the time compared with those early in their journey.

How are they achieving these results? Are they picking different use cases? Do they have a unique approach? As we probed deeper in our interviews and dug through the data in our survey, the real drivers of outcomes and performance were more nuanced.

From insight to autonomy

Our research looked at five different ways that companies are using data and analytics to improve the speed, agility, and performance of operational decision making. This evolution of digital maturity begins with simple tools, such as dashboards to aid human decision making, and ends with true MI, machines that can adjust their own performance autonomously based on historical and real-time data.

The journey from digital dashboards to full-fledged MI depends on multiple factors. The first of these is reliable and available data¹: 92 percent of the companies classified as leaders in our survey had a process in place to track incomplete and inaccurate data. As companies move toward MI, the opportunity and stakes get higher—not only in terms of the required skills and investments, but also in the potential for error. Many companies will only leverage machine learning or AI to the extent that its results are intuitive and transparent and, in some industries, explainable to external regulators. The importance of this transparency came through in our survey: 67 percent of leaders track traceability of outcomes, compared with only 14 percent of the rest.

But the impact from autonomy (taking the human out of the loop) in well-defined use cases can be substantial, as we found at many companies we talked to. Take, for example, forecasting analytics, where leaders frequently use far more sophisticated approaches. These companies report an average service-level and demand-accuracy improvement of

About the research

We reached out to the combined networks of the Massachusetts Institute of Technology (MIT) and McKinsey to study 100 North American companies in sectors from automotive to mining; reviewed literature produced by the users of machine intelligence (MI) technologies, vendors, and other industry participants; and conducted in-depth off-the-record interviews with dozens of industry practitioners. Our survey and research covered nine categories, 17 desired business outcomes, and how 17 different use cases had an impact on 21 KPIs. Given the complexity of the data, we even applied several different MI approaches in our own analysis.

Respondents were refreshingly candid in answering questions about their use of a wide range of digital, data analytics, and MI approaches. We asked companies about the goals of their digital investments, the metrics they tracked, and the progress they have made so far. The interviews allowed experts from McKinsey and MIT to really probe the challenges, opportunities, and key success factors in digital operations today.

¹ Tadhg Nagle, Tom Redman, and David Sammon, “Assessing data quality: A managerial call to action,” *Business Horizons*, May–June 2020, Volume 63, Number 3.

13 percent, whereas those earlier in their journeys typically use dashboards and see average improvements of only 3 percent.

Leading the pack: What makes the difference?

The real goal of this work was to look beyond the top-level numbers to explore the underlying drivers of success in the use of these emerging technologies. Are some companies just naturally better at applying advanced digital technologies, or do they adopt strategies and approaches that others could replicate?

To understand those drivers, we looked at three groups of attributes for each company in our survey. First, we evaluated the *results* they have achieved through their digital programs, measured by factors including the number of functions applying these technologies, the number of KPIs improved, and the size of those improvements. Second, we looked at *deployment*, measured by factors including expenditure, the number of years the companies have been working with these technologies, the maturity of their implementations and analytical

methods, and the number of partner organizations involved. What's interesting is that leaders continually invest more in their deployments, with budgets 30 to 60 percent larger (as a percent of total) across manufacturing, customer service, product design, and supply chain. In fact, leaders' budgets are expected to increase by another 10 to 15 percent next year, while those of other companies stay largely stagnant.

Third, we looked at a set of *enablers*, including the maturity of the organization's digital strategy, its data infrastructure, and the presence of centers of excellence or structured capability-building programs. Some of these enablers may come as no surprise: leaders have clearly defined strategies and utilize centers of excellence at twice the rate of others. Leaders are also five times more likely than the rest to train frontline employees on digitization and the Internet of Things (IoT). As they become more mature, these companies tend to cast a wider net, looking to improve business outcomes in everything from equipment efficiency to talent acquisition, instead of being overly focused on a handful of use cases (see sidebar "How one company increased the ambition of its use cases").

How one company increased the ambition of its use cases

Mobility-technology company Magna International's application of machine intelligence on electronic clutches provides a useful example of the type of impact achieved. The automotive supplier's plant had no shortage of data: the facility measured and recorded virtually all aspects of production. During assembly, the dimensions of components were recorded, along with processes such as press forces and press distances. And all physical, process, and test data for every clutch assembly were stored in a relational database with data on hundreds of thousands of clutches.

