

Received October 1, 2021, accepted November 9, 2021, date of publication November 11, 2021, date of current version November 30, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3127548

# aiSTROM—A Roadmap for Developing a Successful AI Strategy

**DORIEN HERREMANS<sup>1b</sup>**, (Senior Member, IEEE)

Department of Computer Science and Design Pillar, Singapore University of Technology and Design (SUTD), Singapore 487372

e-mail: dorien\_herremans@sutd.edu.sg

This work was supported in part by the Singapore Ministry of Education under Grant MOE2018-T2-2-161, and in part by SRG ISTD 2017 129.

**ABSTRACT** A total of 34% of AI research and development projects fail or are abandoned, according to a recent survey by Rackspace Technology of 1,870 companies. In this perspective paper, a new Strategic RoadMap, aiSTROM, is presented that empowers managers to create an AI strategy. A comprehensive approach is provided that guides managers and lead developers through the various challenges in the implementation process. In the aiSTROM framework, the top  $n$  potential projects (typically 3-5) are first identified. For each of those, seven areas of focus are thoroughly analysed. These areas include creating a data strategy that takes into account unique cross-departmental machine learning data requirements, security, and legal requirements. aiSTROM then guides managers to think about how to put together an interdisciplinary artificial intelligence (AI) implementation team given the scarcity of AI talent. Once an AI team strategy has been established, it needs to be positioned within the organization, either cross-departmental or as a separate division. Other considerations include AI as a service (AIaaS) and outsourcing development. Looking at new technologies, one has to consider challenges such as bias, the legality of black-box models, and keeping humans in the loop. Next, like any project, value-based key performance indicators (KPIs) need to be defined to track and validate the progress. Depending on the company's risk strategy, a SWOT analysis (strengths, weaknesses, opportunities, and threats) can help further classify the shortlisted projects. Finally, one should make sure that the strategy includes continuous education of employees to enable a culture of adoption. This unique and comprehensive framework offers a practical tool for managers and lead developers.

**INDEX TERMS** Artificial intelligence, computational and artificial intelligence, engineering management, project management, system implementation, technical management.

## I. INTRODUCTION

Artificial Intelligence systems are steadily becoming more and more integrated into our daily lives: from autopilots in our cars [2], to personalised suggestions in music streaming app [3], and AI diagnostics in the hospital [4]. To survive as a company, it is essential to leverage these technologies to improve one's business processes and offer competitive products or services. But how does one about developing a successful artificial intelligence (AI) strategy? AI systems are unlike traditional software as they have specific requirements and risks [5]. According to a study by TechRepublic [6], 85% of AI projects do not manage to bring their intended service to the business. While there are scattered reports

and research papers on individual aspects of AI projects, this report aims to bring some of these together from the perspective of how to come to a successful AI project for the business. This paper offers a structured framework to guide an organization's thought processes during the development of an AI strategy, from a multidisciplinary perspective. The proposed framework stems from the experience of the author who has led multiple AI projects in the last decade.

The emergence of new technologies such as the PyTorch and Keras software, as well as hardware (e.g. GPUs), have made it much easier to develop AI applications [7]. This evolution has created an immense amount of opportunities for businesses and organizations. When deciding on which opportunities to implement, however, it is important to consider the risks and requirements that come along with such a project, e.g. big data storage, ethical concerns

The associate editor coordinating the review of this manuscript and approving it for publication was Derek Abbott<sup>1b</sup>.

and privacy, to name a few. AI projects are more than just software projects, as they come with their own unique challenges. According to a recent survey of 1,870 companies by Rackspace Technology [1], 34% of AI research and development projects fail or are abandoned. As a reason for that failure, a poorly conceived strategy is listed by 31% of companies. The proposed aiSTROM roadmap framework will help AI managers to make key decisions to facilitate successful AI projects, by bringing forward the unique challenges that will be faced along the way.

The AI strategic road map tool (aiSTROM) presented in this article offers a concrete roadmap for managers who want to start integrating AI in the organization, as well as developers doing smaller-scale assignments, and consultants overseeing AI projects. In the next section, some background history of AI is introduced, together with why companies need to think about AI strategies now. This is followed by an outline of the aiSTROM framework, which starts by creating a list of potential AI projects, evaluating data strategies, organizing AI development, being aware of technological choices, considering risks and benefits, defining KPIs, and finally, enabling a cultural shift to stimulate AI adoption.

## II. A.I. – WHY NOW?

The concept of artificial intelligence (AI) is an old one. The first mention of AI is generally attributed to McCarthy in 1956, when he organized the “Dartmouth Summer Research Project on Artificial Intelligence”, a workshop attended by what turned out to be some of the leading researchers in the field [8]. Research on AI systems has since advanced from developing rule-based chess computers to deep learning-powered chat agents, and self-driving vehicles. Needless to say, the concept of artificial intelligence has gained immense popularity in the last years [9]. During the first five days of March 2021, a total of 728 academic papers in the category ‘Artificial intelligence’ were published on the popular computer science preprint repository ‘arXiv’ from Cornell University.<sup>1</sup> The field of AI has existed for quite a while, yet, it has never been so popular as now. Before we dive into how to tackle AI projects in organizations, we ask the question: what exactly is artificial intelligence?

Some definitions equate artificial intelligence to self-awareness [10]. Looking at the term through a practical, modern-day lens, one notices that the term is used more often to describe models that learn patterns from data and use this to predict future behaviour or instances. Tremendous progress has been made in the last few years in terms of narrow AI systems, i.e. systems with one concrete task: detect if this image is a dog or a cat; or: change the style of this photograph as if it was painted by Monet. The harder task is that reaching a general AI [11], or human-level intelligence, whereby the system can solve any new task it is presented with [12]. This objective has not yet been reached. In this article, the term AI is used to refer to narrow AI systems that are widely

implemented in consumer technology and back-end systems today.

While the first neural network was already proposed in 1944 [13], its popularity has gone through periods of reduced interest and funding, i.e., the so-called AI winters [14]. The first AI winter started in the mid-1970s, due to a combination of circumstances, including disappointing results in machine translation, as well as the Lighthill report, which stated: “... in no part of the field have discoveries made so far produced the major impact that was then promised.” (James Lighthill, 1973, as quoted by [15]). This report was created at the request of the UK’s Science Research Council and was meant to evaluate the state of AI’s progress. It was not until the 1980s that neural network research became popular again, this time fueled by the work of connectionists such as John Hopfield and David Rumelhart. The last AI major AI winter happened roughly from the late-1980s until the mid-1990s. Its recovery was stimulated by the arrival of the Internet and the availability of large storage devices, both of which led to the age of big data. In addition to this, the current AI boom is further stimulated by the invention of novel, powerful technologies such as convolutional neural networks, along with dedicated processing devices called GPUs. These three factors together, are core contributors to the current AI spring [16], which is offering unprecedented opportunities to businesses.

A report by Price Waterhouse Cooper indicates that 72% of business leaders think that AI will offer the business advantage of the future [17]. By using AI technologies, one can better understand what happens in an organization down to an unprecedented level, as well as offer novel services and products. Furman and Seamans [18] show that AI technologies are already having a huge impact on the economy. The consultancy agency Accenture estimates that, by 2035, AI will have an added economical value of 8,305 billion USD in the United States alone; 1,079 billion USD in Germany; and 2,068 billion USD in Japan [19]. We are thus at the point when AI is reshaping the global economy. Rapid technological disruption is happening through AI technologies, thus empowering companies such as Salesforce, The Trade Desk, and many others. For some companies, however, the adoption of AI is not so straightforward, and it can pose a threat to the continued existence of the organization. Many companies, big and small want to jump on the AI bandwagon, without necessarily knowing what that means to their organization. Gartner [20] argues that companies should strategize and plan for AI as it can ‘deliver clever business outcomes’. This calls for an urgent need for companies to implement an AI strategy. AI has become a driving force in the quick growth of some of the biggest technology companies. Traditional organizations may ask themselves how and when they should start implementing an AI strategy. The general rule of thumb seems to be to start quick, so as to not miss critical opportunities, while newcomers take over the market. A recent report by Gartner also states that these new technologies could lead to a gap between the early adopters and the non- or late-adopters [21].

<sup>1</sup><https://arxiv.org/list/cs.AI/current>

That being said, one should keep in mind that predictive data systems are not entirely new and they have already provided significant advantages to companies, which can be quickly leveraged. Examples of this include business intelligence systems [4], customer churn prediction [22], and APIs for routing [23]. In the next section, we outline the different elements of the proposed roadmap.

