

Drivers and Inhibitors for Organizations' Intention to Adopt Artificial Intelligence as a Service

Konstantin D. Pandl
Karlsruhe Institute of Technology
konstantin.pandl@kit.edu

Heiner Teigeler
Karlsruhe Institute of Technology
heiner.teigeler@kit.edu

Sebastian Lins
Karlsruhe Institute of Technology
sebastian.lins@kit.edu

Scott Thiebes
Karlsruhe Institute of Technology
scott.thiebes@kit.edu

Ali Sunyaev
Karlsruhe Institute of Technology
sunyaev@kit.edu

Abstract

The adoption of artificial intelligence promises tremendous economic benefits for organizations. Yet, many organizations struggle to unlock the full potential of this technology. To ease the adoption of artificial intelligence for organizations, several cloud providers have begun offering artificial intelligence as a service (AIaaS). Extant research on AIaaS exhibits a strong focus on technical aspects and has opposing views on what drives or inhibits the adoption of AIaaS within organizations. In this research, we synthesize extant research on AIaaS adoption factors and conduct semi-structured interviews with practitioners. Our research yields 12 factors that drive and another 12 factors that inhibit the adoption of AIaaS in practice. We thereby close a gap in scholarly knowledge on adopting this emerging service technology, especially on inhibiting factors, and help guide future research on related behavioral and technical aspects.

1. Introduction

Artificial intelligence (AI) is one of the most actively debated information technologies in research and practice today. Following advances in its subfield of deep learning in the early 2010s [1], applications of AI have shown tremendous potential for automation and efficiency gains. This includes, for example, self-configuring robots for logistics [2], or AI-supported medical diagnoses [3]. The adoption of AI is expected to have a positive impact on many industries and the world's economy as a whole. The McKinsey Global Institute, for example, predicts that the utilization of AI could yield an additional worldwide economic output of USD 13 trillion by 2030 [4].

However, organizations still struggle with implementing and integrating AI into their corporate IT en-

vironments. Reasons for this include a lack of AI strategy, the scarcity of AI talent [5], a lack of organizations' abilities and budgets to set up and maintain the extensive IT resources needed [6], or the challenge to collect and process data appropriately [7]. As a result, most organizations often still fail to harness the promising values of AI. A recent survey supports this view as 65% of the surveyed executives are not yet seeing value from AI investments made in recent years [8].

To address these issues of AI adoption, several cloud providers, like Amazon, Microsoft, Google, IBM, Salesforce, or SAP started to offer access to web services that enable organizations and individuals to train, develop, deploy, and manage AI algorithms in the cloud. These services became known as *Artificial Intelligence as a Service* (AIaaS) [9]. In its essence, AIaaS aims at making AI accessible and affordable, whether or not an organization is big, technologically advanced, or has large budgets to spend on AI. Recent debates in research and practice propose that AIaaS could be a valuable alternative for organizations that face difficulties with implementing in-house AI.

Whereas AIaaS comes with unique and innovative features, such as complexity abstraction and pre-trained and customizable AI models, it also introduces new challenges, such as opaque data processing and privacy infringement risks. Extant research on AIaaS mostly takes a technical perspective and investigates, for example, how to design AIaaS systems to overcome such challenges [10, 11], or how to detect security flaws in AIaaS systems [12]. However, while market research indicates that organizations are quick to adopt AIaaS (e.g., it is expected that the AIaaS market will grow by more than 42 % in 2020 [13]), research provides inconsistent rationales for AIaaS adoption. For example, extant research argues that organizations may adopt AIaaS to achieve benefits, such as gaining access to pre-trained models [14], or using automatic hyper-parameter tuning [15], whereas other research-

ers also put emphasis on the downsides of AIaaS, including privacy risks inhibiting AIaaS adoption [16]. As a consequence of this observation, we still lack an in-depth understanding of what actually drives or inhibits organizations to adopt AIaaS. Thus, to address this gap in scholarly knowledge, we ask the following research question: *What drives or inhibits organizations to adopt AIaaS?*

To answer our research question, we employ a two-staged research approach. We first review prior research on AIaaS to better understand and synthesize the inconsistent findings. Second, we conduct semi-structured expert interviews to learn more about the adoption of AIaaS in organizations and streamline opposing opinions. We build on the diffusion of innovations (DOI) theory [17] as a theoretical lens to analyze and structure our data in both stages. Our approach yields 12 factors that drive the adoption of AIaaS, and another 12 factors that inhibit its adoption. These factors relate to the AIaaS's relative advantage (e.g., higher-quality AI models), observability (e.g., transparent pricing models), compatibility (e.g., risks of privacy infringements), complexity (e.g., simple integration of AI services), and trialability (e.g., trial periods).

