
Advancements in ML-Enabled Intelligent Document Processing and How to Overcome Adoption Challenges in Enterprises



Tomasz Janasz, Peter Mortensen*, Christian Reisswig*, Tobias Weller*, Maximilian Herrmann, Ivona Crnoja and Johannes Höhne*

Abstract: The ability to automatically extract and process information from business documents is crucial to many business processes. With the advent of powerful machine learning systems, specifically triggered by deep learning, big data, and today's computing resources, it has become possible to automate document processing tasks. Although AI-enabled business document processing can be an important driver in the digital transformation of businesses, the adoption by enterprises is considered rather low. This paper presents new advancements in business document processing based on machine learning and introduces business scenarios which benefit from its application. It also offers a holistic view on challenges faced by suppliers and buyers of ML applications. We derive a range of critical success factors and discuss the interrelationship between them in the context of ML.



Keywords: Artificial intelligence, machine learning, intelligent document processing, business value, ML adoption, critical success factors



Fortschritte bei der ML-basierten intelligenten Dokumentenverarbeitung und wie die Herausforderungen bei der Einführung und Nutzung in Unternehmen bewältigt werden können

Zusammenfassung: Die Möglichkeit, Informationen aus Geschäftsdokumenten automatisch zu extrahieren und zu verarbeiten, ist für viele Geschäftsprozesse von entscheidender Bedeutung. Mit leistungsstarken Machine-Learning-Systemen, die speziell durch Deep Learning, Big Data und heutige Rechenressourcen ermöglicht werden, können Dokumentenverarbeitungsaufgaben mithilfe auf künstlicher Intelligenz basierender Technologien automatisiert werden. Auch wenn die KI-basierte Verarbeitung von Geschäftsdokumenten ein wichtiger Faktor für die digitale Transformation von Unternehmen sein kann, wird deren Einführung und Nutzung als eher gering angesehen. Diese Publikation stellt neue Ansätze für die Verarbei-

* Main contributors of the publication.



tung von Geschäftsdokumenten auf der Grundlage von KI sowie Geschäftsszenarien vor, die von deren Anwendung profitieren könnten. Außerdem bietet sie eine ganzheitliche Sicht auf die Herausforderungen, mit denen sich sowohl Lieferanten als auch Käufer von ML-Anwendungen konfrontiert sehen. Wir leiten eine Reihe kritischer Erfolgsfaktoren ab und diskutieren die Wechselbeziehungen zwischen ihnen im Kontext von ML.

Stichwörter: Künstliche Intelligenz, maschinelles Lernen, intelligente Dokumentenverarbeitung, Geschäftswert, ML-Einführung/-Nutzung, kritische Erfolgsfaktoren



1 Introduction

Artificial Intelligence (AI) is considered one of the most exciting domains in the technology landscape. Analysts are convinced that AI will play an integral part in any business sector and industry (Manyika et al. 2017). Machine learning (ML), the major subset of AI¹, has vast potential for value creation across the entire economy (Brynjolfsson/Mitchell 2017; Brynjolfsson/McAfee 2017). It has effects on the consumer level, but it also changes how entire enterprises operate. In the recent *Future Predictions Report for AI* (Jyoti et al. 2019), analysts firm IDC observe a growing demand for enterprise-ready ML along with an expansion of ML-based automation across enterprises. The report predicts that as much as 75 % of enterprises will embed some form of intelligent automation into their technology and process development by 2022. As outlined in this work, intelligent document processing (IDP) is a particularly promising application area for ML. This paper aims at identifying the potential of IDP and it also covers critical success factors regarding the adoption of ML technology by enterprises.



The first part of this paper focuses on the latest advancements in IDP and their benefits for enterprises. The second part presents the results of a qualitative research study on the challenges of ML adoption based on expert interviews. The topic is approached holistically by deriving critical success factors for suppliers and buyers while simultaneously investigating the interrelationship between them.



With our contribution, we seek to extend the existing body of management knowledge by identifying key patterns from business practice in the field of applied AI. The paper is targeted towards forward-looking practitioners from all sectors and industries intending to facilitate ML adoption in their enterprise context as well as researchers working on AI in organizations.

¹ ML is the most important subfield of AI. When talking about AI, most people refer to ML. Hence, we use the more technical term ML.

2 Literature Review

Business value of intelligent document processing

The possibilities of electronically exchanging structured information are continuously growing. However, many business processes still revolve around the sending, receiving, and processing of unstructured documents. Invoices alone account for an estimated yearly volume of 550B documents globally, which is expected to quadruple by 2035, with other invoice-like documents adding an additional volume of 5–15 times the invoice volume (Koch 2019). Around 80 % of data in enterprises is considered to be unstructured and most of it is locked in documents (Prasad 2020). Staffing and labor cost reductions of over 50 % are reported by companies that adopted an automated document processing solution (Hyland Software Inc. 2019; Clark 2020), while reducing processing time to less than 20 % of the manual processing time (Canon Business Process Services 2019). Thus, automation of business document processing is a key driver for businesses in the reduction of costs and the increase of efficiency.

Using ML technology for automatic information extraction

When extracting information from a document using machine learning, its layout structure must be considered. In recent times, this issue has been addressed by several ML researchers (Liu et al. 2019; Xu et al. 2019; Garncarek et al. 2020; Wei et al. 2020; Yu et al. 2020). For example, additional spatial layout information was incorporated to facilitate the task either by working on the document text plus spatial position information or by directly using the document image as input (Zhang et al. 2020). In this context also a two-dimensional document representation, which retains the document layout, has been introduced recently (Katti et al. 2018; Denk/Reisswig 2019).

Challenges in the adoption of ML

Even though previous research suggests that a competitive advantage can be achieved through advanced analytics and ML (Chen et al. 2012; Sanders 2016; Côte-Real et al. 2017; Reis et al. 2020), enterprises still face serious challenges in the adoption of ML. The lack of a clear strategy for AI, the lack of talents and the lack of data are the most significant barriers. Soft aspects such as leadership, decision making, and company culture are also considerable challenges (McAfee/Brynjolfsson 2012; Chui/Malhotra 2018). Literature extensively covers the outside-in understanding of the adoption barriers on the buyer side (e.g. Chui/Malhotra 2018; Lorica/Nathan 2019). Yet, the supplier perspective is only partially covered, with research focusing mainly on technology or developer efficiency (e.g. Amershi et al. 2019). The consideration of the supplier perspective regarding the adoption challenges can be of significant importance as value is created through continuous product and process innovations (Mansfield et al. 1977). Hence, it is important to establish a holistic understanding of the interrelationship between both perspectives. With this paper we aim at closing this gap.

