



Applied Artificial Intelligence

An International Journal

ISSN: (Print) (Online) Journal homepage: <https://www.tandfonline.com/loi/uaai20>

Enterprise AI Canvas Integrating Artificial Intelligence into Business

Ulrich Kerzel

To cite this article: Ulrich Kerzel (2021) Enterprise AI Canvas Integrating Artificial Intelligence into Business, Applied Artificial Intelligence, 35:1, 1-12, DOI: [10.1080/08839514.2020.1826146](https://doi.org/10.1080/08839514.2020.1826146)

To link to this article: <https://doi.org/10.1080/08839514.2020.1826146>



© 2020 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 04 Oct 2020.



Submit your article to this journal [↗](#)



Article views: 4517



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 9 View citing articles [↗](#)



Enterprise AI Canvas Integrating Artificial Intelligence into Business

Ulrich Kerzel

Fernstudium, IUBH Internationale Hochschule GmbH, Erfurt, Germany

ABSTRACT

Artificial Intelligence (AI) and Machine Learning have enormous potential to transform businesses and disrupt entire industry sectors. However, companies wishing to integrate algorithmic decisions into their organization face multiple challenges: They have to identify use-cases in which artificial intelligence can create value, as well as decisions that can be supported or executed automatically. Furthermore, the organization will need to be transformed to be able to integrate AI-based systems into their human work-force. In addition, the more technical aspects of the underlying machine learning model have to be discussed in terms of how they impact the various units of a business: Where do the relevant data come from, which constraints have to be considered, how is the quality of the data and the prediction evaluated? The Enterprise AI canvas is designed to bring data scientist and business expert together to discuss and define all relevant aspects which need to be clarified in order to integrate AI-based systems into a digital enterprise. It consists of two parts, where part one focuses on the business view and organizational aspects, whereas part two focuses on the underlying machine learning model and the data it uses.

1. Introduction

Artificial intelligence (AI) and Machine Learning have enormous potential to transform businesses and disrupt entire industry sectors. When referring to artificial intelligence, the distinction between general (or strong) and narrow AI is generally made: General AI systems (Turing 1947) (McCarthy 2007) exhibit intelligent behavior en par or even surpassing human intelligence in a wide range of scenarios. While general AI systems are still elusive, it is expected that they would interact and reason much like we humans do. Narrow AI systems on the other hand aim to solve a particular problem and specialize in the execution of a specific, singular task. This approach has seen spectacular successes such as playing the game “Go” without human knowledge (Silver et al. 2017) or cancer detection (Esteva et al., 2017; Ardila et al. 2019).

CONTACT Ulrich Kerzel ✉ u.kerzel@iubh-fernstudium.de IUBH Internationale Hochschule GmbH, Erfurt D-99084, Germany

© 2020 The Author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

While the performance of narrow artificial intelligence is indeed astonishing in these and other examples, they mainly focus on a specific research question. Businesses on the other hand first and foremost need to address the question how to generate value for their customers and stakeholders. P. Sondergaard, Senior VP at Gartner said already in 2015 that “algorithms are where the real value lies [because] algorithms define action” (Levy 2020). This statement captures the essence of algorithmic decisions quite well: Algorithms allow to create value and can be tied to decisions: The output of an algorithm is typically a prediction, either a classification or a regression, which can then be translated into operational decisions. While many companies also engage in fundamental research, the driving factor for most innovations and offerings can be traced back to the question: How does the company create value?

Integrating artificial intelligence into a business requires more consideration than traditional analytics: While integrating data-driven decisions is not new for many companies, these typically focus on establishing or expanding teams of specialists with sophisticated tools. As these activities are still based on human teams, they can be integrated into organizations in familiar ways including management and organizational structures. However, integrating decisions taken by artificial intelligence also needs to address the questions how human and artificial intelligence work together and how and by whom decisions are taken. From an organizational perspective, AI systems are not “managed.” Furthermore, for the foreseeable future AI systems are not suitable for all kind of decisions: Some decisions are better taken by humans, others are more suitable to be delegated to an artificial intelligence. Deciding which category a particular business idea falls into requires extensive technical expertise beyond the level of most business experts.