With our research, we contribute to the existing body of knowledge in three ways. First, we find that prior literature might have a too optimistic view of the adoption of AIaaS and show that the adoption of AIaaS is a promising avenue for future research. Second, our results enable future technology-oriented research on AIaaS to address adoption drivers and to combat inhibiting adoption factors. Third, our factors help AIaaS providers, consultants, and organizations alike to adopt well-suited AIaaS solutions. Ultimately, our research helps to accelerate the adoption of AIaaS and AI in organizations, in general.

The remainder of this paper is structured as follows. In Section 2, we provide a brief introduction to AIaaS, related research, and the DOI theory. We then delineate our research approach in Section 3. Section 4 describes factors impacting organizations' adoption intention while we discuss our results and elaborate on the potential limitations of our study in Section 5. We end this paper with a brief conclusion in Section 6.

2. Background

2.1. Foundations of AIaaS

AIaaS not only relates to AI services available on-demand, such as chatbots using natural language processing, but also covers tools and resources needed to develop, operate, and maintain the AI service. In line with the typical cloud architecture, we highlight that AIaaS can be divided into three layers, hierarchically

organized as a stack according to the abstraction level of the capability provided: (1) AI services, (2) AI developer environments, and (3) AI infrastructures (Figure 1).

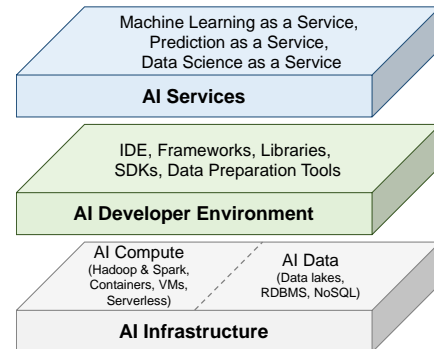


Figure 1. AIaaS stack [adapted from 18]

The most prominent and frequently used type of AIaaS are AI services that are ready-to-use building blocks or purposeful applications; relating to the conventional software as a service model. AI services are typically offered through an application programming interface (API), enabling a simple integration into existing products using various programming languages [19]. Currently, machine learning-based AI services, where users can create and use customized machine learning (ML) models, or prediction service, where users can access pre-trained machine learning models, are most prevalent in practice. For example, prediction services offer on-demand access to language services (e.g., text analytics or translation), or computer vision services (e.g., analyzing of images and videos in order to find and identify objects, text, and labels). AIaaS may also provide an easy-to-use AI developer environment, giving access to tools assisting developers with implementing code to bring out AI capabilities; relating to the conventional platform as a service cloud models. These tools not only comprise AI developer environments enabling faster coding and easier integration of APIs (e.g., PyCharm, Jupyter, or MATLAB) but also AI frameworks and software development kits comprising of diverse AI algorithms, libraries, and tools for effective exploiting of respective algorithms. Particularly, data preparation tools are offered that assist in extracting, transforming, and loading data. Finally, services offer AI infrastructures referring to the raw computational power used to build and train AI algorithms, and network and storage capacities for storing and sharing (training and inference) data; relating to the conventional infrastructure as a service model. AIaaS users have typically a wide choice of provisioning physical servers, virtual machines, containers, or AI-specialized hardware such as using GPU for computations. For example, applying

complex deep learning and neural networks might demand to complement CPUs with GPUs to enable faster calculations. Likewise, the AI infrastructure typically provides access to relational or NoSQL databases, or the capability to upload and integrate external data lakes as input to train AI models.

