

AI, Labor, and Economy Case Studies

# Compendium Synthesis



**PARTNERSHIP ON AI**

# Preface

---

The AI, Labor, and Economy Case Studies Compendium is a work product of the Partnership on AI's "AI, Labor, and the Economy" (AILE) Working Group, formed through a collaborative process of research scoping and iteration. Though this work product reflects the inputs of many members of PAI, it should not be read as representing the views of any particular organization or individual within this Working Group, or an entity within PAI at-large.

The Partnership on AI (PAI) is a 501(c)3 nonprofit organization established to study and formulate best practices on AI technologies, to advance the public's understanding of AI, and to serve as an open platform for discussion and engagement about AI and its influences on people and society.

One of PAI's significant program lines is a series of Working Groups reflective of its [Thematic Pillars](#), which are a driving force in research and best practice generation. The Partnership's activities are deliberately determined by its coalition of [over 80 members](#), including civil society groups, corporate developers and users of AI, and numerous academic artificial intelligence research labs, but from the outset of the organization, the intention has been to create a place for open critique and reflection. Crucially, the Partnership is an independent organization; though supported and shaped by our Partner community, the Partnership is ultimately more than the sum of its parts and will make independent determinations to which its Partners will collectively contribute, but never individually dictate. PAI provides staff administrative and project management support to Working Groups, oversees project selection, and provides financial resources or direct research support to projects as needed.

# Key learnings and synthesis document

## Objectives and scope

The impact of artificial intelligence (AI) on the economy, labor, and society has long been a topic of debate — particularly in the last decade — amongst policymakers, business leaders, and the broader public. Estimates of its current and imminent impact have varied widely, often reaching contradictory conclusions. One major question for the public and policymakers has been AI's impact on the workforce, both in the changing nature of work and net job loss or creation. Another major question is whether AI can enable an acceleration of productivity growth, which has stagnated in many economies. At the same time, a salient question for managers, strategists, and economic analysts has been whether large investments in artificial intelligence and machine learning are warranted: Can the promises of AI be realized, and if so, what are their potential impacts on the various stakeholders involved?

To help elucidate these various areas of uncertainty, the Partnership on AI (PAI)'s Working Group on "AI, Labor, and the Economy" (AILE)<sup>1</sup> conducted a series of case studies across three geographies and industries, using interviews with management as an entry point to investigate the productivity impacts and labor implications of AI implementation. We, PAI's aforementioned AILE Working Group, publish these case studies detailing varied applications of AI with the following objectives:

- **To tell a detailed range of stories about the contexts under which organizations deploy AI implementations, and how they manage these**
  - *How do country and regulatory complexity affect decisions to proceed with AI implementations, and in what forms?*
  - *How do companies build AI into their culture and processes, and reconcile its implementations with existing programs?*
  - *What common themes exist regarding enablers and impediments of AI implementations, and how can other organizations and audiences learn from these themes?*
- **To consider AI's impacts at the level of an organization, especially on workers, business results, and processes, in addition to the macroeconomic view so often discussed<sup>2</sup>**
  - *What specific challenges do organizations face when integrating AI and what steps do they take to address them?*
  - *Where do we see evidence of organizational changes from AI in individual organizations that may foreshadow broader economic impacts?*

<sup>1</sup> See Acknowledgements for more information

<sup>2</sup> Researchers have argued for the need for "more systematic collection of the use of these technologies at the firm level." The case study project intends to provide quantitative and qualitative data at the firm level. For more, see "AI, Labor, Productivity and the Need for Firm-Level Data," Manav Raj and Robert Seamans, April 2018.

- **To provoke and inform conversation about AI's productivity and workforce impacts**
  - Are there examples of AI systems enhancing productivity?
  - Whether AI enhances productivity or not, what are the implications for workers?
  - What factors can engender worker trust and broader organizational buy-in of AI systems and solutions?

Through this synthesis document and the accompanying case study materials, we aim to ground the conversations around workforce impact and productivity in real-world examples of AI implementations. We recognize the methodological and scope-related limitations of this body of work, which we hope will inform ongoing conversation and provoke further inquiry.

This *Key Learnings and Synthesis Document* focuses on key observations and broader trends across three distinct case studies. It introduces the three subject organizations along with their applications of AI and presents our key observations across broader themes. For further detail, please see the individual case studies, which present a more thorough perspective into the AI implementations at these organizations and the resulting impacts on processes, business results, and the workforce. Terms and techniques used, as well as methodology for this report are also presented below.