Magna started using machine learning to study its production data and determine which of the thousands of variables had the most impact on manufacturing quality product. The company learned that adjusting one parameter on one process step could mitigate all that variability. Operators would adjust this parameter periodically, but not knowing its relative impact on all the other possible adjustments and settings on the assembly line, it was difficult to know when and by how much to adjust it, and the impact of changes on scrap rates was hit or miss.

As a small internal team of experts and data scientists worked through the problem, they were able to train a computer model to become a process expert and focus its full attention on determining how to set this one parameter. Today, the model learns how to adapt to the continual small changes in the 41 variables that matter most, especially within the last 25 to 30 units produced. The process is now adjusted *every 20 minutes*, and the number of clutches needing to be reworked or scrapped has fallen by 80 percent.

Using these sets of attributes, we applied a clustering algorithm² to our data to identify four groups of companies with distinct characteristics (Exhibit 1).

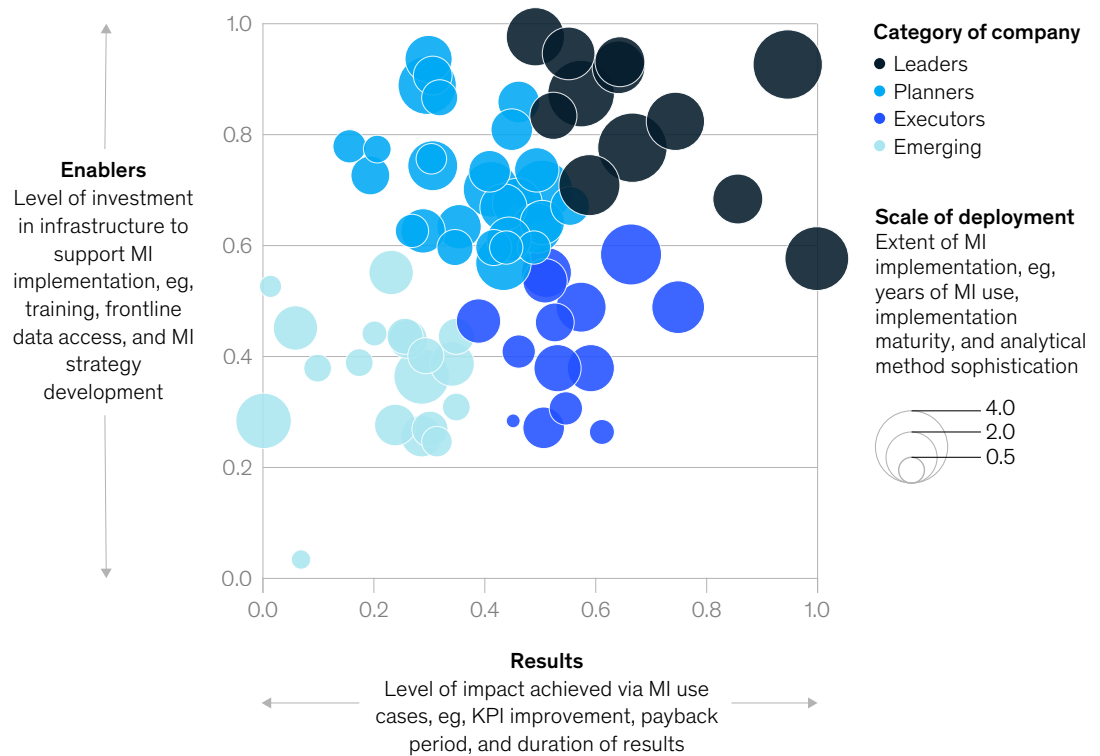
At the top right of the chart are the *leaders*, companies that have captured the largest gains from their MI efforts and which have the highest level of supporting infrastructure and capabilities.

At the bottom left is a group of *emerging* companies, which have the lowest level of enabler maturity and saw the smallest gains. Some of these firms have seen moderate success with select use cases. Many have just started on their journey, but they don't yet have the scale or infrastructure of other, more mature MI practitioners. Some emerging companies struggle with identifying the right places to invest. For instance, an emerging company we

Exhibit 1

We identified four distinct groups of companies that vary in both their approach to machine intelligence and the impact of their efforts.