### III. STRATEGIC AI FRAMEWORK (aiSTROM)

To guide the strategic decisions when integrating new AI technologies into an organization, this perspective paper offers a roadmap. The aiSTROM framework (see Figure 1 and described below) aims to provide concrete thinking points and strategic decisions that should not be overlooked when tackling AI projects. Taking some of these items into account will help the organization to successfully select the best AI projects to focus on and give them more insight into the strategic decisions and challenges involved.<sup>2</sup>

#### A. IDENTIFY OPPORTUNITIES

The myriad of possibilities offered by AI does not come without challenges for organizations: how should they adapt strategically to stay competitive? In other words, which project should the organization focus on first? The aiSTROM framework proposes to first create a list of possible goals and problem statements, each of which can be tackled by an AI system. This will result in a list of different possible projects with a clear task definition, each with varying levels of risk and different requirements. This list should be created by a team consisting of members with both domain/organization knowledge (e.g. c-level and employees) as well as AI knowledge (e.g. AI consultants or developers). As the organization forms a strategic vision for integrating AI across functional borders in the organization, this will be deeply rooted in the goals that have been outlined in this step. Having clear goals is an essential step in any successful IT strategy [25], [26], it is no different for AI projects.

There are a few things that one should keep in mind when identifying these problem statements. One could identify internal efficiency gaps: where could AI make the organization perform smoother and more efficient? Secondly, what new AI technologies exist and how these could be of potential use to the organization (see Section III-E)? Looking at competitors may also provide some insight into what is possible; what should the organization aim to emulate?; or what is the organization competing with? Are there new business models on the rise that could replace the currently offered services, and should AI technology be leveraged to adopt these strategies as well? For instance, a traditional taxi company might do well to implement an online booking platform in

<sup>2</sup>The categories below were identified based on the author's extensive domain knowledge as well as a (non-exhaustive) literature search. While all the statements below are backed by references, this paper hence does not offer an exhaustive literature review such as in a meta-review (as categorised by Grant and Booth [24]), instead, it provides an overview of (categorised) challenges that arise when implementing AI systems.

light of competitors such as Grab and Lift. A similar inspiration for finding AI projects that create value comes from Baccala *et al.* [17], who state that AI can create value in three different manners: 1) by automating processing; 2) identifying trends from historic data; and 3) by making intelligent predictions that empower human behaviour. Based on current trends one can assume that AI will increasingly form part of business solutions, be it in terms of supporting internal processes or as part of services/products [27].

By forming a team that includes both experts with AI expertise, as well as domain experts, the most interesting project ideas can be thought of. Once this list has been made, the top  $n$  projects are selected. Here,  $n$  could for instance be 3 to 5, unless of course, in the case of a huge AI-based company. When choosing these projects, one most likely will want to maximize the potential impact, given the lowest effort. These top  $n$  projects will then be further evaluated and described using the aiSTROM framework laid out in the next subsections. This will result in an analysis that will provide insight into how the implementation of AI projects can be successfully managed, fit into the corporate strategy, as well as ensure that it provides value. By identifying which projects to focus on, time and resources are used efficiently.

#### B. DATA

We emerged from the last AI winter not just because of the development of new tools (e.g. convolutional neural networks) and fast hardware (i.e., graphic processing units or GPUs), but also due to the prevalence of data [16]. Ever since the emergence of the Internet, data is collected everywhere, and the amount of data in warehouses is growing exponentially. Since data is essential for training any successful machine learning algorithm, this is one of the boosting factors of the growing AI industry. Press [28] predicted that the market for big data analytics alone will have over USD 203 billion in revenues worldwide.

The first mention of 'big data' can be traced back to Cox and Ellsworth [29]: "... provides an interesting challenge for computer systems: data sets are generally quite large, taxing the capacities of main memory, local disk, and even remote disk. We call this the problem of big data. When data sets do not fit in main memory (in core), or when they do not fit even on local disk, the most common solution is to acquire more resources.". These days, big data is often characterised by a set of V's. For instance, Russom *et al.* [30] and Laney [31] define the 3Vs of big data as Volume (it requires a lot of storage space), Velocity (data comes in quickly), Variety (diversity and messiness). These original 'Vs' have later been expanded to include Value (how much data contributes to business value) [32], and even Veracity (biases, noise, and abnormality in data) [33]. While data is an essential resource for any AI algorithm, it can come with some challenges. Some of the issues that might arise from storing/collecting/using data are discussed below. Considering these issues for each of the above-defined AI problem statements and preparing a proper data strategy is an essential step in the AI framework,



FIGURE 1. Strategic roadmap (aiSTROM) for AI implementation.

or aiSTROM, document. For more details on how to set up a comprehensive data governance strategy, the reader is referred to [34].

According to Rackspace Technology [1]’s survey, the reasons for the failure of AI research and development include a lack of data quality (34%), and a lack of production-ready data (31%), thus making data strategy a key consideration in any AI project. In what follows, the focus will be on a few AI-specific considerations related to data strategy.

### 1) DATA SOURCES

A first question that needs to be asked for each of the problem statements is: does the organization already have the necessary data? Maintaining a centralized overview of which data that the organization collects can help the machine

learning engineers and designers to ideate which data might be useful for training an AI system. If the organization does not collect data, it is a good idea to start doing so, even before any machine learning models are built. This historical data might prove useful later on for currently undiscovered AI opportunities [35].

### a: BIG DATA OR NOT?

The common belief is that the more data is fed to a machine learning algorithm, the better the result. This is partly true, however, it is not always the case that having massive amounts of data results in the best predictive models [36]. For instance, having large amounts of ‘sparse’ input features will cause an exponential increase in necessary data and computing resources for obtaining a useful output [37],

a phenomenon termed ‘the curse of dimensionality’ by Bellman [38]. In addition, the characteristics of big data cause it to come at quite a cost. Volume necessitates buying or renting large quantities of storage space. While the cost per GB has gone down drastically over the last few years, the cost of *reliable* (secure, backed-up) storage is still higher due to required data replication, continuous administration, and higher hardware requirements.

#### *b: COLLECTION VERSUS ACQUISITION*

Collecting data to create a dataset typically requires time. It may also need ethical board approval, quality checks and preprocessing, as well as someone to oversee the process. With this in mind, one might be tempted not to collect all necessary data themselves, but to purchase existing data, provided that it is available to buy. Another challenge in data collection is that it can be slow to build a sufficiently large dataset. This process may rely on gathering historic data, such as customer transactions, which may mean waiting until enough transactions have happened. Another type of data collection can be done by annotation, either by a select group of experts or using popular crowdsourcing platforms like Amazon Mechanical Turk [39]. In this case, an annotation cost is paid, which is often related to reliability. It is worth keeping in mind that crowdsourcing annotation platforms, even though there are measures to counteract this, often produce noisy, unreliable data. This means that this type of data typically needs more cleaning, preprocessing, and in general, a larger volume of annotations to be useful [40]. The old saying ‘garbage in, garbage out’, also applies to machine learning models, so the cost of cheap data annotation might in reality be much higher.

Even when a company’s AI strategy is not yet set up or finalized, gathering historic data can be beneficial for the (as of yet undefined) future uses. Start collecting as soon as possible seems to be the message [41]. Using historic data collected over time can not only be used to train future AI models that have not yet been thought of, but also to facilitate exploratory data analysis. When exploring available data sources, data scientists and domain-experts ask specific questions and try to uncover if the data may support hypotheses that can then be further explored with an AI model. This can be a time-consuming process, but may lead to interesting problem statements or ideas for potential AI systems [42].

## 2) LEGAL ISSUES: PRIVACY AND SECURITY

Through scandals like that of Cambridge Analytics [43], the public is becoming aware that companies are storing their personal data, and may not be keeping it to themselves. Fundamental privacy rights are becoming ever more important. Whenever an organization is storing data, legal requirements should be carefully considered so as to keep this data safe and private.

Along with the use of data come certain risks and considerations. While the general notion is that more data enables better AI systems, one has to consider that this can cause privacy

issues. By implementing a data sharing policy in the organization, one can reduce the need to create duplicated data and allow engineers come up with ideas and build systems more quickly. The benefits of having data widely available within the organization, are countered by the increased risk of data breaches and privacy issues [44]. An example of a data breach is the case of Ashley Madison [45], a dating website enabling extramarital affairs. Needless to say, their data was highly personal and such breaches should be avoided at all cost. If the organization implements greater transparency (of data) between departments to empower AI development, it will require a shift of the traditional organizational boundaries. As an organization, you have to find the balance between quick access to data without asking for permission, so as to enable new ideas and quick prototyping of AI systems, versus adequate security to minimize the risk of data breaches. Even when the internal permissions are secured, one should also protect against malicious external attacks such as exfiltration attacks [46], in which unauthorized malware or a malicious actor performs an unauthorized data transfer.