## Case studies: subject diversity and common motifs

We present a rich group of stories from three organizations, spanning different industries, geographies, sizes, and use cases of AI and machine learning (ML) techniques.<sup>3</sup> AI, a notoriously broad and nebulous term, refers to the science and engineering of systems that could be described as making machines intelligent, able to function appropriately and with foresight in an environment.<sup>4</sup> ML, presently the most successful sub-field of AI, enables computer systems to learn from data to improve performing a task rather than designing a static algorithm to perform a specific task. Of note, each case study varied in the AI/ML techniques used to achieve organizational ends, and organizations often used a combination of approaches (See section "Terms and AI techniques used" for further details.)

The Tata Steel Europe case, based in the Netherlands, describes how an incumbent steel manufacturer,

facing significant industry headwinds, launched an internally-developed AI strategy (referred to internally as "Advanced Analytics") to optimize its production processes with both ML models and workforce retraining to build and integrate them. The company reports significant economic benefits and productivity gains with a lower investment than alternative, higher-cost capital investments. Notably, the company reports avoiding immediate layoffs by retraining many in its workforce to develop and implement the AI models - though this may not always remain the case, nor may it be the equilibrium that emerges across the entire industry.<sup>5</sup>

The Axis Bank case, based in India, explores how an AI chatbot addressed a growing need to improve customer service in a cost-effective and scalable way through automation. In addition to examining AI's impact on Axis Bank's internal workforce, the

<sup>3</sup> In business circles, many pre-established techniques such as pattern-matching heuristics, or linear regression and other forms of statistical data analysis, have recently been rebranded as "AI" (and in the case of statistical regression, also as "ML"). We accept these expansive definitions not because they are fashionable, but because they are more useful for understanding the economic consequences of present forms of automation. See section "Terms and AI techniques used" for more details on how these terms are defined.)

<sup>4</sup> Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, (Cambridge, UK: Cambridge University Press, 2010).

<sup>5</sup> Quoted figures are reported from subject organizations, not independent analyses i

case discusses how labor impacts cascaded outside the bank and across sectors to third-party customer service providers and technical vendors.<sup>6</sup>

The Zymergen case examines how a biotech startup in the San Francisco Bay Area is applying machine learning and automation as an alternative to conventional R&D and scientific experimentation practices. The case includes observations on the productivity impact of AI in that setting, as well as the labor implications for a highly educated workforce (e.g., scientists with Ph.D.'s).

Each of the organizations profiled has different strategic motivations and thus different approaches to developing AI/ML-related applications. Tata Steel Europe is an incumbent trying to build an AI capability in-house through training its existing labor resources. Axis Bank, another incumbent, chose to bring in a third-party AI vendor to help develop its chatbot capability. Finally, Zymergen is an "AI-native"<sup>7</sup> startup developing material science R&D services.

## Common Motifs

Despite the variety in sector, geography, size and AI applications of the firms we studied (See Figure 1 below) three common motifs emerged:



1. Successful adoption of new AI systems required buy-in from management and the workforce alike. In doing so, having intelligible ML models was of paramount importance.
2. Each firm reported productivity or ROI growth in the short term, but in all cases, these effects could not be attributed solely to AI, as other process changes always accompanied the introduction of AI technologies.
3. There were varying impacts on the local workforce due to the adoption of AI systems and no reported layoffs to date, but the effects were difficult to measure internally and in some cases, cascaded beyond the firms we studied into their third-party vendors and broader business ecosystems.

<sup>6</sup> As the case illustrates, the social and labor impacts can often cascade beyond the location of the AI implementation. Kate Crawford and Vladan Joler explore this concept extensively as it relates to the "vast planetary network" of labor, energy, and data to support small interactions with an Amazon Echo. See more at [www.anatomyof.ai](http://www.anatomyof.ai).

<sup>7</sup> An 'AI-native' refers to a company that was founded with a stated mission of leveraging artificial intelligence or machine learning as a key enabling technology. 'AI-natives' can build infrastructure from the ground-up without the need to shift from legacy systems (e.g., on-premise to cloud-based storage).

# Figure 1. Overview of case study attributes

Figure 1: Overview of case study organizations

			
<b>Sector</b>	Steel	Retail Banking	Biotech
<b>Geography</b>	Europe	India	U.S.
<b>Size of Organization</b>	~9,000 employees (subsidiary)	~6,000 employees (private)	~700 employees (private)
<b>Key AI/ML Applications</b>	Use cases optimizing production processes (e.g., raw material recipe optimization)	Chatbot as new automated customer service channel	Experiment recommendation engine and automated wet lab
<b>Process of AI Development</b>	In-house (with retrained resources)	Outsourced (external AI vendor)	In-house (AI-native company)

Source: Case study interviews

# Themes and observations

## Organizational culture and change management

Organizational change management is a common theme across all case studies, with variations depending on the maturity of the AI integration within the company and the level of comfort with the introduction of new processes or technologies. The subject organizations collaborated with both internal and external stakeholders to ensure trust and buy-in and used various strategies to manage change.