Machine-intelligence (MI) impact matrix, variables are normalized to a 0 to 1 scale



Note: Direct measures of MI benefit weighted 50% more than other variables.
Source: MIT Machine Intelligence for Manufacturing and Operations (MIMO) program; McKinsey Machine Intelligence Survey

² We used k-means clustering, in which companies are grouped based on distance to the nearest cluster centroid, which is the average of all points in the cluster.

As companies move toward MI, the opportunity and stakes get higher—not only in terms of the required skills and investments, but also in the potential for error.

spoke to went after predictive maintenance due to its seemingly broad use (not necessarily because it was a business priority). The effort had little to no impact. But, when the company focused more on its biggest business priorities (such as improving service levels), it hit a home run.

The middle ground is more nuanced. In the center of the chart, we identified two distinct clusters. One group of companies, which we call *planners*, exhibits a level of maturity in enablers that is on a par with companies in the leader group. Planners have yet to achieve the same results as their leader counterparts, but some are just on the cusp of being leaders. We spoke to several planners that had very successful implementations and are well on the way to deploying use cases on the scale of leaders. But we also observed planners that have been unable to achieve the impact of our leader group, even though they have invested significant time and resources into digital projects.

Another group, which we call *executors*, has already achieved significant gains, despite building less infrastructure than the leaders or planners and even less than some companies in the emerging group. Executors typically find a few high-impact use cases, partner with others to create a solution, and then start on a path to implement these point solutions as broadly as they can. One executor we interviewed shared that a robust governance structure that quickly identified and then broadly deployed promising use cases played a critical role in its success.

This pattern reflects how MI has been industrialized in recent years. A decade ago, any company wishing to apply MI to its operations had little choice but to build the necessary infrastructure and systems in-house. Today, organizations have the option to buy tools and expertise from a huge ecosystem of vendors. That makes it possible for executors to achieve significant gains in specific use cases without doing as much groundwork up front.

Meanwhile, the large number of companies in the planner group indicates that plenty of organizations have yet to crack the code on applying MI technologies at scale in the use cases that really count. Some of them may still be struggling to escape from the pilot purgatory that our McKinsey colleagues first described back in 2018.

That doesn't mean planners are taking the wrong route, however. The companies in our study that have achieved the best results from their MI efforts have also reached high levels of maturity in infrastructure and capabilities. Our research suggests that any company with aspirations to reach the frontiers of MI-enabled performance will need to put the right enablers in place.

Emerging companies, meanwhile, have a strategic choice to make. Do they follow the planner route, investing early in enablers that will support their long-term journey? Or do they pursue fast gains, buying technology and know-how from partners and building only the necessary enablers to support specific use cases?

What's the size of the prize for leaders?

It turns out it's substantial. In our sample, we isolated the top performers and compared them with the bottom 50 percent across several dimensions

The dozen highest-performing companies in our survey achieved significantly larger improvements than the rest in 20 of the 21 KPIs we evaluated (Exhibit 2). In 11 of those KPIs, the average improvement achieved by leaders was 10 percent or more. And their overall average improvement, at 9.5 percent, was almost three times the 3.5 percent achieved by lower performers.

The results of even select actions can be substantial. For instance, Wayfair, an e-commerce company,

was early to move its data to the cloud and use machine learning in multiple functions. During the COVID-19 crisis, as consumer demand changed rapidly across regions, the company has been able to optimize container ship logistics, continually adjusting what goods were sent to which ports. With shipping surcharges becoming more prevalent, Wayfair saved about 7.5 percent in inbound logistics costs as a result.