The Enterprise AI canvas is designed to bring data scientists, AI specialists and business experts together and aims to ask all relevant questions that need to be addressed to first identify the most promising use-cases which bring value to the company and its customers, as well as more technical details regarding the underlying data and machine learning model.

2. Relation to other Work

Many approaches exist to identify new business opportunities. A very successful approach is the Business Model Canvas (BMC), proposed by Osterwalder and Pigneur (Osterwalder and Pigneur 2010; Osterwalder 2004). The Business Model Canvas used nine building blocks that focus on the value proposition, key partners, activities, resources, as well as customers, and cost and revenue streams. The BMC is a helpful tool for business experts to identify new value propositions and how they bring value to the company and customers, however, as a general canvas, the BMC does not address the additional

requirements of evaluating whether the proposition is suitable for an AI-based system or integrating artificial intelligence into the organization. Implicitly, the BMC assumes that some team will be responsible of realizing the value proposition – but the details do not need to be evaluated at this stage. The Machine Learning Canvas (MLC) proposed by Dorard (Dorard n.d.) on the other hand focuses mainly on the technical aspects of developing a use-case suitable for machine learning algorithms, covering model related aspects such as data sources, data collection, features, as well as making predictions and the resulting decisions or actions. While the MLC also contains the “value proposition” as central element, most of the blocks of the canvas are related to technical details which are more relevant to data scientists than to business experts, which makes the machine learning canvas mainly relevant for the data science team developing and implementing a business idea. The AI Canvas proposed by Agrawal *et al.* (Agrawal, Gans, and Goldfarb 2018) uses AI as a general term for machine based predictions, synonymly with machine learning. This canvas captures the main ideas of a proposed new project such as what should be predicted and the resulting action, as well as some general aspects of required data and training but doesn't go into neither the technical details like the machine learning canvas nor how businesses would integrate this in their operations. The AI Canvas proposed by K. Dewalt (Dewalt n.d.) provides a high-level mix of business and data science centric aspects. The business-centric aspects include discussions about *opportunity*, *solution*, and *customers*, similar to *value proposition* in the machine learning canvas by L. Dorard. Similar to the Machine Learning Canvas, this canvas also includes fields for *data sources*, *model development*, and *success criteria*, though they remain on a higher level compared to the machine learning canvas. Finally, the AI Project Canvas suggested by J. Zawadzki (Zawadzki 2019) is based on the original business model canvas and focuses mainly on the development of the business idea such as the *value proposition* cost and revenue, as well as who are the end customers and stakeholders and required skills. However, the AI Project Canvas doesn't go in more technical details and is hence mainly aimed at project managers.

3. The Enterprise AI Canvas

At first glance, it might be tempting to hand the Business Model Canvas (Osterwalder and Pigneur 2010; Osterwalder 2004) or the AI Project Canvas (Zawadzki 2019) to business experts or project managers and then the Machine Learning Canvas (Dorard n.d.) to data scientists in the next step. This is indeed what happens in industry and application quite often, although the various canvases are seldom used in practice and instead potential projects are discussed and derived in presentations or flip-charts. However, in many practical situations this will not lead to satisfactory results: Developing data-

driven applications using machine learning or (narrow) artificial intelligence requires a detailed understanding of available or obtainable data for a given use-case as well as an estimation of the feasibility developing a model in reasonable time. While these aspects can be well understood by (senior) data scientists, business experts, or project managers are often too far detached from these details. On the other hand, data scientists have typically only limited insight into how to create business value and optimize operational decisions in a corporate setting, as well as which organizational change or change management would be required to integrate an AI project in a company. Project ideas and details, technical aspects, and organizational implications are intertwined and closely related. Discussing the business aspects first and getting them approved by the senior management before engaging deeply with the data scientist can either lead to unnecessary loops re-defining projects at best and implementing something which fulfills the project plan but doesn't meet the requirements of the customer at worst. In addition, business experts and data scientists have a very different background and use a different language or jargon.