### Organizational buy-in

At Zymergen, a young company with AI/ML central to its founding mission, employee buy-in was generally assumed from the beginning. At Tata Steel Europe and Axis Bank, teams needed to approach changes in culture and practice more proactively with relevant internal stakeholders. In all three cases, executive support or sponsorship and early investment in employee training and awareness-building helped to engender trust internally and led to broader organizational buy-in. For instance, Tata Steel Europe recognized the value of bringing plant operators into the ML model development and began holding “office hours” in which they and project managers could discuss the analytics models with data scientists. Data scientists at Zymergen also found that upfront and frequent involvement of physical scientists in the machine learning model was critical to building trust.

### The Critical Importance of Technical Explainability

Another important element that affected the adoption of AI within these organizations was the explainability of the AI models and techniques. For example, helping operators at Tata Steel Europe understand the early use cases and visualize the models was especially important to ensure their buy-in. Visualization and explainability also played an important role in Zymergen’s implementation of its experiment recommendation engine, as scientists wanted to understand why the model recommended certain strain improvements. Explainability was also critical to AI implementation in the Axis case, where the bank needed to be confident that the chatbot would not mislead customers, that regulatory compliance would

be met, and that the chatbot would not complete fraudulent transactions.

Accordingly, transparency, explainability, and trust of AI systems played an important role in driving ML adoption and deployment at all three organizations studied, perhaps in part due to the types of problems and processes where AI was used. AI was often integrated in heuristic-driven processes where workers previously relied on years of formal education (e.g., Ph.D’s. in biology) or on-the-job experience (e.g., steel plant operators), which could have created friction due to a sense of pride in acquired domain knowledge. Implementation was often most effective if there was close collaboration between developers, subject-matter experts, and AI systems.

### External stakeholders

Managing change for external stakeholders required more attention for Zymergen and Axis Bank than for Tata Steel Europe. Zymergen needed to convince its customers of the value of its differentiated approach, and Axis Bank was required both to comply with domestic regulations and to drive customer adoption of its chatbot. The visibility of benefits and financial impact were key drivers for Zymergen, as illustrated by its adoption of performance-based revenue elements in contracts with clients. Meanwhile, Axis Bank stress-tested its chatbot through third-party agencies and ran extensive reviews to ensure accuracy and security to build trust with both regulators and consumers. In addition, Axis Bank ran customer campaigns to drive adoption and showcase the value of its AI Chatbot.

## Impact on productivity and business results

Artificial intelligence's impact on the broader economy and productivity has long been a topic of debate. Contemporary economists have also posed questions about the possible recurrence of a past computing technology "productivity paradox," or the question of why the recent advancements in AI appear not to have translated into measurable productivity gains at the macroeconomic level. Although ongoing AI and broader digitization efforts have not yet produced clearly measurable productivity gains for the broader global economy,<sup>8</sup> productivity improvements at the micro-level were reported to us across all three subject organizations in this case study project. While the case studies are limited to three organizations, the micro-trends reported on productivity and labor would be consistent with the previous Solow Paradox, where macroeconomic benefits emerge in the decade or two after the technology was deployed.

One hypothesis is that recent progress in machine learning will be subject to large diffusion and implementation lags before effects are seen at the level of entire economic sectors.<sup>9</sup> The AI applications at all three subject organizations seem to have generated observable productivity gains, although their contribution to wider economic and labor effects may be harder to isolate.

### Integration with business strategy

In all three cases, AI/ML-related applications were part of a broader business strategy driven by automation and digitization. The broader strategies often involved particular attention to digital enablement (e.g., a shift to cloud-based storage), organizational restructurings, or workforce training. For instance:

- Tata Steel Europe has positioned its Advanced Analytics Program as part of its ongoing broader digitization effort.
- Axis Bank launched the AI Chatbot within a broader set of customer service automation initiatives across

the bank to address its growing need for improved customer service.

- Zymergen coupled a machine learning-based recommendation engine with its automated wet lab operations as a key point of differentiation.

These strategies suggest the importance of complementary innovations, both technical and organizational, to enable productivity gains observed from AI-related implementations.