Given the above considerations, it is essential that the data requirements necessary to achieve the set-out goals in Step 1 (subsection III-A) are explored thoroughly and a dedicated internal data privacy manager or external consultant is put in charge of correctly dealing with data given the local privacy laws such as the California Consumer Privacy Act (CCPA) [47], Singapore’s Personal Data Protection Act 2012 (PDPA) [48] and the European Union’s General Data Protection Regulation (GDPR) [49]. The latter, for instance, grants customers access to data that has been collected about them. It should also be noted that if the organization is active in multiple countries, the privacy laws of those respective countries should all be respected.

## 3) STORING DATA

As part of the data governance strategy [34], an organization must also think about where to store data. If a lot of (possibly sensitive) data will be stored, a significant infrastructure investment may be needed along with the manpower needed to administrate the data storage (e.g. data engineer). Instead of organizing internal storage, an external data centre may offer a quick way to get reliable storage with guaranteed security. One should also keep in mind that transferring data is not always fast, hence data should best be stored as close to the user of that data as possible [50]. For instance, if the company will be serving many customers in Singapore, but is based in Iceland, it may be worthwhile considering using a Singaporean data storage service.

The data that is collected or purchased may be stored in its raw form in a data lake (unstructured), or it may be preprocessed for indexing into a structured form in a data warehouse (i.e., suited for relational databases such as SQL). Newer types of databases such as MongoDB allow for an in-between type of data, semi-structured, which could, for instance, be an XML file with well-defined fields that may differ between instances [51]. A database specialist should be engaged to

make the right decision on which system to implement. This choice should take into account data management principles such as the CAP theorem and PACELC theorem [52]. The latter refers to the trade-off that different database management systems make between data consistency and data availability in the case of a partition. For more details on the challenges of big data, we refer the reader to Chen *et al.* [53].

### C. THE AI TEAM

When preparing to integrate AI technology into an organization, one should first put together an AI team. In what follows I will discuss some considerations about what the qualities of the AI team may be, as well as possible hiring strategies.

#### 1) HIRING CONSIDERATIONS

Recently, the demand for technical AI talent has risen starkly [54], [55]. Tech giants such as Google and Amazon provide attractive work environments (and compensation) for AI experts. In addition to competing with the tech giants, numerous AI startups are also in the market for hiring from the scarce AI talent pool [56]. A survey by Rackspace Technology [1] uncovers that 34% of failed AI research and development projects are due to a lack of expertise within the organization and 28% due to a lack of investment in the right people. Before delving into strategies of hiring and keeping AI talent, it is pertinent to first explore what to look for in an employee.

Developing AI systems is technologically challenging as it involves the use of cutting-edge, complex models and architectures. In addition, it necessitates ensuring that the required performance is achieved, both in terms of computing time as well as accuracy, and that the system is scalable. Implementing AI systems includes scientific challenges which require both a domain-knowledge and a technical (AI) skill set. Some of the required skills in an AI development team include [57], [58]:

- Mathematics and statistics: to properly understand neural network technologies and built novel solutions, a deep understanding of mathematics is essential, often at the PhD level.
- Data mining: detecting patterns in large datasets, not just through running a simple model, but by intelligent feature engineering and thinking about what data may be useful.
- Big data engineering: handling and storing data, being aware of privacy and security of the data, scaling up models robustly.
- Pattern extraction: it is not just the models that perform pattern extraction, but the developers need to be able to communicate complex technology to less technical supervisors or stakeholders.
- Visualisation and user interface (UI) design: so that data can be more easily understood.
- Software engineering: developing models in such a way that they are optimized for speed, they are secure, and can scale up.

- Communicating: across a team that has members from a different background.
- Continuous learning: new AI technologies are constantly being developed. To stay up to date they have to be curious and motivated.

Given all these different skills required for the development of AI systems, there is a need for a diverse, multidisciplinary project team. These days, AI talent is scarce [59] and in high demand [54]. Hence, organizations must come up with a strategy to attract and keep AI experts.

One possible strategy for obtaining AI talent, other than hiring, is to up-skill existing employees, in technical AI skills as well as vertical AI knowledge, i.e. how to build solutions that target specific customer or business needs [17]. Through internal education, such as workshops and courses, existing employees can be pivoted to become AI experts. Another strategy is to target recent AI graduates. Due to the scientific challenges involved in developing AI models, PhD level developers may be desired. One way to achieve this is by collaborating closely with academic institutions, or by relocating to strategic places with ample PhD-level AI experts such as the San Francisco Bay area and Boston.

#### 2) ACQUIHIRE

There are alternatives to the hiring or up-skilling approach. A possible strategy is that of the acquihire. In an acquihire, a company buys another company, not so much for the product or service they provide, but for the manpower that comes along with the acquisition [60], [61]. An example of this is the acquisition of Android by Google in 2005, before any product was developed and when Android still employed less than ten people. The employees had experience designing mobile devices as well as excellent ideas [60]. In addition, by centering the AI development in the acquired company, the risk of failure can be kept separate from the main business (see also next Section). The strategy of acquihiring has been heavily used by some of the leaders in the AI race: the FAMGA companies (Facebook, Amazon, Microsoft, Google, Apple) [59].

Once the AI team has been built, an important consideration is to focus on retaining employees through continuous education, benefits, and a desired working environment. In conclusion, the importance of finding trained, capable and diligent manpower for building AI models cannot be overestimated, given the scientific approach that needs to be taken during development.

### D. ORGANIZING AI DEVELOPMENT

Once an AI team has been put together, it is imperative to consider how to position it in the organization. Given the difficulties mentioned above in attracting AI talent, alternatives to in-house development can be considered such as AI as a service (AIaaS).

#### 1) POSITIONING THE AI TEAM IN THE ORGANIZATION

The management of an AI project can typically be centralized, or decentralized. According to Andrew Ng, a leader in

the AI field, a centralized AI team is perfectly positioned to develop an organization-wide AI system [62]. To develop an AI application, experts need knowledge about multiple aspects of the organization, and need to understand and access data from many departments. In a centralized approach, a chief AI officer might be appointed, who can work across departments with business leaders to drive AI projects forward. Centralized teams can take many forms, such as centres of excellence or research labs. The latter often has strong ties to universities. In Section III-B, the importance of enabling access to organization-wide data was mentioned. This cross-functionality within the AI team will be crucial for them to develop AI applications by working with all business leaders. Thinking about how to best integrate the AI team in the organization is critical to avoid delay. The decision power of managers and teams will need to be considered, i.e. if any departments need to be involved in decision making, they should be included in the team. A short reporting line will speed up development [17].

Others, like Neil Jacobstein (as per Boulton [63]) argue for a decentralized approach. In this setup, each department or business team is responsible for developing their own AI systems, supported by the CEO. Fountaine *et al.* [64] emphasise that there is no one way of organizing the AI team, as it depends on the organization's individual needs. They refer to the AI team in the centralized approach as a 'hub', and business units in the decentralized approach as 'spokes'. A third, hybrid approach is pointed out by Fountaine *et al.* [64], the 'hub-and-spoke' approach, whereby the AI and data analytics capabilities are distributed over both the hub and the spokes. Companies who have done well with this latter approach include Amazon and Google [65]. In the end, the organization's decision-making structure should be considered. The best approach, given the organization's needs and readiness should be chosen, so as to best allow people to collaborate on unique multidisciplinary projects.

## 2) PORTFOLIO APPROACH

There are some similarities between AI projects and financial assets. Different projects will entail different levels of risk, due to internal knowledge, the novelty of the project, the adoption readiness of the organization, and data availability. Just like stocks, an organization can implement a portfolio of different AI projects to mitigate risks and maximize potential returns.

This portfolio can be spread across time [21]. In the short term existing, or low-risk technologies (e.g. AIaas) can be implemented to enable quick adoption. In the medium term new, emerging technologies can be explored, together with the value they may bring to the organization. As AI technology might significantly alter the way the organization operates, a culture of adoption as well as openness to change should be ensured. Finally, in the long term, a competitive advantage may be obtained through implementing risky AI projects.

To mitigate the risk, a portfolio can not only be spread over time, but also across different teams. In this approach, which is reminiscent of venture capital (VC) firms, the (centralized) AI department is given a budget that is divided across its projects or teams. These then have to develop prototypes and proof of concepts, and in each internal funding round, the organization will build upon the projects that have the biggest potential, as well as explore new ones.