The Enterprise AI canvas proposed in this work is specifically designed such that all relevant communities such as business and domain experts as well as data scientists can work together on a common framework. This allows to capture the expertise from a diverse project planning team and ensures that all relevant aspects of the whole project are considered. Special focus has been given to the consideration of decision making, decision optimization, and impact on organizational structures. The technical details such as defining the prediction target are often discussed among data scientists only, these specifications have often far-reaching consequences on the business side: As operational decisions are taken by machines, how should the company deal with the people who made these decisions earlier? What will these people work on and how will their job role change? A further important aspect is to consider how operational decisions are optimized. Prior to the AI-based project, deriving the optimal decision has typically been the responsibility of a manager or expert, as decisions are taken automatically, how is the optimization process changed to incorporate the predictions from the AI?

The Enterprise AI Canvas consists of two parts, the first part is mainly aimed at the business perspective of integrating an AI system, the second part focuses more on the underlying machine learning and data model. Both parts are shown in detail in [Figure 1](#) and all elements are discussed below in detail. Each element of the canvas can be considered separately in any order, however, it is advisable to start with the central elements first and then move toward the most closely related elements. Hence it is helpful to start with *Value* in part one and first consider how the potential new business opportunity can create value. Consequently, it is helpful to start with *Prediction & Action* in

Enterprise AI Canvas – Part 1: Business View		
<p>Decision & Optimization: How are decisions optimized? By whom?</p>	<p>Value: How does the use-case generate value? What can be offered? Which problem can be solved?</p>	<p>Success: How is success (metrics, KPIs) defined? What makes the use-case “good” or “bad”? Who decides this?</p>
<p>Organization: How are decisions & actions currently formed? How will this change with this use-case? Change Management? Additional training?</p>		<p>Domains: Which domain expertise is needed? Which business units are involved?</p>
<p>Sponsor: Which senior manager is responsible?</p>		
Enterprise AI Canvas – Part 2: Model & Data View		
<p>Domains and Data Quality: Which domain expertise is needed? How and who can decide whether the data are “good” or “bad”? How to measure & improve data quality?</p>	<p>Prediction & Action: What should be predicted? How are predictions (as probabilities) transformed into actions?</p>	<p>Features: Which features are likely important?</p>
<p>Data Sources & Processing: Which data sources are required, which are already available? Which types of data are needed? Which systems are required and already available to handle the data?</p>		<p>Constraints: What is the typical time-frame for a prediction? Which granularity is required? How often does the model need to be re-trained? How are the models served? Edge, on-premise or Cloud?</p>
<p>Evaluation & Monitoring: Which metrics are used? How will the project be monitored? What happens in case of incidents?</p>		

Figure 1. Enterprise AI Canvas.

part two to discuss what needs to be predicted exactly and how a prediction by a machine learning model can be turned into an operational action.

3.1. Part 1

The first part of the Enterprise AI canvas focuses on the business view of the opportunity: How does the opportunity bring value, what is meant by “success,, how does it impact the organization to deliver the opportunity?

- **Value:** The central item of the potential business opportunity is how it creates value for the organization and the customers. The main relevant

questions are focused on the potential offering and which problem can be solved with the new business opportunity. The potential use-case may be a new offering or a substantial improvement on previous approaches.

- **Success:** The central part of a new business opportunity or a significant improvement of an existing offering is connected to its success. However, what *success* means needs to be carefully defined. From a technical perspective, data scientists may be tempted to focus on metrics that focus on the quality of the individual predictions made by the machine learning model. However, as these predictions are embedded into operational decisions, a “good” prediction does not necessarily translate into a “good” operational decision. Business and data science experts need to work together and identify specific metrics that allow to evaluate the quality of the operational decisions made by the AI model. These are likely not technical metrics but Key Performance Indicators (KPIs) related to the business opportunity. Note that this does not imply that the technical metrics should not be used, however, there are often many layers between the prediction of a machine learning model and the resulting operational decision that need to be taken into account. What matters in the end is that the best possible decision is taken.