3 Intelligent Document Processing

3.1 Document Information Extraction

Understanding and extracting information from unstructured documents is part of an emergent area in Natural Language Processing (NLP) known as IDP or simply ‘Document Intelligence’. It typically consists of two steps: (i) converting a document image to text (if the document is not already available as text), and (ii) interpreting the text (see Figure 1). Identifying the information and interpreting it within the context of the document text, including country-specific parsing of units, dates, currencies, and amounts, is challenging.

Katti et al. (2018) and Denk/Reisswig (2019), introduced a new type of two-dimensional document representation – Chargrid (see Figure 2) – which works with the textual information of a document while retaining the document layout. Hence, a downstream machine learning system has access to both textual and layout information. The layout may carry semantically meaningful information. Leveraging this information results in higher system accuracy which is advantageous compared to previous NLP approaches such as named entity recognition (Lample et al. 2016; Palm et al. 2017).

For documents where the layout is not relevant, one can make use of ‘traditional’ sequence tagging approaches powered by a variety of deep neural networks such as recurrent networks and transformers. Transfer-learning and pretrained language models boost the accuracy even in the few-shot regime when only few data samples are available for training.

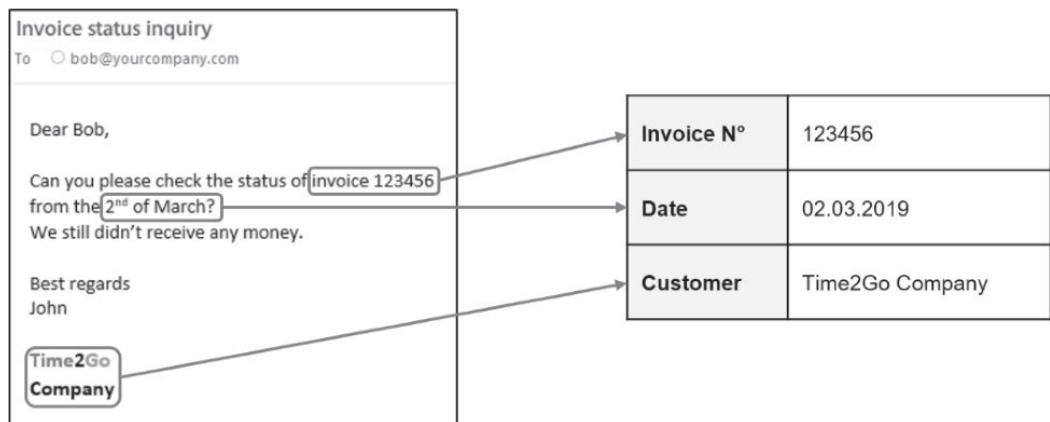


Figure 1: Exemplary information extraction from an email.

When converting document images into a text, the Optical Character Recognition (OCR) ML approach based on the Chargrid representation can be applied (Reisswig et al. 2019). This outperforms traditional and more complicated approaches in terms of accuracy and robustness, especially for poorly scanned or photographed documents.




Figure 2: Image of a document page (left) and its character grid (Chargrid) representation (right). Each pixel overlapping with the bounding box location of a character on the document page is mapped to the vocabulary index of that character. The Chargrid representation incorporates textual as well as layout information such as text position and font size.

3.2 Document Classification

Classifying incoming documents is a common entry point to document related business processes. It is used to determine how to proceed with these documents, whom to direct them to, and how to process them. When classifying – be it manual or automated – it must be possible to distinguish between different document types.

Besides identifying the type of the document, classification can also be used to identify other categories such as the responsible processor, criticality, or language of a given document (see Figure 3).

Several machine learning techniques can be used to perform classification of documents. To extract the text from image-based PDF documents, OCR can be applied. Subsequently, the text can be classified by using text mining techniques such as term frequency – inverse dense frequency (TF-IDF), N-grams, and supervised learning techniques such as random forest.