2.2. Related research

Research on AI is advancing rapidly and focuses, among others, on the design of in-house AI applications as well as AIaaS, typically taking a technical perspective. Also, prior research has looked at why stakeholders implement and adopt in-house AI and began discussing why organizations may use AIaaS instead. Table 1 separates related work based on their study focus (technical design, or understanding stakeholders' adoption) and the service model context (in-house AI, or AIaaS). Based on this separation quadrant A categorizes related work that focuses on the technical design of in-house AI. These studies have developed and evaluated technical concepts about in-house AI in general, for example, in the context of healthcare to improve treatment strategies by predicting and eliminating treatment failures [20], and in the context of prediction markets by examining the causal effect of social audience size on prediction accuracy [21]. Quadrant B shows that researchers have also started to develop and evaluate the technical capabilities of AIaaS. For example, related studies deployed and tested prototypes of AIaaS in the context of software-defined infrastructures and showed new business opportunities for networks by using AIaaS [11].

In regard to understanding why organizations are willing to adopt in-house AI, several studies have analyzed organizations' adoption intention (quadrant C). For example, researchers have analyzed why organizations are not willing to adopt AI and revealed several barriers like 'lack of skills' or 'lack of leadership support' [22]. Because AIaaS leads to a new way of using AI and thus also to new possibilities of adoption, extant research in the field of implementing in-house AI is valuable but not sufficient. Likewise, research on AIaaS started to discuss reasons why organizations

may adopt these novel services (quadrant D). Yet, researchers' suggestions remain inconsistent and lack empirical validation. For example, it is assumed that some organizations are concerned whether a provider has implemented adequate data governance and protection mechanisms to ensure that collected as well as AI-generated data about individuals is not used to impede their privacy [16]. While such factors are inherited from cloud adoption, which are well-studied [e.g., 23], other factors are AIaaS-specific, such as organizations value gaining access to pre-trained models [14].

2.3. Diffusion of innovation theory

To address the different views on AIaaS adoption, we aim to gain a deeper understanding of the factors driving and inhibiting AIaaS adoption by building on the (DOI) theory. Proposed by Rogers [17], it has been widely used in information systems and related research disciplines. According to DOI theory, potential users make decisions to adopt or reject an innovation based on beliefs they form about the innovation. A central concept of the DOI theory is the diffusion process, in which an innovation is communicated through certain channels, over time, among the members of a social system. An innovation is any idea, object, or practice that is perceived as new by the members of a social system. The social system consists of individuals or organizations that share a common culture and are potential adopters of the innovation. Communication channels are the means by which information is transmitted to or within the social system. Time relates to the relative speed with which an innovation is adopted by members of the social system.

DOI theory proposes five major characteristics of an innovation that influence organizations' adoption intentions: (1) relative advantage, (2) compatibility, (3) complexity, (4) observability, and (5) trialability [17]. Relative advantage refers to the degree to which an innovation is perceived as being better than its precursors. Compatibility is the degree to which an innovation is perceived consistent with the values, experiences, and needs of the potential adopters. Complexity is the degree to which an innovation is perceived as

Table 1. Related literature on artificial intelligence

		Service model context	
		In-house AI	AIaaS
Focus of study	Technical design	A <u>Typical research question:</u> „How can AI solutions be designed?“ <u>Example Studies:</u> [20, 21]	B <u>Typical research question:</u> „How can an AIaaS be designed?“ <u>Example Studies:</u> [10, 11]
	Understanding stakeholders' adoption	C <u>Typical research question:</u> „Are organizations adopting AI solutions?“ <u>Example Study:</u> [22]	D <u>Typical research question:</u> „What drives or inhibits organizations to adopt AIaaS?“ <u>Example Studies:</u> This study

relatively difficult to understand and use. The lower the complexity, the more likely organizations and their employees are to adopt a new innovation [24]. Observability is the degree to which the results of an innovation are visible to others. Finally, trialability is the degree to which an innovation can be tested or experimented with before an adoption decision [25].

3. Research approach

The research approach comprises two consecutive stages. We first conducted a literature review, followed by expert interviews. Identified AIaaS adoption factors were grouped in line with the five characteristics of an innovation according to the DOI theory.