- **Decisions & Optimization:** This aspect focuses on the operational aspect of the decisions related to the potential use-case. If an existing approach is improved substantially using an AI-based system, business specialists need to understand which decisions are taken at the moment and how they are formed. Which reasoning goes into making decisions? Ideally, decisions are already based on a clear approach or data-driven analytics, although many decisions may still be taken with “gut feeling.. If a new business opportunity is explored, the associated decisions need to be made explicit and the objectives according to which the decisions are optimized need to be specified.

- **Organization:** This item is related to the point *Decision & Optimization* which focuses on which decisions need to be made in the context of a new business opportunity and how these decisions are optimized to achieve the best possible outcome. The item *Organization* focuses on who currently makes the decision: How are they formed? Who does that? What will happen once the AI-based system is integrated in the process? Which decisions will be taken by humans, which by the AI? What happens with the humans currently involved? Will they need additional training? Does part of the organization or even the whole organization need to be changed to accommodate the fact that operational decisions are now at least partly taken by a machine? Who will prepare the organization for this change and how? Failing to address these points can create significant anxieties among the employees, such as the fear of loss of their jobs or similar.

- **Sponsor:** Any new substantial change in the operations and organization of company requires careful consideration by the senior management team. This is particularly true in the case of systems based on artificial

intelligence which may likely have a large impact on the way the business works when human and artificial intelligence work together. It is therefore paramount to identify the senior management team that is responsible for this project to guide the company through the necessary change process.

- **Domains:** In order to generate value, a new business opportunity has to be embedded in a specific industry domain or application area. When evaluating the prospects of a new opportunity it is important to integrate the relevant domain expertise from an early point on. This can relate both to internal expertise in various business units, as well as external expertise.

3.2. Part 2

Part 2 of the Enterprise AI canvas focuses more on the technical aspects of the implementation of the machine learning model as the foundation of the artificial intelligence based decision system.

- **Prediction & Action:** This item relates to *Decision & Optimization* and *Value* in part one. Whereas the aspects in part one were mainly concerned with the business opportunity and how decisions are derived and implemented, the part *Prediction & Action* in part two focuses on the concrete prediction the AI system needs to make to fulfill the objective, i.e., which quantity exactly needs to be predicted. Is this a classification problem or should a quantity be predicted (regression)? Fundamentally, all predictions are probabilities (in case of classifications) or probability density distributions (in case of regression). How are these probabilities transformed into actions? If, e.g., a threshold needs to be defined, how is this value determined?

- **Features:** Once it has been defined what exactly the model needs to predict, extensive consideration should be given to potential feature variables that are likely needed by the machine learning model to calculate the individual predictions. Although not every aspect can be considered at this stage as more feature variables might be deemed important during the development phase, this step helps to understand which data sources will be required.

- **Data Sources & Processing:** This part relates to *Features*: Once the major feature variables have been identified, the next step is to investigate which data sources can be accessed to extract the relevant data. Potential data sources may cover a wide range of internal and external sources. Ideally, the list should be ranked according to a first estimate of the priority of access of these system as extracting data from each system takes time and effort and the most important data sources should be handled first. Furthermore, the AI system itself also requires adequate processing capabilities both for handling the data and for the calculation of the predictions. The required computing and storage capacity need to be considered as part of the project planning, as well as where these systems are operated, such as a private data-center or a cloud-based system.

- **Domains & Data Quality:** Ultimately, an AI-based system is only as good as the data it can use to learn from. Consequently, the quality of the data is paramount and significant effort should be spent to clean the data before they are used to train an AI system. Although often mentioned in popular belief, data don't speak for themselves but deep domain expertise is required to decide whether a particular data point is correct or whether an outlier is likely due to a technical issue or is related to a rare but valid event. Which expertise is needed to assess data quality? Which criteria can be defined to measure and quantify the quality of the data and which processes can be established to improve the quality of the data throughout the project?

- **Constraints:** In most practical applications, a number of constraints have to be met. For example, the predictions have to be calculated within a specific time or particular requirements have to be met regarding data security and data privacy. In some cases, this may imply that cloud-based services are the best approach, in other cloud-based services may be ruled out and dedicated data centers or edge devices have to be used.