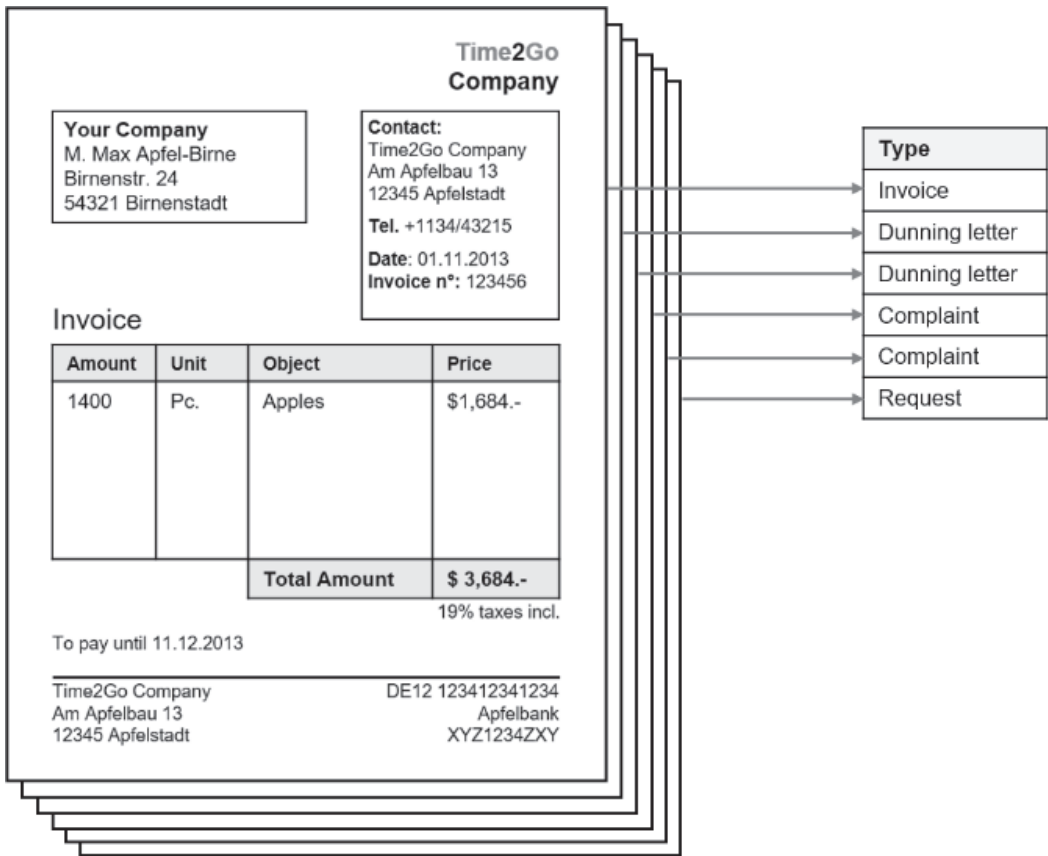


Figure 3: Exemplary use case of document classification. The classifications (here: the document type) are custom and could also represent other criteria such as criticality of documents.

3.3 Document Enrichment / Matching

Extraction focuses on converting documents into a structured machine-readable format. However, additional information might be required before posting the document in an Enterprise Resource Planning (ERP) system. In this enrichment step, supporting data from an ERP system is used to add context to the documents and to validate it (see Figure 4).

Typical enrichment scenarios are matching address information to business partners such as customers or suppliers, matching invoice documents to their corresponding purchase orders, purchase orders and their line items to the material catalogue, or payment advice items to open receivables. This leads to a significant improvement in data quality and confidence while simultaneously identifying errors or incorrect documents before they are posted (Level Research 2019).

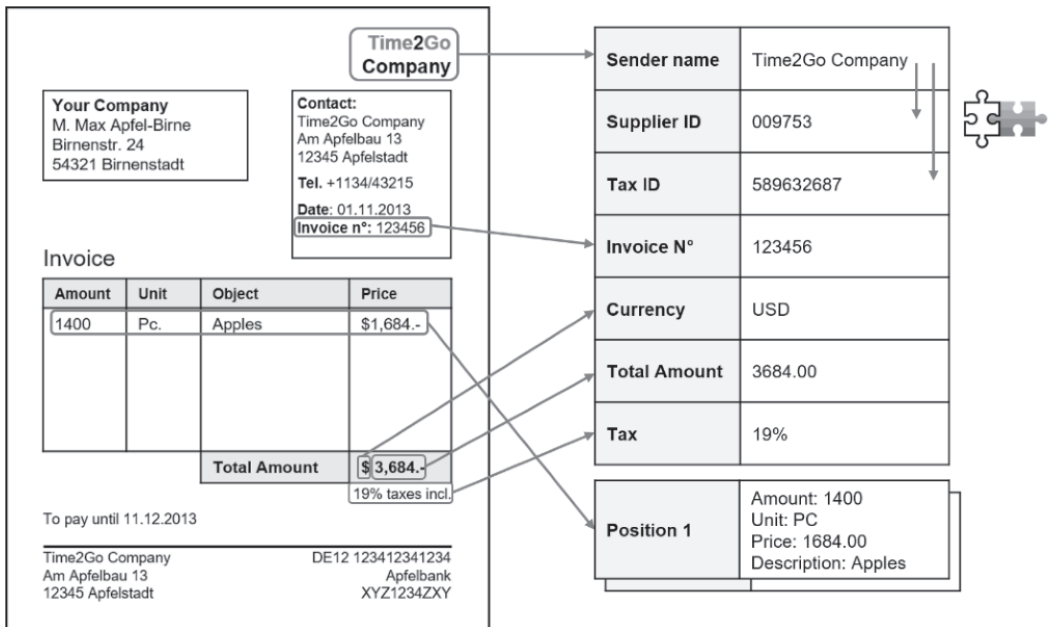


Figure 4: Process of adding supplier ID to extracted fields via an enrichment functionality based on the extracted sender name.

The matching process can be implemented with a pairwise classification model based on decomposable attention (Parikh et al. 2016). Given a query string (or document), it classifies whether there is a match with a target string (or document).

4 Critical success factors for ML adoption by enterprises

4.1 Research Methodology

Based on the technical approaches to ML-based business document processing, this section turns to the business side and explores the adoption challenges of ML technologies such as IDP. The literature review revealed that the previous research on the adoption challenges focused mainly on the buyer side. Our goal is to extend the existing body of management knowledge by providing a holistic view, from the supplier and the buyer perspective. The analysis of the interrelationship between both can disclose additional insights, as opposed to an isolated viewpoint. The aim is to discover key patterns from business practice by applying qualitative research based on interviewing methods to acquire empirical data in reference to Corbin/Strauss (2014).

The research question at hand is: "Which factors are critical for ML adoption by enterprises?". We conducted semi-structured, in-depth interviews with five selected SAP experts responsible for AI Business Services. Three of the experts were chosen from the leadership team. Two experts were responsible for the productization, commercialization, enablement, and related marketing processes. The interviews were based on a semi-structured interview guide to provide a schematic presentation of questions and topics (Corbin/Strauss 2014). For that purpose, we applied the *Business Model Canvas* by Osterwalder/

Pigneur (2010). The questionnaire design was based on the nine building blocks of this model: value proposition, customer segments, customer channels, customer relationships, revenue streams, key activities, key partners, key resources, and cost structure.²

All interviews were conducted via video conferences in August 2020. The qualitative data analysis was based on transcripts of interview notes and on three coding cycles of these transcripts: open, axial, and selective coding (Strauss 1987). We defined and discussed themes and concepts that emerged during the analysis. The main outcomes are the critical success factors (CSFs) for suppliers and buyers as well as the interrelationships between them.