## 3) AlaaS

While it is certainly an option to build AI systems within the organization, the adage 'do not reinvent the wheel' applies here. Specialized AI companies may have already developed a piece of software that can be useful for the organization. In this case, it is worth considering a Software as a Service (SaaS) or AI as a Service (AIaaS) approach [66]. Especially in the case of non-core technologies, it may not make sense to invest in developing cutting-edge AI models. For instance, Google Maps API used by an insurance company to create a daily schedule of their salesmen. The company does not need to become an expert in routing algorithms, yet it can leverage existing technologies as a service. One thing to keep in mind, is that the organization would be dependant upon potential external price changes of the AIaaS.

## 4) DEVELOPING IN-HOUSE OR NOT?

Even if the technology does not exist yet, there is an alternative to in-house development. The organization can engage an external development company (i.e., outsource the burden of development and maintenance). According to a recent survey by Rackspace Technology [1], 41% of the interviewed companies work with trusted partners/specialist vendors to develop AI applications, whereas 38% develop them solely in-house.

The decision of whether to develop in-house or not depends on the type of organization, and if it already has the knowledge and manpower in-house. Ransbotham *et al.* [67] identify four types of organizations depending on their AI readiness level: pioneers, investigators, experimenters, and passives. These categories have varying levels of understanding and experience with AI. Understanding which category an organization falls into, can help build understanding of the organizational context and thus decide the approach towards implementing AI technologies. While building in-house is a significant commitment, it does ensure that the organization's AI know-how does not fall behind any further.

Finally, another option is to acquire AI technology is through 'Growth by acquisition'. This strategy has the added benefit that it may allow organizations to stay the market cap leader, by buying their competitors. Four companies have been leading AI acquisitions: Google, Apple, Amazon, and Facebook [68].

## 5) AGILITY

Looking at implementation strategies, an iterative, agile approach may be best suited for developing AI

systems [69]–[71]. This involves creating prototypes first [72], which are then tested and analysed without spending too many resources. These provide steppingstones to more polished and well-designed software systems. The insights from initial user tests are then used to reduce risk, and change designs. This reduces the risk overall when scaling up the prototype. It is important to engage the stakeholders early on in the process because what the users may want may differ from the managers' opinion.

Another recent set of best practices is that of MLOps [73], which arose from the fusion of machine learning in DevOps practices [74]. The practices included in DevOps offer an agile approach to combine an organization's IT development and operations, thus allowing software to be brought into production at a higher velocity.

### E. TECHNOLOGIES

Ever since John McCarthy organized the first workshop on AI, i.e. the Dartmouth Summer Research Project on Artificial Intelligence, in the summer of 1956 at Dartmouth College [75], numerous machine learning technologies have been developed such as convolutional neural networks, word vector embeddings, reinforcement learning, and transformer networks [16], [76]. Giving an overview of these technologies is outside of the scope of this article, hence the reader is referred to Goodfellow *et al.* [16]. The reader should keep in mind though, that even though deep learning technologies are extremely popular, a hybrid approach, that integrates 'older' AI technologies such as rules, grammar, and fuzzy logic may be even more efficient [77]. In what follows, three important considerations when implementing AI techniques are discussed.

#### 1) ACCURACY VERSUS BLACK-BOX

An important business consideration regarding AI models is: does the model need to be explainable/comprehensible? This concept is also referred to as XAI [78], [79]. While deep neural networks have a reputation for being able to reach high performance in terms of prediction accuracy (given enough data), they are not interpretable. This may be an issue if the result of the model needs to be understandable. For instance, a bank doing credit scoring is legally obliged to use understandable rules [80]. In this scenario, simpler models such as decision trees will have to be used, or rule extraction should be performed on black-box-models [81]. The cost of having an explainable model is usually a lower accuracy.

#### 2) HUMAN-IN-THE-LOOP

To evaluate and measure a model's performance, ground truth labels are typically used. For example, when doing emotion prediction from music, experts need to label music files with emotions [82]. In a traditional context, supervised learning is used with a split training and test. The former is used to train the model, and the latter to evaluate it. When labelled data is scarce, recent techniques can leverage a small amount of labelled data through semi-supervised or self-supervised

learning [83]. Other strategies, such as reinforcement learning, put a 'human-in-the-loop' and allow new labelled data to be added to the model while it is in production [84].

#### 3) REPLACING OR AUGMENTING HUMANS?

Does the organization want the AI system to replace the human previously doing the task, or augment their performance? The latter is a popular approach that aims at creating collaborative/assistive systems instead of autonomous systems, with the intent to make the human's job easier and more effective. For example, one might think that music generation systems can put composers out of business [85]. In reality, those systems can empower composers, to compose quicker, inspire them with new ideas, and let them control generated music on unseen levels, such as through emotion [86].

#### 4) CLOUD VERSUS IN-HOUSE HOSTING

In Section II, one of the enablers of the current AI spring was discussed, namely the existence of fast GPUs. GPUs are processing units optimized for the matrix calculations needed when training neural networks. The GPU equipment requirements should be considered for each AI problem statement. Depending on the data size and problem score, a significant number of powerful GPUs may be needed to train the models. There are two broad strategies to follow: 1) renting cloud servers such as Elastic Compute Cloud<sup>3</sup> or Google Compute Engine,<sup>4</sup> versus 2) buying equipment and running the models in-house. A detailed cost analysis such as that proposed by Emeras *et al.* [87] can offer valuable insights for making an informed decision. Emeras *et al.* [87] point out that running the equipment onsite comes with two types of costs: CApital EXpense (capex), and OPerating EXpenses (opex). The former consists of buying machines, storage, interconnecting them, purchasing other room equipment such as cooling, and security. The latter includes the manpower to administrate and operate the machines, energy use, support, and software licences. It is no small investment to set up GPU equipment within the organization. Hence, cloud services may offer a quick and cheap path to powerful and scalable GPU power, especially in the early phases of AI projects.

Finally, organizations do not always need the most powerful GPUs in terms of processing power. Recent GPUs, like NVIDIA's Jetson [88], are ultra-low-power, small devices that can be integrated in Internet of Things (IoT) devices. Depending on the organization's products and use-cases, such devices may be considered to train AI models on small consumer IoT devices.

### F. KPIs

Like any other project, clear success measures need to be defined from the onset, so as to keep track of the project's performance. These clearly defined measures, or Key Performance Indicators (KPIs) will help navigate threats and

<sup>3</sup><https://aws.amazon.com/ec2>

<sup>4</sup><https://cloud.google.com/compute>



provide a reporting tool to management in order to justify continued investment in the project. Good success measures are linked to the organization's strategic goals, and may be reviewed often.

Key performance indicators can be used to measure how effectively the organization is reaching its goals, more specifically in this case, the goals influenced by the AI project. It is outside of the scope of this article to describe the entire process of choosing the right KPIs, for that, the interested reader is referred to [89] and [90]. It is important to point out, however, that an AI implementation might not only affect financial outcomes, but can equally result in value creation for the customer, e.g. better service such as transactional experience [91]. In the words of A. Einstein as quoted by [92]: "Strive not to be a success, but rather to be of value". Steve Jobs was another proponent of considering value over profit [93]. A similar thought was uttered by Barbara Grosz, Higgins Professor of Natural Sciences at Harvard University, as quoted in [94]: "it is important to take on the challenge of identifying success measures for AI systems by their impact on people's lives". As a result, Kiron and Schrage [95] states that KPIs can centre around customers, costs, processes, or investors. Having clear and simple KPIs, gives the development team, and the management concrete objectives to track and aim for.

In a case study of 500 companies who have adopted AI, Wamba-Taguimdje *et al.* [96] lists a large number of reported added value through AI such as 'Reduce costs by allowing increased asset utilization, as 24/7 executions are possible' in the Transport section; 'Save time, improve accuracy and help consumers to find things easily without wasting time' in Sales; 'Detect and combat fraud and money laundering' in Banking; 'Gain new customers and increase revenue with a real-time analysis solution' in Healthcare; and 'Improve the unstructured analysis of data capabilities for the existing solutions' in IT/Telecom. Some of these are value-based in terms of novel services and innovative products. Much of the created value is hidden, however, in the form of more efficient internal processes such as quicker labour time or better customer understanding which can result in better decision making [97]. Looking at the KPIs related to AI systems from the survey by Rackspace Technology [1], it becomes clear that the most two popular KPIs are profit margins (52%) and revenue growth (51%). Value-based KPIs, however, are also popular: customer satisfaction/net promoter scores (46%), process enhancement/improvement (41%), NPD (New Product Development) (27%), and time to insight (16%).