- **Evaluation & Monitoring:** This part is related to *Success* from part one of the canvas. After defining how the project is evaluated it now needs to be discussed how the relevant metrics will be calculated and presented. Ideally, a first plan is made how to handle any incidents or deviations from the allowed range of metrics.

4. Example

In the following example, a fictitious supermarket chain was analyzed and the Enterprise AI canvas filled in accordingly (Figure 2). The example was chosen because the different elements are easily relatable to everyday experiences and because of the author's experience in this specific area. In general, supermarkets are offering a wide range of products to end customers where the range of products (assortments) varies from hundreds or thousands of different products for discounters to tens of thousand different products for general supermarkets. Typically, supermarkets get new goods once per day or every few days. Many products have a long shelf-life, such as frozen, canned, or dry food, and hence replenishing these products is generally not critical as irregularities in the shipment can be compensated in the next cycle as long as sufficient stock is available on the shelf. However, a wide range of products has only a short shelf-life and needs to be disposed of once the products perish. These include, for example, fresh produce, flowers, bakery products, dairy, meat, and similar. Some of these products (e.g., fresh bread) can only be sold during a single day. These products, called fresh or ultra-fresh, are much more challenging from a replenishment perspective: They can't be stocked to compensate for an unexpected surge in demand or a missed shipment and keeping a large safety stock leads to high level of waste disposal which is both costly and

Enterprise AI Canvas – Part 1: Business View		
<p>Decision & Optimization: How are decisions optimized? By whom?</p> <p>Final decision: Order quantity of a given article (per store, per day), including lot sizes and delivery schedule.</p> <p>future approach: Automatic determination of optimal order quantity and transmission to supplier including all constraints, human experts supervise and modify in exceptional circumstances. Each modification is tracked.</p>	<p>Value: How does the use-case generate value? What can be offered? Which problem can be solved?</p> <p>customer:</p> <ul style="list-style-type: none"> - fewer stockout situations - fewer frustrated customers due to stocking issues <p>retailer:</p> <ul style="list-style-type: none"> - reduced waste of perishable goods - reduced overall inventory - fewer stockout situations - fewer frustrated customers due to stocking issues 	<p>Success: How is success (metrics, KPIs) defined? What makes the use-case "good" or "bad"? Who decides this?</p> <p>Relevant business metrics: stockout rate (< xx%), waste of perishable goods (< xx%), overall inventory (needs to fit into currently available space).</p> <p>Decided by: Data Scientist: possible values considering constraints from lot-size, delivery schedule, Senior Management: KPI values according to strategy within possible range.</p>
<p>Organization: How are decisions & actions currently formed? How will this change with this use-case? Change Management? Additional training?</p>		<p>Domains: Which domain expertise is needed? Which business units are involved?</p> <ul style="list-style-type: none"> - Product Management: Product characteristics, history - Marketing: Past and future marketing campaigns, marketing characteristics - Supply Chain: Constraints and schedules, vendors and alternatives
<p>Sponsor: Which senior manager is responsible? CEO: Core business processes are affected, needs to be approved by CEO. Senior Management (e.g. Chief Procurement Officer) responsible for the project and buy-in from relevant other departments.</p>		
Enterprise AI Canvas – Part 2: Model & Data View		
<p>Domains and Data Quality: Which domain expertise is needed? How and who can decide whether the data are "good" or "bad"? How to measure & improve data quality?</p> <p>Required expertise:</p> <ul style="list-style-type: none"> - IT from operational systems (PoS, Marketing, Supply Chain) to extract data and verify data quality issues. - Product managers to provide details of past, current and future products, product changes, discontinued products, etc and verify data quality issues 	<p>Prediction & Action: What should be predicted? How are predictions (as probabilities) transformed into actions?</p> <ul style="list-style-type: none"> - predict: future demand per day per store for various forecast horizons: 1 day into the future: operational planning. 2-30 days into the future: create awareness of longer term developments, backup in case 1 day horizon becomes unavailable. - Predict the full PDF (regression) - Simulation: how do the quantiles of the predicted PDF translate into KPIs such as stockout, waste? - Senior Management to choose best operational point on customer-centric metrics 	<p>Features: Which features are likely important?</p> <ul style="list-style-type: none"> - historic sales data (per day, per store), maybe using time-series as additional input. - product properties - store characteristics - weather - calendar information
<p>Data Sources & Processing: Which data sources are required, which are already available? Which types of data are needed? Which systems are required and already available to handle the data?</p> <ul style="list-style-type: none"> - Operational systems: PoS data, ERP - external: Weather, calendar information (e.g. school/bank holidays, special events such as Christmas, Easter but also e.g. major sports competitions) 		<p>Constraints: What is the typical time-frame for a prediction? Which granularity is required? How often does the model need to be re-trained? How are the models served? Edge, on-premise or Cloud?</p> <ul style="list-style-type: none"> - predictions required daily. PoS data available at midnight, orders need to be placed by 4 am at supplier to be fulfilled in the next cycle. - data and calculations can be done in the cloud – no on-premise solution required.
<p>Evaluation & Monitoring: Which metrics are used? How will the project be monitored? What happens in case of incidents? Live and long-term monitoring of customer centric metrics (stockout, waste). MAD as indicator for prediction quality of point estimators. Profile histogram of point estimator (prediction) vs observed sales to monitor for any bias.</p>		