4.2 Results of the expert interviews

Our analysis revealed two primary CSFs: business value and adoption. Adoption is the primary CSF and at the same time a goal of the supplier. Whereas, business value is the primary CSF and a goal of the buyer. Additionally, we identified a range of reinforcing CSFs on the supplier and the buyer side that have direct or indirect impact on business value and adoption. Figure 5 depicts the interdependencies between the primary CSFs and the reinforcing CSFs.

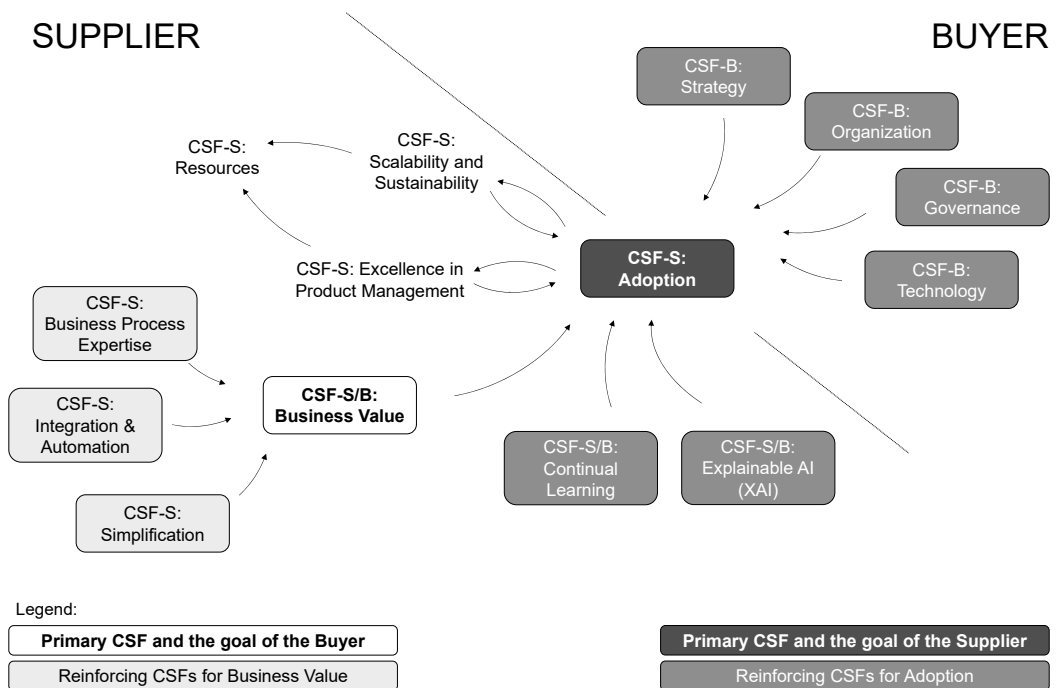


Figure 5: Identified interdependencies between the revealed Critical Success Factors.

² See the questionnaire in the appendix.

4.2.1 Critical Success Factors for Suppliers

Ten CSFs relevant to the supplier side (see Table 1) were identified. The three key findings for the supplier are discussed below.

Table 1: Critical Success Factors for Suppliers.

#	Category	Critical Success Factor	Definition
1	CSF-S	Adoption	It encompasses implementation, usage, and feedback. Adoption can be achieved by creating a perceivable business value for the buyer.
2	CSF-S/B	Business value	It corresponds with the determination of the relevant business use case, automation potential, and the expected business outcome.
3	CSF-S	Business process expertise	It refers to the combination of the technology expertise (concepts, data requirements, ML) and deep process domain expertise (business use case). It is essential to providing business value.
4	CSF-S	Integration and automation	Embedding ML into standard products through end-to-end automation creates business value.
5	CSF-S	Simplification	It leads to higher quality (product standards, data compliance) and lower entry barriers to ML for enterprises.
6	CSF-S/B	Continual learning	Ensure gradual improvement of ML models without neglecting the obtained knowledge from preceding trainings. It is expected to be embedded into the product by the supplier.
7	CSF-S/B	Explainable AI (XAI)	Consider XAI as a compliance requirement in terms of auditability, local regulations, data security, and protection. It is expected to be part of the product provided by the supplier.
8	CSF-S	Excellence in product management	Establish strong product management to define the vision and the strategy for ML as well as to enable buyers, and to collect early feedback from buyers.
9	CSF-S	Scalability and sustainability	Adoption generates scale effects. With increasing adoption, suppliers need to consider sustainable resource consumption.
10	CSF-S	Resources	Key resources comprise the data science and engineering teams, data, IP rights, process expertise, co-innovation partners, and infrastructure. Successfully managing these resources is essential to develop superior capabilities.

Finding 1: *Business process expertise, integration, and automation as well as simplification generate business value.*

ML needs to be driven by adoption, which can be defined as the implementation and usage of the technology in the business context at scale. Adoption can be achieved by

creating a perceivable business value. However, the main challenge is finding the right purpose for it (relevant business use case) as well as delivering the technology in a compliant way, as part of product or process innovations. In this context, profound process expertise and understanding of business requirements are crucial for the generation of business value. Regarding the value drivers, all interviewees focused on efficiency and effectiveness. Efficiency is achieved through automation and increased speed of business process execution (cost savings). Effectiveness refers to higher business process quality (fewer errors and compliant process execution).

It is expected that ML addresses clear business needs and is integrated into a business process so that it can deliver efficiency, preferably in form of an end-to-end automation, simplifying the access to the technology for business users. In that context, the buyers expect 'everyday AI' that offers capabilities without long data preprocessing and training activities. This simplification can be achieved through the embedding of the technology into the process (where data resides) and improving the user experience. According to the interviewees, this leads to higher user satisfaction, i.e., the fulfillment of user expectations.

These findings are supported by Chui/Malhotra (2018) who state that most digitized firms have embedded AI into standard business processes. Thus, embedding ML directly into standard business processes (such as the core of a standard ERP-system) can generate business value and help overcome adoption challenges.