### Timeframe of transformations

Finally, we observed that executives and managers want quick results, be they financial benefits or operational improvements. In the Tata Steel Europe case, executives prioritized use cases with higher value potential and lower implementation complexity to realize some financial benefits early on. Similarly, Zymergen offers performance-based contracts to its clients where the contract payout depends on results demonstrated and benefits realized through the strain improvement program. We also observed that elevated executive attention led Axis Bank to release versions of its AI chatbot early on and improve its performance through ongoing iterations rather than doing more internal development and delaying its release until a later date.

### Productivity

The business results stemming from AI take various forms in the three case studies. Tata Steel Europe experienced productivity gains through improved production yield, reduced raw material expenses, and enhanced product quality. TSE also increased productivity overall by delivering roughly 13 percent higher EBITDA (a common proxy for cash flow) through its AI-related initiatives, achieved with the same production staffing levels. Axis Bank reports it was able to handle its growing customer service volumes with fewer customer service agents, driven by

<sup>8</sup> For more, see "Is the Solow Paradox Back?", McKinsey Quarterly, June 2018.

<sup>9</sup> Erik Brynjolfsson, Daniel Rock, and Chad Syverson argue in a 2017 working paper that "lags have likely been the biggest contributor to the [modern productivity] paradox." Interestingly, the authors suggest that the economy could be in the midst of a broader economic transformation due to technological changes. The lags in the measurable effects of AI could be due to the time required to invest in and "build the stock" of the new technology as well as the "complementary innovations" required to enable AI, such as cloud-based computing, organizational changes, and new business models.



a migration in volumes from human-enabled to automated customer service channels, implying an increase in the labor productivity of the overall customer service function.<sup>10</sup> Lastly, in the Zymergen case, management reported higher labor productivity, driven by a high degree of automation in the wet lab (reportedly allowing 10 times higher experimentation throughput) and accelerated project durations (eight to ten years was reduced to three to five years) compared with conventional R&D labs.

### Return on investment

Case subjects also reported that AI/ML-related applications generally offered and generated a higher return on investment (ROI) compared to other more capital-intensive alternatives. As an example, Tata Steel Europe reported higher ROI in the first year of AI implementation compared to regular capacity investments, such as an investment in new plant equipment. Similarly, the payback period for a particular piece of plant equipment is four to six years, whereas Advanced Analytics use cases at Tata Steel Europe usually break even in one to two years. Zymergen offered accelerated strain improvement processes to its customers that may allow them to postpone or forgo large capital investments in downstream operations (e.g., fermentation manufacturing facilities). This evidence suggests that adoption of AI-related technologies could implicate broader trends of capital efficiency or contribute to shifts in organizations' investment patterns.

## Impact on labor and the workforce

Debate around the consequences and implications of AI-related technologies is often closely related to its potential impact on the workforce. Many say that broader adoption and advancement of AI will lead to profound workforce changes, suggesting that productivity gains from AI-related technologies may manifest themselves through labor reductions and growing inequality.<sup>11</sup> In this limited sample set of case studies, we examined whether productivity gains from AI-related integrations came at the expense of jobs, or whether the impact was more nuanced. We also explored what types of workforce skills may be required to adopt AI-related technologies and the challenges associated with the introduction of new technologies.

In the three cases, we observed varying degrees and forms of labor impacts, both direct and indirect, demonstrating that labor effects are often multi-layered and can extend beyond the core organization implementing AI. The case studies also demonstrated that organizational and societal contexts play a role

in how AI-related initiatives impact labor. We noted distinct labor characteristics and dynamics in each case depending on how AI was integrated within the company's operations. For instance, Tata Steel Europe operates in a mature and highly unionized industry with a relatively inflexible labor market. At Axis Bank, the impact of AI on labor was often indirect, as the majority of its customer service operations were outsourced and thus were not conducted by direct employees of the bank. Zymergen started out as an AI-native company and communicated the importance of automation to its workforce from the very beginning through its mission statement - though its business model may have downstream impact on its customers' R&D teams, shifting to new types of work or potentially leading to lower hiring rates in the future. As these cases demonstrate, the social and economic impact of AI and other automating technologies goes beyond the immediate sites of implementation and extends externally across supply chains, partners, and customers.

<sup>10</sup> We do not have a measure of hours worked to estimate the increase in labor productivity precisely.

<sup>11</sup> Some have argued that inequality could increase with the proliferation of AI in the long term. While we do not address this question, please see [Joseph Stiglitz and Anton Korinek's paper](#) for more: "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," December 2017.