In addition to more traditional KPIs, we should keep in mind that 'success' or more concretely 'performance accuracy' is not guaranteed for AI systems due to their stochastic nature. While we can have the best intentions to create a model that can identify customers that are most likely to buy a product, the model might simply not be accurate when looking at new data. Model performance metrics such as accuracy, confusion matrices, or F1-scores give us quick insights into how the model development is advancing and

how much potential it will have to be successfully integrated into the business. It is dangerous to just assume that a model will perform well, instead it is better to be open to adapt and willing to pivot in case the data simply does not have any detectable patterns, hence, metrics can provide us with useful insights.

Finally, while KPIs are useful to measure the performance and impact of AI systems, AI can equally be leveraged to calculate better and more insightful KPIs. For instance, Natural Language Processing systems can be used to get real-time feedback on the sentiment of customers [98]. Business intelligence systems can be enhanced by AI, and provide easy to interpret metrics that empower decision-makers.

## G. CHOOSING A RISK LEVEL

### 1) RISKS

Risk management is essential when managing any project [99]. Unlike traditional software, where detailed use case diagrams can describe an entire system in a fairly deterministic and established way [100], the nature of AI algorithms creates particular kinds of risks and challenges.

#### a: STOCHASTIC NATURE

When creating AI software, be it a voice assistant, an autonomous driving system or a music album cover generator, the developed models always need to be tested with new, unseen data. A manager can 'order' developers to make a model that predicts the customers' next action or to generate music that matches a video, however, there is no guarantee that the model will work accurately. Through several iterations, the development team will need to test how the model performs on different, new datasets. During this process, the model architectures is often adjusted. In addition, since the technology is still rapidly advancing, time needs to be reserved for the continuous education of developers and testing of novel/better technologies that may come out mid-project. Such uncertainties need to be taken into account during development as they can create longer development time, a need for more/better data, research time, and even the need to pivot to a more feasible problem statement. Given this uncertainty, it makes sense to develop a prototype first, before investing a lot in an AI idea. Organizations can also look at competitors to see if they have working models. Both approaches can give an idea as to how feasible the desired task is and alert the organization on time if a pivot is needed.

#### b: BIASES AND ETHICS

As machine learning models learn from data, they are only as good as the data they learn from. The models should therefore be tested for issues such as 'learned biases' [57]. A famous example of a biased model is the AI recruiting model from Amazon, which was biased against women [101]. Such biases arise because historical datasets are imperfect and contain human biases. In this case, there had been a historical hiring preference for men, which was reflected

in the AI predictions. Great care should be taken when building AI models that might inadvertently capture such hidden, unethical biases. These biases are not limited to the field of recruiting. For instance, in healthcare, researchers have established a correlation between wealth, ethnicity, and health [102]. In Rackspace Technology [1]'s survey, 27% of failed or discarded AI research and development projects were due to a bias in the algorithm. There is an ethical responsibility that comes with utilizing datasets that contain discrimination based on wealth or ethnicity.

Thoroughly testing AI technologies for biases makes the technology more ethically acceptable. This sort of testing goes beyond looking at loss functions and will require a domain expert or ethics expert to dig up potential problem areas for review and stress testing. There are a few other considerations to make regarding ethics, including ethical dilemmas such as the famous self-driving car example where a vehicle needs to choose between hitting an old lady or a young child. These are its only two options and they have an equal probability. Do you let the algorithm decide, or do you set rules?

#### c: SECURITY

In addition to traditional security issues related to data, a particular kind of security concern that arises only with AI models is related to adversarial attacks [103]. In such an attack, the user feeds deceptive input to the system to cause a malfunction in the machine learning model and thus fool it. Like many things in AI, the idea came up in movies first, e.g. in *Zero History* (2010), the character wears a t-shirt that is decorated such that it renders him invisible to surveillance technology. A few years later, in the real world, Xu *et al.* [104] designed an 'adversarial t-shirt', with up to 74% success rate. Their shirt introduced 'adversarial patches' in the image, that fool person detectors. Adversarial attacks do not just happen on image recognition systems, even simple linear spam classifiers can fall prey to evasion attacks, by spammers who insert 'good words' into their spam emails [105]. More detailed information and examples on adversarial attacks can be found in [106].

A separate category of AI systems consists of those used for biometric identification such as face and voice identification systems. They have increased security concerns and will be the target of different types of spoofing algorithms. For instance, voice recognition systems might encounter playback attacks, whereby the perpetrator plays a recorded voice of the rightful user. Spoofing detection AI algorithms try to counter these attacks [107]. Since these are highly specialized systems, the reader is referred to [108] for more in-depth information.

#### d: OTHER STRATEGIC DECISIONS

Risk can also result from the chosen strategy, as discussed in the previous sections. Are we using new, hard to implement technologies? Will the system be developed in-house and does the organization have the necessary know-how to do so?

If AIaas will be used, is there a price guarantee for the future? Is the AI team integrated into the organization, and do they have access to data across departments? Is a new task being tackled that nobody has tried before, or is it an established task, like face identification? Finally, for each identified AI system that the organization wants to implement, the trade-off between risk and benefits can be plotted in a graph for each potential project.

#### 2) BENEFITS

Any digital business strategy should keep value-creation in mind [109]. The aim of developing any AI system should be to create value for either the customer or internally for the business. In this phase, the organization should explore what value will be created by each of the AI projects. Examples of added value include a more secure banking experience, personalized shopping recommendations, and health alerts based on biometric feedback from a smartwatch. Ideally, these are captured by the KPIs, but they may be broader and harder to measure.

#### 3) SWOT

Just like when weighing financial investments, one can opt to implement high-risk, high-value AI models (e.g. highly innovative services and products), or choose to stick with low-risk, low-value models (e.g. automatic internal repetitive processes). Alternatively, some technologies offer a blend of the two. High-value, low-risk models include those that automatise complicated mathematical decision support systems; and low-value, high-risk include those that try to predict complex customer behaviour.

Depending on the risk strategy of the organization, a portfolio of different models may be adopted to have a spread of risk, as discussed above in the portfolio approach. The information gathered in the aiSTROM roadmap framework can be used in a Strengths/weaknesses/opportunities/threats (SWOT) analysis [110] to assist managers to create an informed decision and put together a balanced portfolio. The risk profile of the organization will also affect its implementation strategy as discussed below.

#### H. ENABLING A CULTURAL SHIFT THROUGH EDUCATION

The employees directly involved with the AI implementation, either as managers or developers, should be knowledgeable of AI technologies. The technical experts in the AI team can also be part of the team that ensures the continuous education of not only the employees involved in the project but the rest of the organization as well. Other employees, as far removed from the AI project as they might be, still contribute to the culture of the organization and thus help form the vision of the products or services, and they may even be the user of the AI system if it is related to internal processes.

One way to achieve this internal education is by setting up a centre of excellence within the organization, which is responsible for creating awareness and knowledge of AI technologies within the organization [111]. According to Wilson

Pang, Chief Technology Officer of Appen [112], these centres can be seen as think tanks, run by experts in AI. Having an AI centre of excellence can help enable a shift of the organization's culture to be inclusive of AI thinking, and hence let people embrace AI. When employees are educated, they can be involved in brainstorming workshops to help think about how to embed AI technologies in the organization. Through this, the organization is not just leveraging AI for competitive advantage, but embracing these technologies as a driving force of the organization, and letting go of legacy systems. Through education, employees can be empowered to be part of the AI strategy of the organization.

In addition, when AI literacy is stimulated among employees, they may replace their potential fears of being replaced, by the knowledge that AI can empower them (as discussed in Section III-E), and how they can work in conjunction with the system in a symbiotic way. This promotes a culture of adoption. Fountaine *et al.* [64] point out that 'only 8% of firms are engaging in core practices that support widespread adoption', which leads to a failure when the organization tries to scale up AI technologies. Creating an AI culture that supports broad organization-wide adoption is essential.

Finally, it is worth noting that employee training is not a one-time thing. AI technologies constantly change, and hence there should be a culture of continuous education. This can work beneficial towards keeping employees engaged with the AI integration in the organization and help facilitate a cultural shift of AI adoption.

#### IV. CONCLUSION

This perspective paper presents aiSTROM, a strategic framework that guides managers through the process of integrating AI in the organization (Figure 1). aiSTROM provides managers with a framework to think strategically about the challenges involved in AI projects. The first step in the proposed framework is identifying a top  $n$  list of potential problem statements/tasks that can be tackled with AI. In the following steps, the potential projects on this list will be analysed according to different factors.