Finding 2: *Continual learning and explainable AI (XAI) lead to higher adoption. Also, scalability and sustainability as well as excellence in product management need to be considered with respect to resource efficiency.*

From the supplier perspective achieving adoption is an essential goal to generate scale effects by serving many consumers and processing high volumes of data at the same time. This leads to higher efficiency of the underlying resources required for the service provisioning, such as data, data scientists, computing, and infrastructure resources. Especially in the early stages of technology diffusion and developing standard applications for enterprises, it is crucial to avoid pitfalls such as focusing merely on value appropriation (extracting profits), extensive marketing campaigns or developing scenarios without a clear business value. Firms need to consider their own strategic emphasis, i.e., the relation between value creation (delivering products to the market) and appropriation as it can have an impact on the firm's valuation (Mizik/Jacobsen 2003). Strong product management can contribute to the right vision and the strategy of supplier's ML portfolio by continuously evaluating the feedback from buyers and exploring further options to generate business value. This results in a much more efficient allocation of the supplier's resources.

Even though the business value aspect is substantial, also other relevant adoption-related aspects need to be considered. The interviewees mentioned two additional drivers: continual learning and explainable AI (XAI). Continual learning refers to the gradual improvement of ML models without neglecting the obtained knowledge from preceding trainings. XAI is an essential compliance requirement in terms of auditability, local regulations, data security, and protection. It can be concluded that both continual learning and XAI need to be addressed via a standard product capability.

Sustainability has been discussed in terms of increasing resource usage resulting from higher adoption and the required scalability of resources. Increasing adoption implies

higher infrastructure requirements, which can result in increased costs. Although scale effects are expected to be achieved, they might be limited. Over time the infrastructure and the underlying ML concepts should be further optimized by implementing new advancements in the infrastructure and data science domain. **Finding 3:** *Develop superior capabilities and resources to achieve competitive advantage.*

Resources are essential for delivering the value proposition (Osterwalder 2004). In the ML domain, resources comprise talents, knowledgeable staff, IP rights, data, and partners. Besides data, which is necessary for training and developing algorithms, also highly educated teams and skills ranging from data scientists, developers, product managers to business process experts are essential. This is challenging as the field of ML is highly contested and characterized by an intense 'war for talents' (Ng 2016). Moreover, it has also been recognized that the engineering efficiency is important when developing ML (see also Amershi et al. 2019).

Developing superior capabilities can lead to competitive advantage. For example, the resource-based view suggests that organizations must develop unique, firm-specific core competencies by doing things differently to outperform competitors (Prahalad/Hamel 1990). This requires various organizational resources, especially in R&D. In this context, interviewees agree that partner collaboration plays a crucial role (e.g. related to the infrastructure, talent, or data). Prior research also suggests that firm-level alliances are reshaped by technological changes (Schilling 2015). Therefore, ML suppliers need a clear strategy regarding their partners to respond to uncertainty and facilitate innovation.

4.2.2 Critical Success Factors for Buyers

Many complex challenges need to be considered regarding the adoption of ML on the buyer side. The most significant are lack of clear strategy for AI, lack of talents with appropriate skill sets, and lack of available data (Chui/Malhotra 2018) as well as mindset of the leadership and the enterprise's culture (Lorica/Nathan 2019). McAfee/Brynjolfsson (2012) focus on the importance of managing the change process while adopting new technology. They identify five crucial areas: leadership, talent management, technology, decision making, company culture.

During the analysis we identified eight CSFs relating to the buyer side (see Table 2).

Table 2: Identified Critical Success Factors for Buyers.

#	Category	Critical Success Factor	Definition
1	CSF-B/S	Business Value	ML must provide a clear business value. Through better performance on process-level, enterprises can address efficiency, effectiveness as well as end customer satisfaction.
2	CSF-B	Innovation	It corresponds with the willingness and openness to innovate and to adopt. Enterprises must be convinced of the technology and the value it creates.
3	CSF-B	Culture	Enterprises need a management and employees with an adequate mindset and an agile organization.

#	Category	Critical Success Factor	Definition
4	CSF-B	Maturity	Enterprises need to be able to consume ML. This presumes the capacity to implement, prepare data, run proof of concepts etc.
5	CSF-B	Data	Data access, availability, and quality is necessary for implementing ML-based business use cases.
6	CSF-B/S	Continual Learning	ML needs to continually internalize the knowledge of employees and ultimately improve over time. Algorithms need to retain this knowledge.
7	CSF-B/S	Explainable AI (XAI)	Buyers need to understand the results provided by ML-services.
8	CSF-B	Compliance and Security	In the overall ML context, aspects of data protection, data access and usage as well as data security and privacy must be considered.

We confirm the previous findings and extend the discussion by elaborating on the importance of explainable AI (XAI) and continual learning. We propose a framework comprising strategy, organization, governance, and technology (see Figure 6).

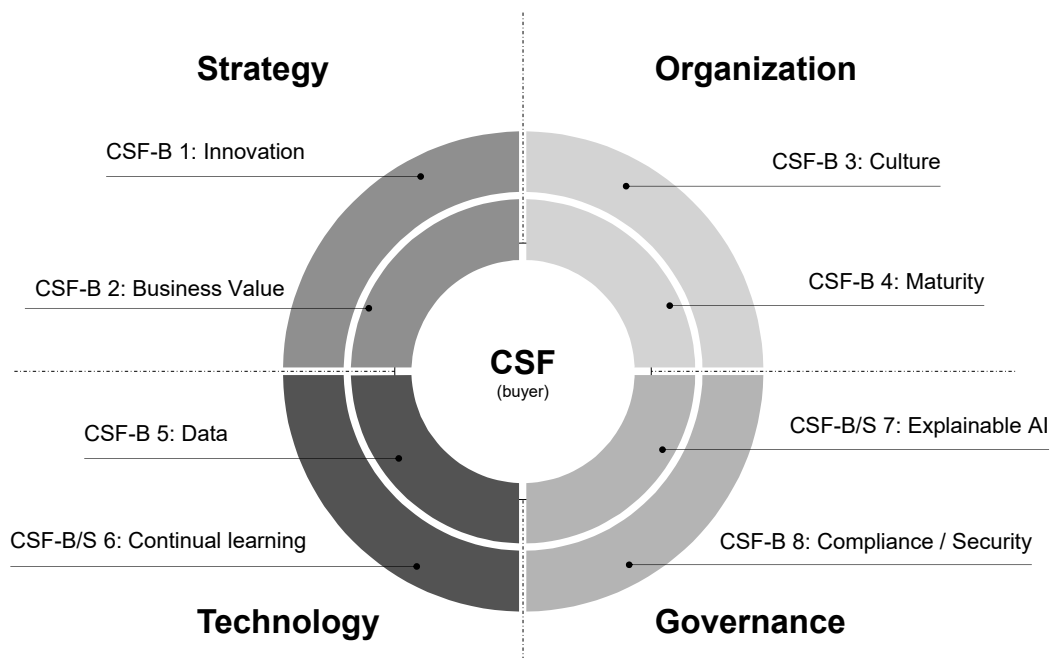


Figure 6: Framework for AI adoption by buyers.