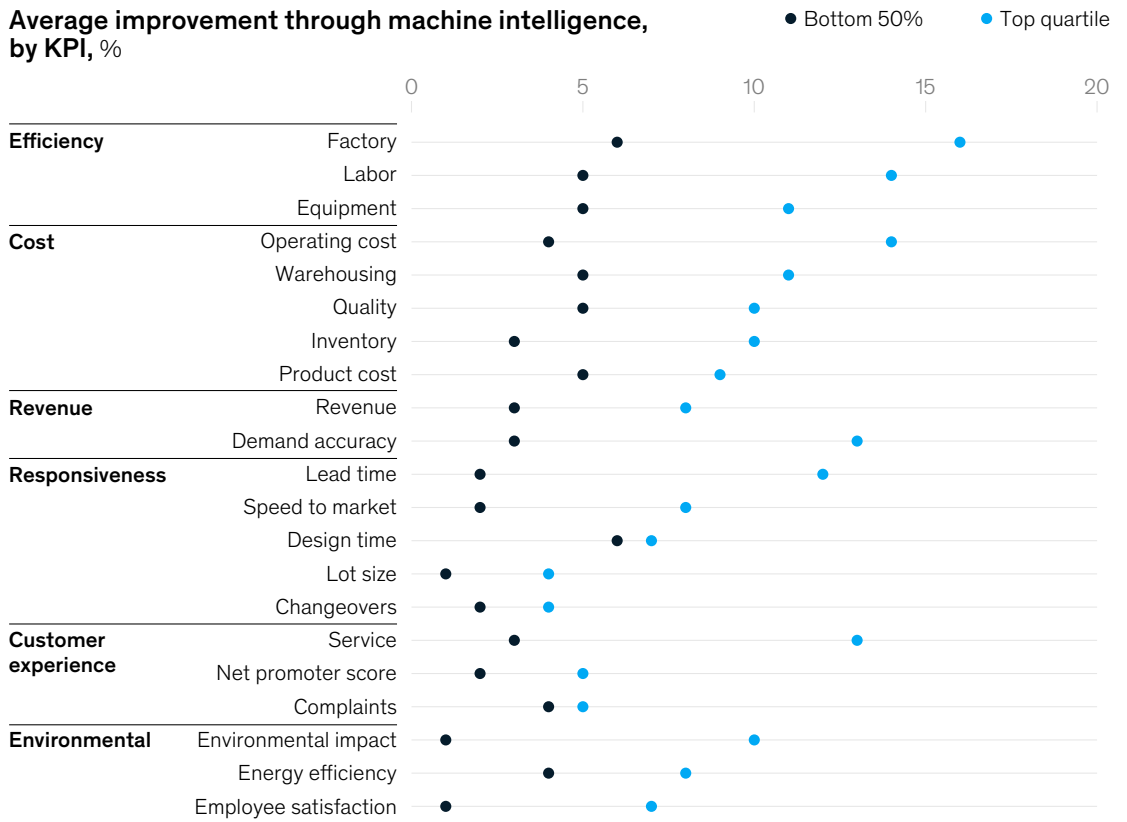
What leaders do well

Whether it is a planner, an executor, or an emerging organization today, any company with ambitions to gain from advanced digital technologies has the opportunity to learn from best-practice approaches. Our research identified five areas where the top-performing companies stand out.

Exhibit 2

Across a broad range of metrics, machine-intelligence leaders achieve triple the improvement of other companies.

Average improvement through machine intelligence, by KPI, %



Note: All variables are normalized to a 0 to 1 scale.
Source: MIT Machine Intelligence for Manufacturing and Operations (MIMO) program; McKinsey Machine Intelligence Survey


Governance: Processes and capabilities to rapidly scale

MI is a strategic priority for leading companies. Many of these organizations have built dedicated centers of excellence to support their MI efforts, either within business units or as a centralized function to support the entire organization (Exhibit 3), ensure standards, and accelerate deployment. One healthcare firm also used a well-documented governance process to deploy multiple use cases at one plant, which it then rolled out across its network, resulting in a revenue lift.

Leading companies also have strong governance in place to keep their digital programs on track. For example, leaders are much more likely than lower-performing companies to have a defined process for the assessment and implementation of digital innovation. They are also more likely to follow that process regularly and to update it continually as the organization learns more about what works and what doesn't (Exhibit 4).

Exhibit 3


More than 90 percent of machine-intelligence leaders have a dedicated center of excellence.



Source: MIT Machine Intelligence for Manufacturing and Operations (MIMO) program; McKinsey Machine Intelligence Survey

Exhibit 4

Leading companies keep a close eye on the performance, reliability, and progress of machine-intelligence implementation.



Source: MIT Machine Intelligence for Manufacturing and Operations (MIMO) program; McKinsey Machine Intelligence Survey

Deployment: Ambitious vision and complex use cases

Leading organizations apply MI more widely and use more sophisticated approaches. For example, every leader in our research is using MI in three areas: forecasting, maintenance optimization, and logistics and transportation (Exhibit 5). Among lower performers, only half are addressing the first two of those, with fewer than a third addressing the last. And of the 17 use cases we identified in our survey, more than three-quarters of leading companies are addressing at least 14. Furthermore, while simple dashboards are the most common digital tool used by lower-performing companies, leaders are much more likely to adopt more advanced approaches—such as predictive and prescriptive methods—or MI. Many choose to move up this ladder of sophistication once they have more of their enablers in place.