Data is the first important factor as it is the throbbing heart of any machine learning system. Strategic issues that are important to keep in mind when selecting data sources are discussed, together with legal issues associated with them, and options related to secure storage. Secondly, due to the scarcity of AI talent, important challenges and options related to composing an AI team are examined. In addition, given the interdisciplinary of AI projects, some consideration is given to where to position the team in the organization, and if the organization should consider alternative solutions such as AIaaS and acquisitions.

In the aiSTROM framework, several considerations related to technology are taken into account. These include black-box versus explainable models, human-in-the-loop models, replacing versus empowering human users, and finally, issues related to GPU equipment. Like any project, AI projects also need proper, value-based KPI metrics to help

measure and justify their progress to management. Finally, risks and benefits that need to be considered when performing a SWOT analysis are discussed, and a case is made for the importance of enabling a cultural shift within the organization to support widespread adoption.

Armed with a thorough analysis of the potential AI projects resulting from aiSTROM, managers can make informed decisions, that are bound to increase the success of the organization's AI strategy.

In future research, the proposed aiSTROM framework should be validated in real-world situations, so as to go beyond relying on data and conclusions drawn from existing surveys and literature.

#### REFERENCES

- [1] Rackspace Technology. (2021). *Global Report: Are Organizations Succeeding at AI and Machine Learning?* [Online]. Available: <https://www.rackspace.com/solve/succeeding-ai-ml/ty>
- [2] S. Kuutti, R. Bowden, Y. Jin, P. Barber, and S. Fallah, "A survey of deep learning applications to autonomous vehicle control," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 2, pp. 712–733, Feb. 2021.
- [3] D. Herremans and C.-H. Chuan, "The emergence of deep learning: New opportunities for music and audio technologies," *Neural Comput. Appl.*, vol. 32, no. 4, pp. 913–914, Feb. 2020.
- [4] M. Alloghani, D. Al-Jumeily, A. J. Aljaaf, M. Khalaf, J. Mustafina, and S. Y. Tan, "The application of artificial intelligence technology in healthcare: A systematic review," in *Proc. Int. Conf. Appl. Comput. Support Ind., Innov. Technol.*, vol. 1174. Cham, Switzerland: Springer, 2019, pp. 248–261, doi: 10.1007/978-3-030-38752-5\_20.
- [5] A. Burgess, "Industrialising AI," in *The Executive Guide to Artificial Intelligence*. Cham, Switzerland: Springer, 2018, pp. 147–164.
- [6] *Managing AI and ML in the Enterprise*, TechRepublic, Mapusa, India, 2019.
- [7] R. G. Smith and J. Eckroth, "Building AI applications: Yesterday, today, and tomorrow," *AI Mag.*, vol. 38, no. 1, pp. 6–22, Mar. 2017.
- [8] B. G. Buchanan, "A (very) brief history of artificial intelligence," *AI Mag.*, vol. 26, no. 4, p. 53, 2005.
- [9] J. Liu, X. Kong, F. Xia, X. Bai, L. Wang, Q. Qing, and I. Lee, "Artificial intelligence in the 21st century," *IEEE Access*, vol. 6, pp. 34403–34421, 2018.
- [10] M. Mitchell, "Self-awareness and control in decentralized systems," in *Proc. AAAI Spring Symp., Metacogn. Comput.*, 2005, pp. 80–85.
- [11] J. McCarthy and P. J. Hayes, "Some philosophical problems from the standpoint of artificial intelligence," in *Readings in Artificial Intelligence*. Amsterdam, The Netherlands: Elsevier, 1981, pp. 431–450.
- [12] A. Lieto, M. Bhatt, A. Oltramari, and D. Vernon, "The role of cognitive architectures in general artificial intelligence," *Cogn. Syst. Res.*, vol. 48, pp. 1–3, May 2018.
- [13] W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bull. Math. Biophys.*, vol. 5, no. 4, pp. 115–133, 1943.
- [14] T. Gonsalves, "The summers and winters of artificial intelligence," in *Advanced Methodologies and Technologies in Artificial Intelligence, Computer Simulation, and Human-Computer Interaction*. Hershey, PA, USA: IGI Global, 2019, pp. 168–179.
- [15] J. Hendler, "Avoiding another AI winter," *IEEE Intell. Syst.*, vol. 23, no. 2, pp. 2–4, Mar./Apr. 2008.
- [16] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep Learning*, vol. 1. Cambridge, MA, USA: MIT Press, 2016.
- [17] M. Baccala, C. Curran, D. Garrett, S. Likens, A. Rao, A. Ruggles, and M. Shehab, "AI predictions 8 insights to shape business strategy," PwC, London, U.K., Tech. Rep. [pwc.com/us/AI2018](http://pwc.com/us/AI2018), 2018.
- [18] J. Furman and R. Seamans, "AI and the economy," *Innov. Policy Econ.*, vol. 19, no. 1, pp. 161–191, Jan. 2019.
- [19] Accenture. (2016). *Why AI is the Future of Growth—Country Spotlights*. [Online]. Available: [https://www.accenture.com/ve-es/\\_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth-Country-Spotlights.pdf](https://www.accenture.com/ve-es/_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth-Country-Spotlights.pdf)
- [20] A. Rowsell-Jones and C. Howard. (2019). *2019 CIO Survey: CIOs Have Awaken to the Importance Of AI*. [Online]. Available: <https://www.gartner.com/en/documents/3897266/2019-cio-survey-cios-have-awaken-to-the-importance-of-ai>