Figure 2. Example: Enterprise AI Canvas for a supermarket chain focusing on store replenishment.

harmful for the environment. Keeping the stock too low on the other hand will result in stock-out situations where customers face empty shelves. Therefore, supermarkets increasingly turn to automatic ordering systems based on artificial intelligence, that predict the expected demand of individual products and optimize the shipments to each store (Robinson 2020). Such a system has the benefit that it can take all available information into account from past sales records to weather information, as well as constraints from lot sizes and shipping schedules.

- **Value:** Optimized replenishment decisions are beneficial for both the customer and the retailer: The customer face fewer stock-out situations in

which they do not find the desired goods. In addition, the retailer can reduce the amount of waste due to perished products and reduce the overall inventory.

- **Success:** The success of the replenishment decisions can be measured by operational key performance indicators (KPIs), in particular the waste rate, i.e., the amount of perished products that have to be disposed of, and the stock-out rate, i.e., the number of times products are not available on the shelves for customers to buy. In principle, the overall inventory is a third metric to optimize, however, in times of low-interest rates the latter optimization is often less of a concern as long as the excess inventory does not perish quickly and fits into existing storage capacities. These metrics need to be decided upon both by the business and data science experts: Business experts need to determine which KPIs are relevant for the project and together with the data scientists they need to identify which range of values are possible due to the constraints from, e.g., lot-size or shipping schedule and which range of KPIs is considered to be successful for the project.

- **Decisions & Optimization:** In this example, the operational decisions taken are the order quantity per product per delivery cycle and store. Following the implementation of the project, all decisions should be taken by the AI-based system with human oversight to react to unforeseen circumstances such as extreme weather, food scandals, etc.

- **Organization:** This part of the canvas has been intentionally left blank as this varies significantly between different retailers. However, typically retailers move from a manual system supported by spreadsheets, where each decision was taken by a human operator, to an automatic system. This generally requires both change management efforts to prepare the company as a whole to this shift, as well as training of the individual operators whose role is changing due to the implementation of the project.

- **Sponsor:** As the change of the replenishment process changes, core business processes are affected which should be generally sponsored and monitored by the CEO. However, a member of the senior management team should be assigned in addition to focus on the operational responsibilities of the project.

- **Domains:** Successful implementation of such a project will require the cooperation and interaction with several departments such as product management, marketing, and supply chain experts. Since several critical departments need to work together, the sponsor needs to be sufficiently senior to bring all required departments together.

- **Prediction & Action:** The starting point of the replenishment is the prediction of the future demand per product per order cycle per store. This quantity is ideally predicted as a full probability density distribution from which the optimal point estimator is derived which forms the basis for the later ordering decision. Typically, these predictions are needed 1 day ahead for operational purposes and 2-30 days ahead or longer to anticipate future trends and as a fail-safe precaution: Should the prediction not be available

for any reason, these longer term prediction can be used as a basis for the subsequent ordering decision instead. In order to derive the optimal point estimator, a simulation can be used in which each quantile of the predicted probability density distribution forms the basis of a simulated order and the effect on the metrics chosen in the element *Success* can be studied.