## Workforce reductions

While no direct AI-related workforce reductions were observed at Zymergen or Tata Steel Europe in the short-term, we did observe direct workforce reductions in one case.<sup>12</sup> Axis Bank specifically aimed to automate its human-enabled customer service offerings and reduce the size of its outsourced customer service team. However, Axis pointed out that while the size of the customer service team may be further reduced, the team is unlikely to be completely eliminated because of the technical challenges and operational complexity posed by multiple languages in India. Similarly, while Tata Steel Europe's workforce size has not changed as a result of the Advanced Analytics program, it is likely to change in coming years. With an aging workforce and maturing AI-related technologies in production, voluntary attrition of the production crew reaching retirement age in the coming five to fifteen years could potentially lower the overall size of the production workforce, even without direct displacement of particular workers.

## Workforce composition

A common theme across all three cases is the changing composition of the workforce and the addition of new skill-sets and profiles to the organizations (e.g., data scientists, automation engineers). Zymergen's core team<sup>13</sup> of researchers and scientists is 50 percent smaller than that of a conventional R&D lab. Zymergen, however, has a higher number of support staff partially dedicated to each team, compared to a conventional R&D lab. Support staff are composed of data scientists, software engineers, automation engineers (about 25 percent of Zymergen's workforce) or lab technicians and personnel to run the wet lab (about 8 percent of Zymergen's workforce). At Tata Steel Europe, upon initiation of the Advanced Analytics Program, the company built a team of data scientists and engineers through internal retraining. Although not as significant, the implementation of the AI chatbot led to re-tasking of key IT personnel at Axis Bank as well.

## Workforce hiring

Shifts in the composition of a workforce also trigger changes in companies' hiring needs, increasing the demand for hybrid profiles with interdisciplinary expertise (e.g., data science and steel or biology). Zymergen's target hiring profile is different than that of a conventional R&D lab; for example, entry-level research associates are expected to have experience in biology and have data science capabilities, such as proficiency in Python or R. Similarly, Tata Steel Europe reminds its employees that it is in the business of making the best steel, not developing the best AI models. Hence, it looks for internal talent with a deep understanding of the steelmaking process who can be trained in data science skills.

## Workforce ramifications outside case organizations

As noted, labor implications are rarely limited to the core subject organization, and impacts usually cascade externally to customers and vendors. Axis Bank's implementation of its chatbot resulted in temporary business gains for third-party vendors such as technology developers and certification agencies, while leading to a reduction in the number of personnel assigned to its account at its third-party customer service provider. Zymergen's strain improvement programs could be considered akin to an outsourced R&D lab, which, if proven successful, could displace a customer's internal R&D workforce or create lower hiring demand for scientists and downstream manufacturing labor.

<sup>12</sup> It is not clear what the net-impact of AI on jobs will be in the near future. The [McKinsey Global Institute estimates](#) that "total full-time-equivalent-employment demand might remain flat, or even that there could be a slightly negative net impact on jobs by 2030," yet demand for new types of jobs may increase, as seen with the advent of the personal computer in the late 20th century.

<sup>13</sup> This only includes scientists and research associates and does not account for data scientists, automation engineers, and lab technicians that support teams with their services.

## Technical implementation overview

### Processes being replaced by AI

The technical advantages of AI/ML technologies allowed all three organizations to replace or enhance conventional processes that had been performed less effectively or efficiently by humans in the past. In particular, AI-related technologies helped these companies process high volumes of data at a faster pace to uncover patterns more reliably, more quickly, or more cost effectively than a human could. For instance, Zymergen leveraged AI in order to better understand the complexity of the microbial genome during fermentation. Tata Steel Europe analyzed thousands of variables throughout its production processes to optimize the material recipes that determine the chemical properties of steel. Axis Bank implemented an AI chatbot capable of handling thousands of customer inquiries simultaneously across multiple banking products with 24/7 availability.

### Data strategy and infrastructure

Subject organizations had differing approaches to generating their training data for the core machine learning systems. Zymergen took a clean slate approach<sup>14</sup> to its infrastructure by investing early on in standardizing and accelerating the speed of its data collection processes at the wet lab. This created high total data acquisition costs, including high upfront investment costs and higher ongoing experiment costs (e.g., raw materials) because they ran more experiments as part of their learning process. Tata Steel Europe had to significantly improve its data infrastructure and systems to consolidate data in the cloud and improve overall data quality. Axis Bank has also put significant effort into generating its data, for instance assembling the chatbot's training data manually, because open conversational NLP<sup>15</sup> datasets are not applicable to banking domains, and in any case need to be customized for an organization's products and business practices. Today, Axis Bank faces ongoing challenges relating to the limited standardization and varying quality of the data generated through new chat conversations from users. To address this challenge, the Digital Banking team at Axis Bank has conducted manual reviews of the chatbot during each release cycle, and it has also received a weekly list of questions the chatbot couldn't answer. These approaches by subject organizations illustrate the costs of building and maintaining conversational AI systems over time.