Strategy

The strategic perspective focuses on innovation and business value. The interviewed experts shared the opinion that the adoption of ML is mainly driven by early adopters

characterized by having innovation embedded into their corporate strategy. Such companies tend to show strong openness and willingness to innovate and to adopt new technologies. Hence, ML implementation and usage require a strategic approach, i.e., setting mission, objective, strategy, and tactics.

From the buyer perspective, ML needs to provide a considerable business value which not only offsets the costs but also provides additional benefits. Through better performance on the process-level, enterprises can address efficiency (costs) and effectiveness (employee satisfaction, quality) as well as higher end customer satisfaction (revenue).

Organization

The organizational perspective encompasses the concepts of culture and organizational maturity. Culture plays a crucial role in ML adoption and thus an appropriate climate (i.e. innovation principles, resources, assets, time, etc.) must be established. This is mainly driven by the mindset of the leadership and the willingness of the employees to accept change. It needs to be derived from and interlinked with the right strategy set-up.

Enterprises with pragmatic and hands-on mentality often seem to be more agile. They understand the business needs, have well-defined use cases, and provide the necessary environment for execution (e.g. fewer organizational hurdles, empowerment of people, no fear of losing control or jobs). Buyers with extensive bureaucracy and management hierarchies without a clear innovation-driven culture seem to be reluctant and much slower in adoption.

Governance

The governance perspective encompasses the concepts of compliance and XAI. To adopt ML, companies need to ensure that it is enterprise-ready, compliant, and auditable. Since ML is data-driven, GDPR-compliance is a *conditio sine qua non*. Companies consuming ML are expected to consider all aspects related to data and its management, e.g.: data access, data storage, data protection, data usage as well as data security and privacy.

The interviewees share the opinion that XAI is one of the major CSFs for both buyers and suppliers. Buyers seek to understand the results provided by ML not only for their own sake, but also to be able to provide information to supervisory authorities and auditors (especially in regulated industries). Lack of understanding of the results prevents users from trusting in the ML-based predictions. Especially in contexts with potentially significant consequences this can lead to the rejection of the system (Rai 2020). The buyer expects that suppliers deliver these capabilities as part of their products.

Technology

The technological perspective covers the concepts of data access and continual learning. The access to data and its quality are substantial prerequisites for the success of ML. On the one hand, aspects such as data acquisition, access and transfer, anonymization, annotation, curation, secure storage, and deletion need to be considered. On the other hand, legal frameworks, and general regulations require strict compliance.

AI needs to learn and evolve over time as more data is generated or acquired. It is expected that the algorithms learn directly from users' behavior to internalize their knowledge. In this context, ML is still perceived as an assistance system, coexisting with humans, and processing the data provided by them. It is expected that this happens

automatically, in a way that human corrections are used to improve the algorithms and that this knowledge is not forgotten by the machine. This can be referred to as continual learning (Parisi et al. 2019). This way, users can expect increasing accuracy over time, build trust, and rely on the outcomes.

5 Conclusion

Recent extensive research related to information extraction in the business context indicates the interest and high potential of ML for businesses. The academic and business audience has started to appreciate and address the challenge of understanding layout information in business documents. This has led to advancements in ML research and in the ML application space. It is exemplified by the novel document representation model of Chargrid and by the proliferation of ML in the field of intelligent document processing (IDP).

In fact, IDP is an important application area in enterprises, which can be applied to many business scenarios. The estimated volume of manually processed business documents and the estimated cost savings indicate the magnitude of the efficiency potential that lies in this field. Since previous research is mainly limited to case studies or reports, we recommend conducting additional business research regarding the business value of IDP.

In this paper we also focused on how to overcome the adoption challenges of ML in enterprises. Our contribution extends the body of management knowledge by providing a holistic view on critical success factors by considering both the supplier's and the buyer's perspective.

We identified *business value* and *adoption* as the primary CSFs, whereas the first is the goal of the buyer and the latter is the goal of the supplier. Adoption means focusing on implementation, usage, and feedback and it can be achieved by creating a perceivable business value for the buyer in terms of increased efficiency and effectiveness. Business value is determined by business process expertise, integration and automation, as well as the simplification and it has to be addressed by the supplier. Suppliers also need to consider scalability of technical resources along with sustainability, resource efficiency, and excellence in product management.

Our results also suggest that buyers require explainable AI as well as continual learning and that suppliers have to include both into their offerings. Furthermore, buyers need to establish an appropriate environment for AI adoption which encompasses a range of aspects related to strategy, organization, governance, and technology.

For many enterprises ML is yet another technology for efficient and effective business execution. Buyers expect enterprise-ready and compliant 'everyday AI' integrated into their business processes. However, we observe that the business process expertise remains undervalued, and the expectations towards the technology are possibly too exaggerated. We conclude that ML in enterprises needs to be driven by business value and adoption, reinforced by additional influencing factors such as XAI and continual learning. We share the opinion that this is a promising way to ensure that ML will not remain just a hype, but it will become an actual value driver helping enterprises achieve efficiency, effectiveness, and competitive advantages.