- **Features:** Previous sales records will be the most important information for future demand, assuming the underlying consumer behavior hasn't changed drastically. In addition, properties of the individual products influence buying decisions as well as the location, size, and other store characteristics due to their effect on customer segmentation. Current information like weather forecasts or special calendar events such as holiday, festive seasons, or major sports events will influence demand as well.

- **Data Sources & Processing:** Historic sales data will be in some form of a data warehouse or database, new records are either added after the store closes and the data are pulled from the till or the tills stream each completed transaction into the relevant network. Marketing information, product data, etc. will be in some Enterprise Resource Planning (ERP) tool. In addition, external information has to be brokered from information providers such as weather forecast, calendar information, etc. All data have to be made available for the training phase.

- **Evaluation & Monitoring:** Customer and business centric metrics as defined under *Success* will be monitored at least on a daily basis, more frequently if the data are available. In addition, further metrics such as the mean absolute deviation (MAD) can be used to evaluate the quality of the prediction. A profile histogram showing the behavior of the predicted point estimator against the observed sales can help to determine if there are any biases.

5. Conclusion

Recent advances in artificial intelligence are set to disrupt a wide range of industries and sectors. However, companies wishing to integrate AI-based systems into their offering or operations face significant challenges how to proceed. The Enterprise AI canvas allows to bring business and data science experts together and systematically evaluate potentially new business opportunities. The canvas consists of two parts, where part one focuses on the value a new AI-based offering or operation would bring to the business, whereas part two focuses more on the technical aspects of the underlying machine learning model. Both parts are intertwined and relate to each other to allow an optimal evaluation of business and data science-related aspects of a new project.

References

- Agrawal, A., J. Gans, and A. Goldfarb. 2018. A simple tool to start making decisions with the help of ai. Accessed 11 September 2020. <https://hbr.org/2018/04/a-simple-tool-to-start-making-decisions-with-the-help-of-ai>
- Ardila, D., A. P. Kiraly, S. Bharadwaj, B. Choi, J. J. Reicher, L. Peng, D. Tse, M. Etemadi, W. Ye, G. Corrado, et al. 2019. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. *Nature Medicine* 25: 954–961.
- Dewalt, K. n.d. Become an ai company in 90 days. ISBN 978-0-692-19233-7.
- Esteva, A., B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun. 2017, Jan. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542 (7639):115 EP. doi: [10.1038/nature21056](https://doi.org/10.1038/nature21056).
- Levy, H. The arrival of algorithmic business. Accessed 11 September 2020. <https://www.gartner.com/smarterwithgartner/the-arrival-of-algorithmic-business/>.
- McCarthy, J. 2007. From here to human-level ai. *Artificial Intelligence* 171 (18):1174–82. Special Review Issue. doi:[10.1016/j.artint.2007.10.009](https://doi.org/10.1016/j.artint.2007.10.009).
- Osterwalder, A., and Y. Pigneur. 2010. *Business model generation*. Hoboken, New Jersey, USA: J. Wiley & Sons.
- Osterwalder, A. 2004. *The Business Model Ontology Proposition in a Design Science Approach*. PhD thesis, University of Lausanne, CH.
- Robinson, E. Accessed 11 September 2020. <https://www.technologyrecord.com/Article/morrison-implements-blue-yonders-ai-stock-replenishment-technology-62728>
- Silver, D., J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. 2017, Oct. Mastering the game of go without human knowledge. *Nature* 550 (7676):354 EP. (Article). doi: [10.1038/nature24270](https://doi.org/10.1038/nature24270).
- Turing, A. M. 1947. Lecture to the london mathematical society. *The Collected Works of A.M. Turing* Mechanical Intelligence, 390.
- Zawadzki, J. 2019. Introducing the ai project canvas. Accessed 11 September 2020. <https://towardsdatascience.com/introducing-the-ai-project-canvas-e88e29eb7024>