A related issue: More sophisticated models and techniques did not always generate better results at subject organizations, since they often required much larger data sets to be effective. This proved to be a challenge for all three organizations, which are relatively new in their experiences with AI/ML applications. More complex models were also more difficult to explain to the end-users (e.g., scientists, plant operators) because of the "black-box" phenomenon whereby the reasons for models' recommendations weren't fully transparent.

<sup>14</sup>Zymergen is an "AI-native" company that was founded in 2013. As such, the company started its data storage in the cloud. All data infrastructure could be built with a clean slate and modern toolchains, making data exportation and analysis on cloud systems easier than it might be for an incumbent (such as Tata Steel Europe). The latter might be dependent on proprietary or embedded on-premise systems that were installed without these objectives in mind.



<sup>15</sup>Natural Language Processing, a popular subfield of AI

# Terms and AI techniques used

In order to understand how the techniques studied in each case study fit into the larger spectrum of AI, we depict an overview of the techniques deployed at Zymergen, Tata Steel Europe, and Axis Bank, as well as one framework for a spectrum of AI techniques. While the figure below tries to capture specific techniques, it should be noted that the research also reflects broad usage of terminology such as 'AI' or 'machine learning' as deployed during interviews. To further illustrate the specific uses of technology, we provide more detail in the individual case studies on how they were utilized and the use cases targeted.

Figure 2. Overview of case study attributes

Figure 2: Overview of case study organizations

			
<b>Techniques<sup>1</sup></b>	<ul style="list-style-type: none"> <li>• Rule-based methods</li> <li>• Linear regression</li> <li>• Polynomial regression</li> <li>• Decision trees</li> <li>• Random forest</li> <li>• Nutral networks</li> <li>• Monte Carlo Methods</li> </ul>	<ul style="list-style-type: none"> <li>• Rule-based methods</li> <li>• Nutral networks</li> </ul>	<ul style="list-style-type: none"> <li>• Rule-based methods</li> <li>• Linear regression</li> <li>• Polynomial regression</li> <li>• Bayeasian inference</li> <li>• Bayesian hierarchical modeling</li> </ul>
<b>Capabilities</b>	<ul style="list-style-type: none"> <li>• Computer vision</li> <li>• Physical automation and robotics (e.g., torpedo tapping)</li> </ul>	<ul style="list-style-type: none"> <li>• Natural Laungage Understanding (NLU)</li> <li>• Natural Language Processing (NLP)</li> <li>• Natural Laungage Generation (NLG)</li> </ul>	<ul style="list-style-type: none"> <li>• Computer vision</li> <li>• Wet lab robotics               <ul style="list-style-type: none"> <li>- Liquide handling systems</li> <li>- Automated plate readers</li> <li>- Acoustic dispensers</li> </ul> </li> </ul>

1.The list of techniques below are non-exhaustive, but show a range of the most commonly used or attempted techniques at the organization.

Source: Case study interviews

## In each case, we captured the range of AI-related techniques that were deployed:

- **Tata Steel Europe:** Tata Steel Europe's Advanced Analytics Program was the company's term for its increased institutional focus on data science, including AI and machine learning, in contrast to "traditional" analytics. It is associated with both cultural and technical changes from earlier practice, including a significant increase in the variables used in the analytics; the use of both structured and unstructured data (e.g., images, audio); more predictive (vs. descriptive) analytics; and use of more advanced techniques, i.e., moving from linear regression and decision trees to techniques such as random forest and some testing of neural networks.
- **Axis Bank:** At the time the case was written, Axis Bank used AI technologies including natural language processing (NLP)/natural language understanding (NLU) and natural language generation (NLG). It used neural networks to extract information from unstructured text, determine a user's intent, select appropriate action, and respond in an appropriate language for the user (e.g., English if the user is querying in English). In addition, the AI chatbot ("Axis Aha!") included rules-based systems as backup responses when the chatbot had problems with the accuracy and relevance of its answers, referring customers to other service channels in these instances; these explicitly rules-based systems are not regarded as AI within the case study, rather, they are supplemental technologies that help to facilitate the use of less-structured systems.
- **Zymergen:** Zymergen uses a range of analytics and AI techniques for its experiment recommendation engine and for its data normalization and cleaning process in the automated wet lab. At the time the case was written, these techniques included linear regression, polynomial regression, Bayesian hierarchical modeling, and convolutional neural networks (CNNs)<sup>16</sup> to explore tradeoffs in accuracy, complexity, implementation time, and change management within its organization. At Zymergen, these AI techniques are accompanied by equipment and tools in the automated wet lab such as liquid handling systems, robotic colony pickers, barcoders, acoustic dispensers, automated plate readers, robotic rule-based scripts, and systems or software used to operate this equipment. Throughout the case, these technologies and equipment are referred to as the "automated wet lab," "robotics," "advanced machinery," or, in aggregate, "high-throughput screening." The terms are meant to encompass the equipment, systems, and computer scripts that support Zymergen's wet lab and are distinct from our direct discussion of AI.