- [21] J. Bughin, J. Seong, J. Manyika, M. Chui, and R. Joshi, "Notes from the AI frontier: Modeling the impact of AI on the world economy," McKinsey Global Inst., New York, NY, USA, Tech. Rep., 2018.
- [22] Y. Zhao, B. Li, X. Li, W. Liu, and S. Ren, "Customer churn prediction using improved one-class support vector machine," in *Proc. Int. Conf. Adv. Data Mining Appl.* Berlin, Germany: Springer, 2005, pp. 300–306.
- [23] P. Miller, "Web 2.0: Building the new library," *Ariadne*, vol. 45, no. 30, p. 10, 2005.
- [24] M. J. Grant and A. Booth, "A typology of reviews: An analysis of 14 review types and associated methodologies," *Health Inf. Libraries J.*, vol. 26, no. 2, pp. 91–108, Jun. 2009.
- [25] P. P. Tallon, K. L. Kraemer, and V. Gurbaxani, "Executives' perceptions of the business value of information technology: A process-oriented approach," *J. Manage. Inf. Syst.*, vol. 16, no. 4, pp. 145–173, Mar. 2000.
- [26] T. Morgan, *Business Rules and Information Systems: Aligning IT With Business Goals*. Reading, MA, USA: Addison-Wesley, 2002.
- [27] S. Ransbotham, P. Gerbert, M. Reeves, D. Kiron, and M. Spira, "Artificial intelligence in business gets real," MIT Sloan Manage. Rev., Boston Consulting Group, Boston, MA, USA, Tech. Rep. 60280, 2018.
- [28] *Predictions for the \$203 Billion Big Data Analytics Market*, Obtido em, Gil Press, Jersey City, NJ, USA, Jan. 2017, p. 6, vol. 30. [Online]. Available: <https://www.forbes.com/sites/gilpress/2017/01/20/6-predictions-for-the-203-billion-big-data-analytics-market/?sh=32a05a052083>
- [29] M. Cox and D. Ellsworth, "Application-controlled demand paging for out-of-core visualization," in *Proc. 8th Conf. Vis. (VIS)*. Washington, DC, USA: IEEE Computer Society, 1997, p. 235.
- [30] P. Russom, "Big data analytics," *TDWI Best Practices Rep.*, vol. 19, no. 4, pp. 1–34, 2011.
- [31] D. Laney, "3D data management: Controlling data volume, velocity and variety," *META Group Res. Note*, vol. 6, no. 70, p. 1, 2001.
- [32] J. Gantz and D. Reinsel, "The digital universe in 2020: Big data, bigger digital shadows, and biggest growth in the far east," *IDC iView, IDC Analyze Future*, vol. 2007, no. 2012, pp. 1–16, 2012.
- [33] M. White, "Digital workplaces: Vision and reality," *Bus. Inf. Rev.*, vol. 29, no. 4, pp. 205–214, Dec. 2012.
- [34] P. Tabesh, E. Mousavidin, and S. Hasani, "Implementing big data strategies: A managerial perspective," *Bus. Horizons*, vol. 62, no. 3, pp. 347–358, May 2019.
- [35] T. Davenport, *Big Data at Work: Dispelling the Myths, Uncovering the Opportunities*. New York, NY, USA: Harvard Business Review Press, 2014.
- [36] E. Junqué de Fortuny, D. Martens, and F. Provost, "Predictive modeling with big data: Is bigger really better?" *Big Data*, vol. 1, no. 4, pp. 215–226, Dec. 2013.
- [37] M. Brown, K. M. Bossley, D. J. Mills, and C. J. Harris, "High dimensional neurofuzzy systems: Overcoming the curse of dimensionality," in *Proc. IEEE Int. Conf. Fuzzy Syst.*, Mar. 1995, pp. 2139–2146.
- [38] R. Bellman, *Adaptive Control Processes: A Guided Tour*, vol. 3. Princeton, NJ, USA: Princeton Univ. Press, 1961, p. 2.
- [39] H. Aguinis, I. Villamor, and R. S. Ramani, "MTurk research: Review and recommendations," *J. Manage.*, vol. 47, no. 4, pp. 823–837, Apr. 2021, doi: [10.1177/0149206320969787](https://doi.org/10.1177/0149206320969787).
- [40] W. Wang and Z.-H. Zhou, "Crowdsourcing label quality: A theoretical analysis," *Sci. China Inf. Sci.*, vol. 58, no. 11, pp. 1–12, Nov. 2015.
- [41] V. Mayer-Schönberger and K. Kukier, *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. Boston, MA, USA: Houghton Mifflin Harcourt, 2013.
- [42] A. Azevedo and M. F. Santos, "KDD, SEMMA and CRISP-DM: A parallel overview," in *Proc. IADIS Eur. Conf. Data Mining (MCCSIS)*, Amsterdam, The Netherlands, 2008, pp. 182–185.
- [43] J. Hinds, E. J. Williams, and A. N. Joinson, "'It wouldn't happen to me': Privacy concerns and perspectives following the Cambridge analytica scandal," *Int. J. Hum.-Comput. Stud.*, vol. 143, Nov. 2020, Art. no. 102498.
- [44] A. S. Figueiredo, "Data sharing: Convert challenges into opportunities," *Frontiers Public Health*, vol. 5, p. 327, Dec. 2017.
- [45] C. Cross, M. Parker, and D. Sansom, "Media discourses surrounding 'non-ideal' victims: The case of the Ashley Madison data breach," *Int. Rev. Victimol.*, vol. 25, no. 1, pp. 53–69, Jan. 2019.
- [46] C. J. D'Orazio, K.-K. R. Choo, and L. T. Yang, "Data exfiltration from Internet of Things devices: iOS devices as case studies," *IEEE Internet Things J.*, vol. 4, no. 2, pp. 524–535, Apr. 2017.
- [47] L. de la Torre, "A guide to the California consumer privacy act of 2018," SSRN, Rochester, NY, USA, Nov. 2018. [Online]. Available: <https://ssrn.com/abstract=3275571>
- [48] B. Wong YongQuan, "Data privacy law in Singapore: The personal data protection act 2012," *Int. Data Privacy Law*, vol. 7, no. 4, pp. 287–302, Nov. 2017.
- [49] P. Voigt and A. Von dem Bussche, *The EU General Data Protection Regulation (GDPR): A Practical Guide*, 1st ed. Cham, Switzerland: Springer, 2017.
- [50] J. S. Plank, A. Bassi, M. Beck, T. Moore, M. Swany, and R. Wolski, "Managing data storage in the network," *IEEE Internet Comput.*, vol. 5, no. 5, pp. 50–58, Sep/Oct. 2001.
- [51] K. Chodorow, *MongoDB: The Definitive Guide: Powerful and Scalable Data Storage*. Sebastopol, CA, USA: O'Reilly Media, 2013.
- [52] D. Abadi, "Consistency tradeoffs in modern distributed database system design: CAP is only part of the story," *Computer*, vol. 45, no. 2, pp. 37–42, Feb. 2012.
- [53] M. Chen, S. Mao, and Y. Liu, "Big data: A survey," *Mobile Netw. Appl.*, vol. 19, no. 2, pp. 171–209, 2014.
- [54] J. Garfinkel, "Gartner survey of more than 3,000 CIOs reveals that enterprises are entering the third era of IT," Gartner.com, Stamford, CT, USA, Tech. Rep., 2020. [Online]. Available: <https://www.gartner.com/en/newsroom/press-releases/2018-10-16-gartner-survey-of-more-than-3000-cios-reveals-that-enterprises-are-entering-the-third-era-of-it>
- [55] A. Cassard and J. Hamel, "Exponential growth of technology and the impact on economic jobs and teachings: Change by assimilation," *J. Appl. Bus. Econ.*, vol. 20, no. 2, pp. 1–6, 2018.
- [56] TalentSeer. (2020). *2020 AI Talent Report: Current Landscape & Market Trends*. [Online]. Available: <https://www.talentseer.com/2020-ai-talent-report>
- [57] S. Dube, "AI and business," in *An Intuitive Exploration of Artificial Intelligence*. Cham, Switzerland: Springer, 2021, pp. 261–266.
- [58] A. Verma, K. Lamsal, and P. Verma, "An investigation of skill requirements in artificial intelligence and machine learning job advertisements," *Ind. Higher Educ.*, Feb. 2021, Art. no. 095042221990990. [Online]. Available: <https://journals.sagepub.com/doi/full/10.1177/095042221990990>, doi: [10.1177/095042221990990](https://doi.org/10.1177/095042221990990).
- [59] *AI Trends Report*, CBInsights, New York, NY, USA, 2019.
- [60] H. R. Varian, "Seven deadly sins of tech?" *Inf. Econ. Policy*, vol. 54, Mar. 2021, Art. no. 100893.
- [61] R. Fantasia, "Acquihiring: A new process for innovation and organizational learning," in *Information and Communication Technologies in Organizations and Society*. Cham, Switzerland: Springer, 2016, pp. 205–214.
- [62] *How to Build a Great AI Team*, Forbes Insights, Jersey City, NJ, USA, 2019.
- [63] C. Boulton. (Mar. 2017). *Think Twice Before You Hire a Chief AI Officer*. [Online]. Available: <https://www.cio.com/article/3177427/think-twice-before-you-hire-a-chief-ai-officer.html>
- [64] T. Fountaine, B. McCarthy, and T. Saleh, "Building the AI-powered organization," *Harvard Bus. Rev.*, vol. 97, no. 4, pp. 62–73, 2019.
- [65] International Workplace Hub. (Oct. 2020). *Five Companies Using the Hub- and-Spoke Model*. [Online]. Available: <https://www.iwgplc.com/MediaCentre/Article/five-companies-using-the-hub-and-spoke-model>
- [66] K. D. Pandl, H. Teigeler, S. Lins, S. Thiebes, and A. Sunyaev, "Drivers and inhibitors for organizations' intention to adopt artificial intelligence as a service," in *Proc. 54th Hawaii Int. Conf. Syst. Sci.*, 2021, p. 1769.
- [67] S. Ransbotham, D. Kiron, P. Gerbert, and M. Reeves, "Reshaping business with artificial intelligence: Closing the gap between ambition and action," *MIT Sloan Manage. Rev.*, vol. 59, no. 1, pp. 1–17, 2017.
- [68] S. Makridakis, "The forthcoming artificial intelligence (AI) revolution: Its impact on society and firms," *Futures*, vol. 90, pp. 46–60, Jun. 2017.
- [69] A. M. Sharif, "Harnessing agile concepts for the development of intelligent systems," *New Gener. Comput.*, vol. 17, no. 4, pp. 369–380, Dec. 1999.
- [70] R. Hoda, N. Salleh, and J. Grundy, "The rise and evolution of agile software development," *IEEE Softw.*, vol. 35, no. 5, pp. 58–63, Sep. 2018.
- [71] S. Leijnen, H. Aldewereld, R. van Belkom, R. Bijvank, and R. Ossewaarde, "An agile framework for trustworthy AI," in *Proc. NeHuAI 24th Eur. Conf. Artif. Intell. (ECAI)*, Santiago de Compostela, Spain, 2020, pp. 75–78.