3.1. Synthesizing prior research on AIaaS

To synthesize extant AIaaS literature, we conducted a structured database search in pertinent scientific databases, which we deemed to be representative for our topic of interest (i.e., EBSCOHost, Proquest, IEEE Xplore, and ScienceDirect) on April 22nd, 2019 with the following search string in the title: ("*Artificial Intelligence*" OR AI OR "*Machine Learning*" OR "*Neural Network**" OR "*Deep Learning*" OR "*Natural Language Processing*" OR "*Computer Vision*") AND (*Servic** OR "*as a Service*" OR **aaS* OR *Cloud* OR *Adopt**). We limited our search to publication titles because preliminary searches with the above search string indicated that many publications referred to AI in their abstracts or keywords without actually focusing on AI.

After removing duplicates, a total of 1,600 publications were left to which we applied several, predefined exclusion criteria. Specifically, we first excluded 363 publications not written in English and publications published before 2010 because the term AI in combination with cloud computing was not widely used before. We also excluded 270 non-peer-reviewed publications to ensure an adequate level of quality. Afterward, we thoroughly screened the remaining 967 publications' titles and abstracts, and excluded those that were unrelated to AI algorithms in cloud environments (in total 902 articles). For the remaining 65 publications, we read their full-texts and excluded 35 articles that did not explicitly examine AIaaS. This step resulted in a set of 30 relevant publications, for which we performed a forward and backward analysis using Google Scholar. The forward search yielded four additional relevant publications, while the backward search yielded ten additional publications, each meeting the above criteria, leading to a final set of 44 articles.

We employed the coding method of Jeyaraj et al. [26] and Lacity et al. [27] to analyze the final set and identify driving or inhibiting factors. In particular, we recorded for each extracted factor a name, a description, and its impact on the adoption intention. A list of master variables was created to aggregate the identified factors. If an identified factor fit into an existing master variable, we assigned it accordingly; otherwise, a new master variable was created. The resulting coding scheme consisted of 308 variables that were aggregated to 50 master-variables.

3.2. Expert interviews

We complemented our literature review findings with expert interviews to deepen our knowledge of AIaaS adoption and better synthesize opposing opinions in prior research. We conducted 8 semi-structured one-to-one expert interviews (Table 2). To recruit potential interviewees, we applied a purposeful sampling strategy that focused on selecting individuals who are especially knowledgeable about our phenomenon of interest (i.e., AIaaS). Consequently, we included only experts who were engaged in AIaaS activities, such as offering or using AIaaS, or consulting organizations when adopting AIaaS, and having mostly a multi-year professional experience in the AI domain.

Table 2. Information on interviewees

ID	Job position	Experience in AI	Industry
i01	CEO	3-5 years (managerial)	Service provider and consultant
i02	CEO	> 5 years (technical)	Healthcare
i03	Chief Research Officer	3-5 years (technical)	Consulting
i04	CEO	1 year (managerial, technical)	Information Technology
i05	Research Vice President	> 5 years (technical)	Technology industry analysts
i06	CEO	> 5 years (managerial, technical)	Service Provider
i07	Software Engineer	1-3 years (technical)	Consulting
i08	CEO	> 5 years (managerial, technical)	Service provider

We applied a semi-structured interview method for different reasons. A certain basic structure was necessary for our research because we aim to gather further information on identified factors from prior research. While providing such a basic structure, semi-structured interviews also leave interviewed experts with a sufficient degree of freedom to talk about aspects that might not have come to our attention during the literature review or preparation of the interview. The interview guide was derived and discussed by two re-

Table 3. Factors driving or inhibiting the adoption of AIaaS in organizations

Cluster	Factor	Source
Relative advantage	Complexity abstraction (+)	Lit, Int
	Data and model sharing capabilities (+)	Lit, Int
	Direct link to cloud storage (+)	Lit, Int
	Higher degree of automation (+)	Lit, Int
	Higher-quality ML models (+)	Lit, Int
	On-demand computing resources (+)	Lit, Int
Observability	Reduced need for in-house AI specialists (+)	Lit, Int
	Transparent service levels (+)	Lit, Int
	Difficulty to estimate total cost (-)	Int
	Difficulty to estimate ROI (-)	Int
	Missing proof of concepts (-)	Int
	Transparent pricing models (+)	Lit, Int
Complexity	Challenging deployment process (-)	Int
	Ease of use of AI services (+)	Lit, Int
	Need for AIaaS specialists (-)	Int
	Opaque data processing (-)	Lit, Int
	Simple integration of AI services (+)	Lit, Int
Compatibility	Uncertain governance and regulation compliance (-)	Int
	High network latency (-)	Lit, Int
	Increased risks of privacy infringements (-)	Lit, Int
	Lack of own training data to develop individual AI models (-)	Int
Trialability	Integration costs (-)	Int
	Trial periods (+)	Lit, Int
	Provider lock in (-)	Int