<sup>16</sup> CNN's were tested as part of Zymergen's broader recommendation engine and were also used in isolated cases within the lab (e.g. computer vision for plate readers).

# Methodology

Subject organizations were recruited from a pool of 100+ candidates that 'AI, Labor, and the Economy' Working Group compiled. The final set of organizations prioritized for study reflects a combination of their willingness and ability to participate in the project and the intention to profile organizations representing a variety of sizes, geographies and industries. To this extent, the subjects of these case studies were not chosen randomly and were sourced through existing relationships within PAI's AILE Working Group. Zymergen has raised \$574M from investors such as SoftBank Vision Fund, Data Collective, ICONIQ Capital, and McKinsey & Company (a co-author of the case studies), among others.<sup>17</sup> Also, the subject organizations may be subject to selection bias. For instance, companies that have had successful experiences implementing AI at their organizations may be more eager to discuss their results, particularly if they have done so without adverse impacts to labor. Yet sourcing these case studies through existing relationships also allowed for an elevated level of trust and candor in probing more deeply about companies' experiences.

The three case studies were developed over the course of six months in late 2018 and early 2019. Key sources of insight included:

- *Interview-based research* with a set of key decision-maker stakeholders from subject organizations such as business unit heads, functional leads, data scientists, and developers. Interviews were generally limited to decision-makers for AI/ML-related applications and their use cases. Interviews thus represented a managerial perspective and did not involve other workers who may have been directly impacted by AI/ML-related applications such as plant operators, call center employees, or R&D lab technicians (and thus should caution over-extrapolation from case observations). Further details of interviewees include:
  - **Tata Steel Europe:** Ten employees were interviewed (three senior executives, one plant manager, one technical director, two data scientists, one data engineer, one automation engineer, one analytics project manager). Interviews did not include steel plant operators, who are more implementation-focused.
  - **Axis Bank:** Five people were interviewed (four senior executives within Axis Bank and one employee from the third-party technology vendor contracted to develop the chatbot). Interviews did not include customer service agents, whose jobs had been outsourced.
  - **Zymergen:** Ten employees were interviewed (three senior executives, one data scientist, two business development representatives, four scientists at different levels of seniority). Interviews did not include lab technicians or Zymergen's customers.
- *External supplemental research* was conducted by representatives from not-for-profit and for-profit organizations affiliated with the Partnership on AI. External research primarily focused on gathering broad industry and macroeconomic insights in the form of expert consultations, literature scans of third-party publications, and general web research.

Representatives from the Partnership on AI's non-profit and for-profit organizations supported the case study development by conducting interviews, drafting case documents, and supplementing the case with external research or expert interviews on industry and macroeconomic dynamics. Research was syndicated and reviewed by various stakeholders throughout development, including the "[AI, Labor, and the Economy](#)" (AILE) Working Group of the Partnership on AI.

We hope these case studies will be insightful additions to the research landscape, and wish to complement these with follow-on research that may help to further address our objectives.

<sup>17</sup> During the time of writing the case study in fall 2018, the company had raised \$174M. On December 13, 2018, the company announced a \$400M Series C round from multiple investors. See coverage of the announcement on [Bloomberg](#) and the [Wall Street Journal](#).

## Limitations and Further Work

We acknowledge that this group of case studies is not comprehensive, and that its chosen methodology and scope has limitations. In some cases, we were unable to collect first-hand the underlying data on the investment in the AI integration to fully understand the productivity impact (this was especially true for Axis). We also did not collect data directly from workers directly impacted by the AI implementation (e.g., production crew at Tata Steel Europe, Zymergen’s customer R&D teams, Axis Bank’s outsourced customer service representatives). As documentation of actual impact of real-world AI applications, albeit partial, the case studies can contribute nuanced pictures of the way that AI and ML are playing out in real workplaces. These observations and insights can and should inform dialogue on the impact of AI on labor and the economy. For a more extensive discussion of gaps in the literature and further areas of potential research, please see our extended case studies.