- [72] J. Shore, *The Art of Agile Development: Pragmatic Guide to Agile Software Development*. Sebastopol, CA, USA: O'Reilly Media, 2007.
- [73] S. Alla and S. K. Adari, "What is MLOps?" in *Beginning MLOps With MLFlow*. Berkeley, CA, USA: Apress, 2021, pp. 79–124.
- [74] I. Karamitsos, S. Albarhami, and C. Apostolopoulos, "Applying DevOps practices of continuous automation for machine learning," *Information*, vol. 11, no. 7, p. 363, Jul. 2020.
- [75] A. Kaplan and M. Haenlein, "Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence," *Bus. Horizons*, vol. 62, no. 1, pp. 15–25, 2019.
- [76] T. Young, D. Hazarika, S. Poria, and E. Cambria, "Recent trends in deep learning based natural language processing," *IEEE Comput. Intell. Mag.*, vol. 13, no. 3, pp. 55–75, Aug. 2018.
- [77] K. Dastipour, M. Gogate, J. Li, F. Jiang, B. Kong, and A. Hussain, "A hybrid Persian sentiment analysis framework: Integrating dependency grammar based rules and deep neural networks," *Neurocomputing*, vol. 380, pp. 1–10, Mar. 2020.
- [78] D. Castelvecchi, "Can we open the black box of AI?" *Nature News*, vol. 538, no. 7623, p. 20, 2016.
- [79] A. Rai, "Explainable AI: From black box to glass box," *J. Acad. Marketing Sci.*, vol. 48, no. 1, pp. 137–141, Jan. 2020.
- [80] B. Waltl and R. Vogl, "Explainable artificial intelligence: The new frontier in legal informatics," *Juslett. IT*, vol. 4, pp. 1–10, Feb. 2018.
- [81] D. Martens, B. Baesens, T. Van Gestel, and J. Vanthienen, "Comprehensible credit scoring models using rule extraction from support vector machines," *Eur. J. Oper. Res.*, vol. 183, no. 3, pp. 1466–1476, 2007.
- [82] K. W. Cheuk, Y.-J. Luo, B. T. Balamurali, G. Roig, and D. Herremans, "Regression-based music emotion prediction using triplet neural networks," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2020, pp. 1–7.
- [83] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "ALBERT: A lite BERT for self-supervised learning of language representations," 2019, *arXiv:1909.11942*.
- [84] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 2018.
- [85] D. Herremans, C.-H. Chuan, and E. Chew, "A functional taxonomy of music generation systems," *ACM Comput. Surv.*, vol. 50, no. 5, pp. 1–30, Nov. 2017.
- [86] D. Herremans and E. Chew, "Morpheus: Generating structured music with constrained patterns and tension," *IEEE Trans. Affect. Comput.*, vol. 10, no. 4, pp. 510–523, Oct. 2019.
- [87] J. Emeras, S. Varrette, V. Plugaru, and P. Bouvry, "Amazon elastic compute cloud (EC2) versus in-house HPC platform: A cost analysis," *IEEE Trans. Cloud Comput.*, vol. 7, no. 2, pp. 456–468, Apr. 2019.
- [88] J. Tang, D. Sun, S. Liu, and J.-L. Gaudiot, "Enabling deep learning on IoT devices," *Computer*, vol. 50, no. 10, pp. 92–96, 2017.
- [89] D. Parmenter, *Key Performance Indicators: Developing, Implementing, and Using Winning KPIs*. Hoboken, NJ, USA: Wiley, 2015.
- [90] M. Badawy, A. A. El-Aziz, A. M. Idress, H. Hefny, and S. Hossam, "A survey on exploring key performance indicators," *Future Comput. Informat. J.*, vol. 1, nos. 1–2, pp. 47–52, Dec. 2016.
- [91] R. Amit and C. Zott, "Value creation in e-business," *Strategic Manage. J.*, vol. 22, nos. 6–7, pp. 493–520, 2001.
- [92] P. E. Juras, "Evaluating success," *Strategic Finance*, vol. 102, no. 1, p. 6, 2020.
- [93] W. Isaacson, "The real leadership lessons of Steve jobs," *Harvard Bus. Rev.*, vol. 90, no. 4, pp. 92–102, 2012.
- [94] T. Hornigold. (Apr. 2018). *How Fast is AI Progressing? Stanford's New Report Card for Artificial Intelligence*. [Online]. Available: <https://singularityhub.com/2018/01/18/how-fast-is-ai-progressing-stanfords-new-report-card-for-artificial-intelligence/>
- [95] D. Kiron and M. Schrage, "Strategy for and with AI," *MIT Sloan Manage. Rev. Mag.*, vol. 60, no. 4, pp. 30–35, 2019.
- [96] S.-L. Wamba-Taguimdje, S. F. Wamba, J. R. K. Kamdjoug, and C. E. T. Wanko, "Influence of artificial intelligence (AI) on firm performance: The business value of AI-based transformation projects," *Bus. Process Manage. J.*, vol. 26, no. 7, pp. 1893–1924, May 2020.
- [97] T. H. Davenport, *The AI Advantage: How to Put the Artificial Intelligence Revolution to Work*. Cambridge, MA, USA: MIT Press, 2018.
- [98] E. Cambria, S. Poria, A. Gelbukh, and M. Thelwall, "Sentiment analysis is a big suitcase," *IEEE Intell. Syst.*, vol. 32, no. 6, pp. 74–80, Nov./Dec. 2017.
- [99] T. Cooke-Davies, "The 'real' success factors on projects," *Int. J. Project Manage.*, vol. 20, no. 3, pp. 185–190, 2002.
- [100] J. L. Whitten and L. D. Bentley, *Systems Analysis and Design Methods*. New York, NY, USA: McGraw-Hill, 1997.
- [101] L. Yarger, F. C. Payton, and B. Neupane, "Algorithmic equity in the hiring of underrepresented IT job candidates," *Online Inf. Rev.*, vol. 44, no. 2, pp. 383–395, Dec. 2019.
- [102] D. C. Hague, "Benefits, pitfalls, and potential bias in health care AI," *North Carolina Med. J.*, vol. 80, no. 4, pp. 219–223, Jul. 2019.
- [103] N. Akhtar and A. Mian, "Threat of adversarial attacks on deep learning in computer vision: A survey," *IEEE Access*, vol. 6, pp. 14410–14430, 2018.
- [104] K. Xu, G. Zhang, S. Liu, Q. Fan, M. Sun, H. Chen, P.-Y. Chen, Y. Wang, and X. Lin, "Adversarial T-shirt! Evading person detectors in a physical world," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2020, pp. 665–681.
- [105] N. Dalvi, P. Domingos, S. Sanghai, and D. Verma, "Adversarial classification," in *Proc. 10th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2004, pp. 99–108.
- [106] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," 2014, *arXiv:1412.6572*.
- [107] B. T. Balamurali, K. E. Lin, S. Lui, J.-M. Chen, and D. Herremans, "Toward robust audio spoofing detection: A detailed comparison of traditional and learned features," *IEEE Access*, vol. 7, pp. 84229–84241, 2019.
- [108] L. M. Dinca and G. P. Hancke, "The fall of one, the rise of many: A survey on multi-biometric fusion methods," *IEEE Access*, vol. 5, pp. 6247–6289, 2017.
- [109] A. Bharadwaj, O. A. El Sawy, P. A. Pavlou, and N. Venkatraman, "Digital business strategy: Toward a next generation of insights," *MIS Quart.*, vol. 37, no. 2, pp. 471–482, Feb. 2013.
- [110] E. Gürel and M. Tat, "SWOT analysis: A theoretical review," *J. Int. Social Res.*, vol. 10, no. 51, pp. 994–1006, 2017.
- [111] T. H. Davenport and S. Dastgupta, "How to set up an AI center of excellence," *Harvard Bus. Rev. Digit. Articles*, pp. 1–11, 2019.
- [112] W. Pang. (Jul. 2020). *Why an AI Center of Excellence is the Key to Success*. [Online]. Available: <https://finovate.com/why-an-ai-center-of-excellence-is-the-key-to-success/>



**DORIEN HERREMANS** (Senior Member, IEEE) received the Ph.D. degree in applied economics from the University of Antwerp. She was awarded a Marie-Curie Fellowship to work at the Centre for Digital Music, Queen Mary University of London. Before that, she graduated as a Business Engineer in Management Information Systems at the University of Antwerp in 2005, after which she worked as a Drupal Consultant and as an IT Lecturer at the world's leading hospitality business

school, Les Roches, Bluche, Switzerland. She is currently an Assistant Professor with the Singapore University of Technology and Design (SUTD). At SUTD, she is also the Director of the SUTD Game Lab and heads both the Audio, Music, and AI (AMAAI) Lab as well as the AI for Finance (AIFI) Group. Her passions include strategic thinking as well as novel applications of AI technologies. She serves on numerous committees and boards and has given invited talks around the world. She is on the Singapore 100 Women in Tech 2021 list, which recognizes and celebrates women based at Singapore who have been inspiring and have made significant contributions to the tech industry.

• • •