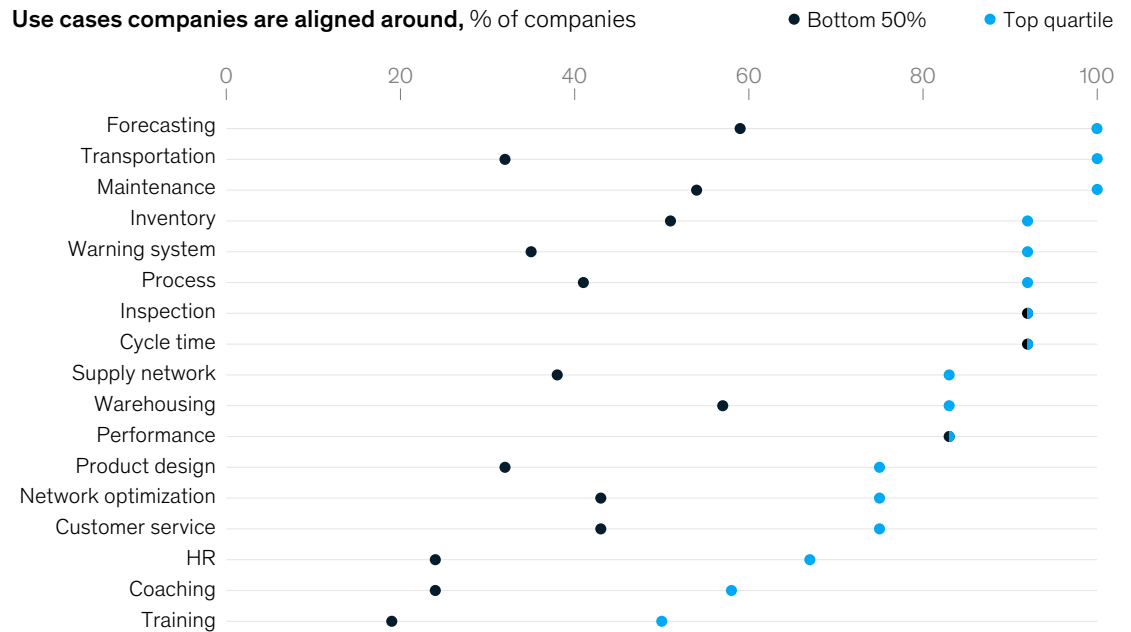
It seems leaders have somewhat paved the way here for others who are earlier in their journeys to follow. There's clear consensus among leaders in terms of not just sophistication but also the top use cases to focus on.

Some companies face many challenges when scaling up their programs, including a lack of historical data on some products and the need to comply with strict regulatory requirements. In some cases, the time required to verify and document new digital tools can be four times that required for their development. Despite these barriers, even heavily regulated companies are seeing significant returns from MI investments.

One notable recent success has been the application of machine vision to product-quality assurance. Biopharmaceutical company Amgen found that

Exhibit 5

Leading companies are largely aligned around the top use cases.



Source: MIT Machine Intelligence for Manufacturing and Operations (MIMO) program; McKinsey Machine Intelligence Survey

visual-inspection-system operations posed great opportunities to automate and used AI technologies to cut batch-inspection time from hours and days down to minutes.

Partnering: A broad ecosystem of partners

The development and implementation of advanced digital tools has become a team sport. Most companies in our research say that they partner with other organizations in their MI projects, with academia, start-ups, existing technology vendors, and external consultants as the most common choices. The reasons vary. According to one chief technology officer, “Manufacturing lines are inherently unique; we don’t have consistently well-structured data as you might get from standard enterprise-resource-planning systems. We found general-purpose AI solutions either lacked the capability or cost too much to customize for our application. Start-ups and academia have better domain expertise. They’re trying to solve a particular problem for us.”

Leading organizations are more likely than others to pursue such partnerships, and they collaborate with a wider range of organization types. More than

80 percent of leaders have partnerships with academic groups and start-ups, for example, while less than a third of other organizations have these relationships (Exhibit 6).