(+): factor drives adoption; (-): factor inhibits adoption
Lit = literature review findings; Int = interview findings

searchers before conducting the interviews. In addition, we made constant improvements to the questionnaire in terms of clarity and comprehensibility of the questions. The interview guide was structured as follows. First, the interviewer introduced himself and explained AIaaS and the objectives of the interview to the interviewees. Then, the interviewer asked interviewees about some basic demographics and their experience with AI projects. Subsequently, the interviewees were asked about which AIaaS adoption factors they had faced in their own projects or daily business. We applied a non-judgmental form of listening, maintained distance, and strived to sustain an open and non-directive style of conversation during the interviews to ensure impartiality and avoid bias. We recorded and transcribed each interview. The interviews lasted 29 minutes on average.

To analyze the interview data, we conducted selective, open, and axial coding [28] using ATLAS.ti 8 to facilitate this process. To determine the labels, we used words that the interviewees suggested. We first started with selective coding by assigning master-variables identified in the literature review to the interview findings to validate findings from prior research as well as gather additional information. Afterward, we performed open coding to identify new drivers or

inhibitors that have been neglected in prior research so far. For example, we coded the phrase *"most people [are] looking for automation benefits"* [i08] as the driving factor of *higher degree of automation*. Finally, axial coding was used to identify the conditions and consequences of each factor.

We also aimed to move beyond a mere description of factors to a more abstract level of conceptualization [29]. We, therefore, synthesized literature and interview findings to group similar master-variables and interview codes to more abstract categories according to common themes, thereby creating hierarchical classifications. In addition, we used DOI theory to cluster our findings into the five adoption dimensions: relative advantage, compatibility, complexity, observability, and trialability [17]. Our results reveal 12 factors that drive AIaaS adoption, and another 12 factors that inhibit its adoption (Table 3).

4. Drivers and inhibitors

4.1. Relative advantage

The combination of cloud computing and AI in AIaaS leads to unique advantages of AIaaS over in-house AI solutions. When deciding for AIaaS, *"most people [are] looking for automation benefits"* [i08]. Automation abstracts the complexity of AI and relates to both, the hardware and the software layer. On the hardware layer, AIaaS offerings typically rely on cloud infrastructures where hardware resources are automatically managed by the cloud provider, for example, for maintenance or scaling [6]. On the software layer, users only need to call (relatively) simple-to-use APIs to access the service and perform tasks, for instance, along the ML pipeline. Some AI services are capable of automatically pre-processing data (which can be *"80 % of the work"* [i08]), selecting appropriate ML model architectures for a given data set, or tuning hyperparameters for an ML model automatically [30, 31]. Thus, organizations need less AI talent, which can be rare and expensive to hire [32]. This is especially beneficial for small and medium-sized businesses because *"a lot of [these] places [...] don't have the internal skills to develop [AI solutions] in-house"* [i05].

Based on automation, the on-demand availability of computing resources is actually the most frequently mentioned driving factor in our interviews. With AIaaS being a cloud service, it inherits the strengths and typical cloud characteristics that have transformed cloud services into a critical information infrastructure for our everyday life. The most dominant among the advantages is scalability since AIaaS providers can elastically provision and release hardware resources available to the platform and thus scale horizontally in accordance with the user-defined configurations and

requirements if the consumption of computing resources for the defined AI model has increased [14]. Scalability of the cloud, combined with the number of available hardware resources, results in a large amount of processing power provisioned by the cloud and enables the AIaaS to respond to extensive requests with scalable and responsive utilization of CPUs and GPUs [33]. This is particularly beneficial, because when using AI, organizations' hardware requirements typically change frequently and quickly. For example, the training of ML models can require powerful GPU resources for a certain period of time (e.g., weeks), while the hardware requirements for the inference of ML models are typically much less. By using AIaaS, organizations can share hardware resources building on a multi-tenant architecture, thus, utilizing the hardware resources more efficiently [34].