6 Bibliography

- Amershi S., et al.* (2019): Software Engineering for Machine Learning: A Case Study, at: 2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP) (pp. 291–300), IEEE.
- Brynjolfsson, E./McAfee, A.* (2017): The business of artificial intelligence, in: Harvard Business Review, 7, 3–11, URL: <https://starlab-alliance.com/wp-content/uploads/2017/09/The-Business-of-Artificial-Intelligence.pdf>, retrieved: March 11, 2021.
- Brynjolfsson, E./Mitchell, T.* (2017): What can machine learning do? Workforce implications, in: Science, 358 (6370), 1530–1534.
- Canon Business Process Services* (2019): Five Steps to Reducing Your Paper Processing Costs, URL: <https://cbps.canon.com/insights/five-steps-to-reducing-your-paper-processing-costs>, retrieved: December 22, 2020.
- Chen, H.-C., et al.* (2012): Business Intelligence and Analytics: From Big Data to Big Impacts, in: MIS Quarterly, 36 (4), 1165–1188.
- Chui, M./Malhotra, S.* (2018): AI adoption advances, but foundational barriers remain, URL: <https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain>, retrieved: March 11, 2021.
- Clark, M.* (2020): Business Impact of Intelligent Document Processing (IDP), URL: <https://www.infrard.ai/blog/business-impact-of-intelligent-document-processing>, retrieved: March 16, 2021.
- Corbin, J./Strauss, A.* (2014): Basics of qualitative research: Techniques and procedures for developing grounded theory.
- Côte-Real, N., et al.* (2017): Assessing business value of Big Data Analytics in European firms, in: Journal of Business Research, 70, 379–390.
- Denk, T. I./Reisswig, C.* (2019): BERTgrid: Contextualized Embedding for 2D Document Representation and Understanding, at: "Document Intelligence" workshop of 33rd NeurIPS 2019.
- Garncaiek, L., et al.* (2020): LAMBERT: Layout-Aware (Language) Modeling using BERT for information extraction, URL: <https://arxiv.org/abs/2002.08087>, retrieved: March 19, 2021.
- Hyland Software Inc.* (2019): 10 Things You Need To Know About Intelligent Capture, URL: <https://www.hyland.com/-/media/Files/hyland/ebooks/brainware-ebook-ten-things-intelligent-data-capture.pdf?la=en>, retrieved: December 15, 2020.
- Jyoti, R., et al.* (2019): IDC FutureScape: Worldwide Artificial Intelligence (AI) 2020 Predictions, IDC Corporate USA, URL: <https://www.idc.com/getdoc.jsp?containerId=US45576319>, retrieved: March 17, 2021.
- Katti, A. R., et al.* (2018): Chargrid: Towards Understanding 2D Documents, at: EMNLP 2018.
- Koch, B.* (2019): The e-invoicing journey 2019–2025, URL: https://compacer.com/wp-content/uploads/2019/05/Billentis_Report_compacer_single-sponsor.pdf, retrieved: March 16, 2021.
- Lample, G., et al.* (2016): Neural Architectures for Named Entity Recognition, in: Proceedings of NAACL 2016.
- Levvel Research* (2019): Pitching ROI for Accounts Payable Automation, URL: https://cdn.levvel.io/reports/PitchingRIO_Report_Rebrand.pdf, retrieved: November 28, 2020.
- Liu, X., et al.* (2019): Graph Convolution for Multimodal Information Extraction from Visually Rich Documents, at: NAACL 2019.
- Lorica, B./Nathan, P.* (2019): AI Adoption in the Enterprise. How Companies Are Planning and Prioritizing AI Projects in Practice.

- Mansfield, E., et al. (1977): Social and Private Rates of Return From Industrial Innovation, in: Quarterly Journal of Economics, 91 (2), 221–240.
- Manyika, J., et al. (2017): A future that works: automation, employment, and productivity, URL: <https://www.mckinsey.com/~media/mckinsey/featured%20insights/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works-Full-report.ashx>, retrieved: March 11, 2021.
- McAfee, A./Brynjolfsson, E. (2012): Big Data: The Management Revolution, URL: <https://hbr.org/2012/10/big-data-the-management-revolution>, retrieved March: 16, 2021.
- Mizik, N./Jacobsen, R. (2003): Trading Off Between Value Creation and Value Appropriation: The Financial Implications of Shifts in Strategic Emphasis, in: Journal of Marketing, 67 (1), 63–76.
- Ng, A. (2016): What AI Can and Can't Do, from <https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now>, retrieved: March 16, 2021.
- Osterwalder, A. (2004): The business model ontology: A proposition in a design science approach, URL: http://www.hec.unil.ch/aosterwa/PhD/Osterwalder_PhD_BM_Ontology.pdf, retrieved: August 21, 2020.
- Osterwalder, A./Pigneur, Y. (2010): Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers.
- Palm, R. B., et al. (2017): CloudScan – A configuration-free invoice analysis system using recurrent neural networks, ICDAR 2017.
- Parikh, A. P., et al. (2016): A Decomposable Attention Model for Natural Language Inference, in: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing.
- Parisi, G. I., et al. (2019): Continual Lifelong Learning with Neural Networks: A Review, URL: <https://arxiv.org/abs/1802.07569>, retrieved March: 17, 2021.
- Prahalad, C. K./Hamel, G. (1990): The Core Competence of the Corporation, in: Harvard Business Review, 68 (3), 79–91, URL: <https://hbr.org/1990/05/the-core-competence-of-the-corporation>, retrieved: March 16, 2021.
- Prasad, S. (2020): How is Intelligent Document processing disrupting traditional workflows?, URL: <https://www.klearstack.com/intelligent-document-processing-technology/>, retrieved: November 29, 2020.
- Rai, A. (2020): Explainable AI: from black box to glass box, in: Journal of the Academy of Marketing Science, 48, 137–141.
- Reis, C., et al. (2020): Assessing the drivers of machine learning business value, in: Journal of Business Research, 117, 232–243.
- Reisswig, C., et al. (2019): Chargrid-OCR: End-to-end trainable Optical Character Recognition through Semantic Segmentation and Object Detection, at: Neur IPS 2019 Workshop Document Intelligence.
- Sanders, N. (2016): How to Use Big Data to Drive Your Supply Chain, in: California Management Review, 26–48.
- Schilling, M. A. (2015): Technology shocks, technological collaboration, and innovation outcomes, in: Organizational Science, 26 (3), 668–686.
- Strauss, A. (1987): Qualitative Analysis for Social Scientists, Cambridge University Press.
- Wei, M., et al. (2020): Robust Layout-aware IE for Visually Rich Documents with Pre-trained Language Models, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2367–2376, URL: <https://arxiv.org/abs/2005.11017>, retrieved: March 17, 2021.