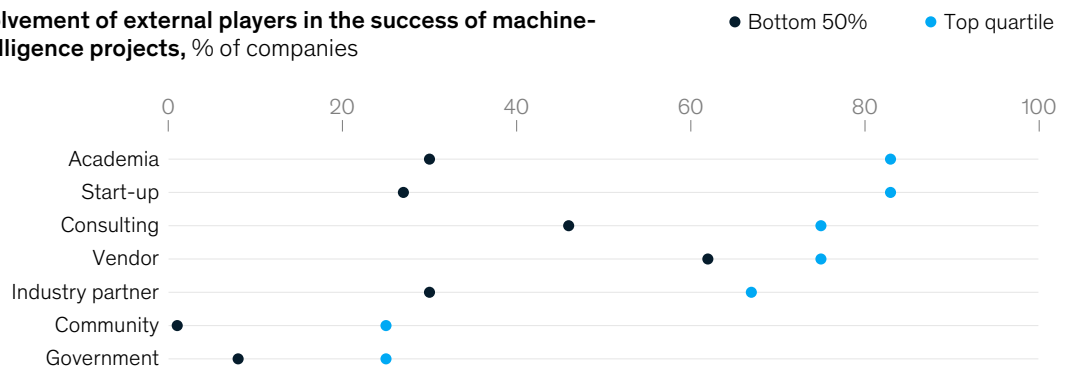
Colgate-Palmolive and PepsiCo-owned Frito-Lay are both working with MI-systems vendor Augury to deploy AI-driven machine-health diagnostics on their production lines. Using the approach to detect and prevent unscheduled stoppages has resulted in significant savings for both companies, often from a single incident; in one case, this prevented an eight-day outage.

Analog Devices’ efforts on real-time quality control demonstrate why leaders may turn to academia and start-ups for deep technical expertise. Faults in semiconductor fabrication processes are extraordinarily costly. Given enough examples of faulty process runs that result in “bad” parts, MI models can be trained to identify those faults in real time. However, semiconductors have very high yield, meaning there are very few faulty runs or bad wafers to train an MI model. Analog Devices manufactures a broad variety of semiconductors, making it even less likely that there would be

Exhibit 6

Leading companies have a broad ecosystem of partners for their machine-intelligence projects.

Involvement of external players in the success of machine-intelligence projects, % of companies



Source: MIT Machine Intelligence for Manufacturing and Operations (MIMO) program; McKinsey Machine Intelligence Survey

enough failure examples for any given recipe or tool to train a conventional model. Therefore, Analog Devices needed something that could learn to identify bad process runs using only the sensor data from producing lots of “good” wafers. With a novel MI method developed in collaboration with MIT, Analog Devices was able to identify which production runs and tools might have a fault. This identification meant that experienced process engineers only had to review 5 percent of the process data than they did before, resulting in better focus and a 20-fold reduction in workload.

People: Employee empowerment

At leading companies, advanced digital approaches are not the concern of only a few specialists. These organizations believe that MI is important for stakeholders across the business, and they take steps to ensure as many of them as possible have the skills and resources they need to make use of these techniques. More than half of them train their frontline personnel in MI fundamentals, for example, compared with only 4 percent of other companies.

McDonald’s, a global and highly decentralized quick-service restaurant, used MI to improve a wide range of operational tasks, from predicting customer response to promotions to forecasting real-time footfall in its restaurants. The company adopted a hybrid approach: a corporate center of excellence tests and develops new modeling approaches, which it packages into easy-to-use tools that are made available for regional operating organizations and restaurants. Ultimately, in-market employees use the tools. This approach helps team members in the field learn the importance of good data, develop a process understanding, and hone problem-identification skills.

Data execution: Data they can use

Leaders make data accessible. All of the leader organizations give frontline staff access to data, compared with just 62 percent of the rest. Each one of the leading companies also acquires data from customers and suppliers, and 89 percent of them share their own data back with these groups. Among other companies, the equivalent figures are 73 percent and 46 percent, respectively. Leading

companies are almost twice as likely as others to enable remote access to data and to store a significant fraction of their data in the cloud.

The democratization of data is a critical aspect at many companies. Automotive supplier Cooper Standard, for example, requires teams to address data strategy early in the organization’s agile development process for new MI applications. That helps to ensure that all use cases are built on robust, well-managed data that provide a foundation for future developments or new use cases.