## Conclusion

From our research, we find evidence of AI impacting processes of individual organizations to varying degrees, with common themes regarding implications for workers, necessary cultural components of AI implementations, productivity gains, and the broader complexities under which organizations make these decisions.

All three of our case study organizations — Tata Steel Europe, Axis Bank, and Zymergen — provide insights into how organizations can best practice cultural and change management components of a successful AI implementation. For instance, the momentum of “organizational buy-in” proved a common theme, whether via support from senior executives or a mission centered around the value of analytics. Similarly, in cases of new AI implementations, helping workers to understand the impetus for the AI implementation early-on, ensuring that workers understand how the new programs work through intelligible and explainable AI, and informing these programs with the perspectives of existing employees proved especially valuable.

Additionally, all three organizations reported varying degrees of productivity improvements from their deployments of AI. In tandem, all three instances of AI implementations entailed changes for the way labor operates: sometimes a shift in the roles humans take on, compared with technology; sometimes a shift in the skills humans must possess to thrive in their roles; and sometimes a shift in the number of humans required to complete a process or deliver a service.

Though labor displacements generally happened outside these implementing organizations, if at all, there may eventually be reductions within the implementing organizations. As skilled workers retire, and as companies face increased industrial pressure to operate leanly, companies may decide to cut jobs, rather than shift their workers into other areas. Yet even in these instances, it is possible this job loss may be offset by internal or external job gains elsewhere, which may be more highly-skilled or -compensated than the jobs they replace.

Though each case subject experienced certain common themes, their particular experiences implementing AI were also informed and constrained by the relevant contexts in which the organizations exist. For instance, we note a greater flexibility for Zymergen - an “AI-native” company - in pursuing its programs, than for longer-established companies, particularly those that may have greater local regulatory complexity or may have unionized workforces. Managers, scholars, and others should not overlook these contextual factors when trying to gauge AI’s impact on the broader economy and on individual organizations. No AI implementation is one-size-fits-all, and the details matter when attempting to understand a topic that is too often discussed in broad strokes. For this, we offer the Partnership on AI’s series of extended case studies.

# Appendix

## Definitions and terms used

While we acknowledge that there is no consensus on the definition of terms such as AI and automation, we would like to explain how these terms are used in the compendium:

**Artificial intelligence/AI** is a notoriously nebulous term. Following the [Stanford 100 Year Study on Artificial Intelligence](#), we embrace a broad and evolving definition of AI. As Nils J. Nilsson has articulated, artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment. (Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, (Cambridge, UK: Cambridge University Press, 2010).

Our definition of **automation** is based on the classic human factors engineering definition put forward by Parasuraman, Sheridan, and Wickens in 2000: <https://ieeexplore.ieee.org/document/844354>, in which automation refers to the full or partial replacement of a function previously carried out by a human operator.<sup>18</sup> Following Parasuraman et al.'s definition, levels of automation also exist on a spectrum, ranging from simple automation requiring manual input to a high level of automation requiring little to no human intervention in the context of a defined activity.

**Explainable AI** or **Explainability** is an emerging area of interest in communities ranging from DARPA to criminal justice advocates. Broadly, the terms refer to a system that has not been “black-boxed,” but rather produces outputs that are interpretable, legible, transparent, or otherwise explainable to some set of stakeholders.

In this compendium, a **model** refers to a simplified representation of formalized relations between economic, engineering, manufacturing, social, or other types of situations and natural phenomena, simulated with the help of a computer system.

<sup>18</sup> Our definition draws on the classic articulation of automation described by Parasuraman, Sheridan, and Wickens (2000): <https://ieeexplore.ieee.org/document/844354>



# Acknowledgements

PAI is deeply grateful for the collaboration of so many colleagues in this endeavor, including project leadership from AILE Working Group member Madeleine Clare Elish (Data & Society) and Working Group Co-Chairs Michael Chui (McKinsey) and Elonnai Hickok (Centre for Internet and Society).

In addition to this holistic leadership, authorship of the Tata Steel Europe, Axis Bank, and Zymergen cases was co-led by Lars Jensen (United Nations Development Programme), Elonnai Hickok, and Tithi Chattopadhyay (Center for Information Technology Policy, Princeton University) respectively, with the support of Arif Cam, Turker Coskun, Travers Nisbet, and Eric Schreiber (all from McKinsey).

The authors are grateful to members of the Partnership on AI for their thoughtful feedback and reflections.

The authors also wish to thank the many additional minds and talents without whom this work would not have been possible - particularly Aayush Rathi (Centre for Internet and Society), Margaret Steen, Sue Wunder, and KNI Design.



**PARTNERSHIP ON AI**