Xu, Y., et al. (2019): LayoutLM: Pre-training of Text and Layout for Document Image Understanding, in: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 1192–1200, URL: <https://arxiv.org/abs/1912.13318>, retrieved: March 17, 2021.

Yu, W., et al. (2020): PICK: Processing Key Information Extraction from Documents using Improved Graph Learning-Convolutional Networks, at: International Conference on Pattern Recognition 2020, URL: <https://arxiv.org/abs/2004.07464>, retrieved: March 17, 2021.

Zhang P., et al. (2020): TRIE: End-to-End Text Reading and Information Extraction for Document Understanding, in: Proceedings of the 28th ACM International Conference on Multimedia, 1413–1422, URL: <https://arxiv.org/abs/2005.13118>, retrieved: March 17, 2021.

7 Appendix

Appendix 1: Interview questionnaire used for the in-depth expert interviews based on the Business Model Canvas by Osterwalder & Pigneur (2010).

Focus Areas	Building Blocks	Questions
VALUE	<i>Value Propositions</i>	1. What value do AI Business Services deliver to customers? 2. Which customer problems are AI Business Services helping to solve? 3. What bundles of products and services are AI Business Services offering to each customer segment? 4. Which customer needs are AI Business Services exactly satisfying?
	<i>Customer Segments</i>	5. For whom are AI Business Services creating value? 6. Who are AI Business Services’ most important customers?
CUSTOMER SIDE	<i>Customer Relationships</i>	7. What type of relationship does each of our customer segments expect AI Business Services to establish and maintain with them? 8. Which ones have AI Business Services established? 9. How are they integrated with the rest of AI Business Services’ business model? 10. How costly are they?
	<i>Channels</i>	11. Through which channels do AI Business Services’ customer segments want to be reached? 12. How are AI Business Services reaching them now? 13. How are AI Business Services’ channels integrated? 14. Which ones work best? 15. Which ones are most cost-efficient? 16. How are AI Business Services integrating them with customer routines?
	<i>Revenue Streams</i>	17. For what value are AI Business Services’ customers really willing to pay? 18. For what do they currently pay? 19. How are they currently paying? 20. How would they prefer to pay? 21. How much does each revenue stream contribute to overall revenues?

Focus Areas	Building Blocks	Questions
DELIVERY SIDE	<i>Key Partners</i>	22. Who are AI Business Services' key partners? 23. Who are AI Business Services' key suppliers? 24. Which key resources are AI Business Services acquiring from partners? 25. Which key activities do AI Business Services' partners perform?
	<i>Key Activities</i>	26. What key activities do AI Business Services' value propositions require? 27. What key activities do AI Business Services' distribution channels require? 28. What key activities do AI Business Services' customer relationships require? 29. What key activities do AI Business Services' revenue streams require?
	<i>Key Resources</i>	30. What key resources do AI Business Services' value propositions require? 31. What key resources do AI Business Services' distribution channels require? 32. What key resources do AI Business Services' customer relationships require? 33. What key resources do AI Business Services' revenue streams require?
	<i>Cost Structure</i>	34. What are the most important costs inherent in the AI Business Services business model? 35. Which AI Business Services key resources are most expensive? 36. Which AI Business Services key activities are most expensive?

Tomasz Janasz, Dr. rer. oec., ist Product Manager für Künstliche Intelligenz und Machine Learning bei SAP.

Anschrift: SAP (Schweiz) AG, CH-8058 Zürich-Flughafen, Tel.: +41 (0) 58/871-6884, E-Mail: t.janasz@sap.com

Peter Mortensen, MBA, ist Industry Expert Banking im Bereich Customer Advisory bei SAP.

Anschrift: SAP Deutschland SE & Co. KG, Tesdorpfstraße 8, 20148 Hamburg, Tel.: +49 (0) 151/4260-2926, E-Mail: peter.mortensen@sap.com

Christian Reisswig, Dr. rer. nat., ist Data Science Manager für Künstliche Intelligenz und Machine Learning bei SAP.

Anschrift: SAP SE, Münzstraße 15, 10178 Berlin, Tel.: +49 (0) 160/9081-9882, E-Mail: christian.reisswig@sap.com

Tobias Weller, B.Sc., ist Development Manager für den Bereich Dokumentenverarbeitung mit Künstlicher Intelligenz.

Anschrift: SAP SE, Dietmar-Hopp-Allee 16, 69190 Walldorf, Tel.: +49 (0) 151/1887-4867, E-Mail: tobias.weller@sap.com

Maximilian Herrmann, M.Sc., ist Project Consultant Development im Bereich Künstliche Intelligenz und Machine Learning bei SAP.

Anschrift: SAP SE, Hasso-Plattner-Ring 7, 69190 Walldorf, Tel.: +49 (0) 160/9084-1478, E-Mail: maximilian.herrmann@sap.com

Ivona Crnoja, M.A., ist Project Consultant Development im Bereich Künstliche Intelligenz und Machine Learning bei SAP.

Anschrift: SAP SE, Hasso-Plattner-Ring 7, 69190 Walldorf, Tel.: +49 (0) 160/9081-9882, E-Mail: i.crnoja@sap.com

Johannes Höhne, Dr. rer. nat., ist Data Science Expert für Künstliche Intelligenz und Machine Learning bei SAP.

Anschrift: SAP SE, Münzstraße 15, 10178 Berlin, Tel.: +49 (0) 160/9081-9882, E-Mail: johannes.hoehne@sap.com

Acknowledgments:

We are grateful to Dr. Johannes Kriebel (the editor) for guidance and support during the publication process. We also wish to thank an anonymous reviewer for his remarks and valuable feedback. Our gratitude goes to all the interviewees for their insights, as well as to internal reviewers, Leonie Wendler, Jessica Koziak, and Dr. Maria Fay.