Vistaprint, a design, digital, and print company serving small businesses, also puts data strategy at the forefront of its analytics efforts. The company’s MI organization reports into the chief data officer, giving it a great starting point with a critical enabler to success. One of its more ambitious and successful applications, on-time-delivery analytics, required tie-ins across multiple functions (including carrier managers, the website team, the product group, and the supply chain) with an aligned data strategy. The analytics platform reduced late deliveries across the business by more than 50 percent and is the backbone for introducing eCommerce-leading shipping options like next-day delivery.

Where next?

The race to leverage data and analytics could be won with multiple coordinated actions rather than any single bold move. All four segments—leaders, planners, executors, and emerging companies—operate in a dynamic space where the bar is rising and the number of machine-learning use cases is set to keep increasing, quickly embedding themselves into business as usual.

For example, planners often have strong people skills and considerable data-execution expertise, while executors get strong results from a laser-sharp focus on the most promising opportunities in their business. Depending on its starting point, each company’s path forward will be different, as it seeks to capitalize on its current strengths while addressing the remaining weaknesses in its MI capabilities.

At leading companies, advanced digital approaches are not the concern of only a few specialists. These organizations believe that MI is important for stakeholders across the business.

Since not everyone can strive to be a leader immediately, they can instead strive to move to the next better state. Across executors, planners, and emerging companies, we see three leverage points: opportunity focus, governance, and budget.

Executors in our study are getting great results and are focusing on the right problems. However, to be more like leaders, they can invest more and look to their governance and budget. Planners have put many of the right processes in place, and can focus more on specific opportunities and driving results across the company at pace through better governance. Emerging firms, meanwhile, can focus on selecting the right early opportunities and obtaining results to become more like executors, quickly stacking wins to build momentum and the case for further investment.

In our interviews, we found that many of the leading companies reached their level of excellence by following four steps:

1. **Identify the organization's North Star** and conduct a visioning exercise that describes an ideal version of its operations in three or four years' time. This will require looking past legacy infrastructure and skill constraints. It will probably involve borrowing best practices from other sectors too. For instance, Vistaprint organized its early efforts around on-time delivery as a paramount objective.
2. **Conduct an honest assessment** of the organization's starting point across dimensions. Companies often overindex on a few dimensions during early successes but find the momentum difficult to maintain without an operating-model change. Many companies find that an objective, external diagnostic is the pragmatic choice in understanding where they really stand versus their competitors—and, more importantly, versus their potential.
3. **Scope out a rough transition plan**, accounting for barriers to change, critical infrastructure (such as data migration to the cloud), and assigning realistic medium-term targets. Most leaders in our survey started by using data to make decisions using tools simpler than MI. As they built maturity and familiarity with their data, they could move to more advanced techniques.
4. **Find a handful of use cases to gain early momentum** while at the same time scoping out (and potentially self-funding) the necessary changes in infrastructure, talent, and related supports. The healthcare company mentioned earlier invested in use cases that had applicability across its production network, helping it build momentum by engaging a broader user base.

Despite recent advances in MI, it is increasingly clear that the full scale of the opportunity from the embedded use of data and analytics in business operations is just beginning to unfold and is set to continue as more use cases evolve from simple dashboards to greater levels of autonomy. As companies understand better how to direct their MI efforts to address their most pressing determinants of value, the gap between leaders and the rest may start to widen: leaders continually invest more in MI and expect their budgets to increase by another 10 to 15 percent next year, while the rest may stay largely stagnant.

At Vistra, eighteen months into the transformation that followed the work that Wayne Brown led, the company identified even more impact potential and has started scaling analytics models across the business, continuing to improve operating efficiency. The next frontier is Vistra's renewables portfolio, including solar and batteries, which could yet turn out to be the biggest impact of MI on the company.

Vijay D'Silva is a senior partner in McKinsey's New York office; **Pete Kimball** is an associate partner in the Chicago office; and **Ingrid Millan** is a partner in the Washington, DC, office. **Duane S. Boning** is the Clarence J. LeBel professor in electrical engineering and professor of electrical engineering and computer science at the Massachusetts Institute of Technology (MIT); **Bruce Lawler** is the managing director of MIT Machine Intelligence for Manufacturing and Operations; and **Retsef Levi** is the J. Spencer Standish (1945) professor of operations management at the MIT Sloan School of Management.

The authors wish to thank Erez Kaminski, Aaron Wang, MIT's Leaders for Global Operations Program (LGO), and all of the companies that supported this research and article.

Designed by McKinsey Global Publishing
Copyright © 2022 McKinsey & Company. All rights reserved.