Another frequently perceived benefit of AIaaS are the high-quality ML models. These high-quality models can actually be reached by three different causes. First, prediction services come with pre-trained models, allowing users to customize these models even with little data instances available for training [10, 14]. As a result, AIaaS can provide a higher accuracy than alternatives, such as training a plain neural network from scratch with limited data. In the case of prediction services, users often do not even need to customize the ML model. For example, commoditized tasks across vertical industries, such as speech to text processing in the English language, are readily available to use in extant AIaaS offerings. Second, relying on automated hyperparameter tuning can result in a higher quality of the ML model [31]. The third cause affects AIaaS based on shared cloud environments: such AIaaS offerings can allow different users to share data or ML models [35]. Combining data silos can increase the accuracy of AI-based systems, or enable the application of AI-based systems in the first place. A further aspect of data management of AIaaS, is the direct link to a cloud storage [36]. For organizations, which already use the cloud or are even cloud-native, AIaaS can directly access the data on the cloud, which decreases the effort required to use AIaaS.

4.2. Observability

AIaaS needs to have a high observability by members of the organization or their customers to foster adoption. Interviews revealed that a fundamental barrier for AIaaS adoption is the estimation of the impact before it gets deployed. Specifically, organizations need to know "*what [they are] going to do with [the] new information*" [i03] that AIaaS provides. Estimating benefits of AIaaS adoption is highly dependent on

the use case and industry. If AIaaS, for example, automates the programming of a robot and thereby reduces the amount of required human labor, the benefits of using AIaaS can be easily quantifiable and observable. If AIaaS, however, should be used to support but not actually replace human radiologists in analyzing medical image scans, the benefit may be less visible because the radiologists' work does not change noticeably. In many cases, therefore, members in an organization demand a proof-of-concept, which demonstrates the value of AIaaS, before they are convinced of the technology's benefits.

Likewise, interviewees report that the costs of adopting AIaaS can be difficult to estimate. This is especially the case for estimating the number of requests sent to an AIaaS beforehand and related project overhead costs. This finding contrasts not only extant research that emphasizes the budget-flexibility and cost-effectiveness when using AIaaS [37] but also cloud providers' efforts ensuring transparent service offerings. Particularly because the cost of the AIaaS offering such as API calls itself and service level agreements are typically described in a transparent manner, for example, on the website of the AIaaS provider [36, 38]. While it is reasonable to assume that adopters should be able to observe what kind of service level they can expect, interviewees stress that the impact of AIaaS adoption and the total cost are still difficult to estimate beforehand, which makes it challenging to estimate the return of investment (ROI). These issues are inherent to conventional cloud computing and related concerns are still discussed in practice today, despite the emergence of cloud computing in 2007.

4.3. Complexity

With a high degree of automation on both, the hardware and the software layer, AIaaS abstracts a large amount of complexity compared with in-house AI. As such, extant research perceives AI services as easy to use in general. In contrast, interviews revealed a more differentiated view on the perceived complexity of using AIaaS, particularly when using AI developer environments and AI infrastructures. Interviewees regard the integration of AIaaS as a double-edged sword. On the one hand, any system connected to the internet is capable of connecting with an AIaaS offering. As such, AIaaS is highly compatible with a large variety of information systems. While AIaaS are mostly offered through an API, most AI service providers also offer a graphical user interface (GUI, e.g., dashboards) to simplify operation. Typically, offered GUIs can be used for selecting, tuning, and deploying an appropriate machine learning algorithm, or to monitor key performance indicators and visualize analytics

performed over the data. On the other hand, interviewees argued that setting up an AIaaS system on provisioned AI developer environments or AI infrastructures, and integrating it into existing information system can be daunting, especially for AI novices. Likewise, legacy systems may not support the connection to or integration of cloud technologies.

Another factor that increases the complexity of AIaaS is its opaque data processing, which makes it challenging to understand and audit [32]. It occurs twofold: First, many extant AI models (e.g., deep neural networks) appear as black boxes and are not easily understandable. This aspect gets further intensified with AI services based on pre-trained models because users do not know which data instances, which model architecture, and which model hyperparameters were used for training that AI model. In the worst case the AI model *"might not actually work on [user's] data"* [i05]. Second, the opaque data processing character relates to the entire pipeline of AIaaS activities (e.g., pre-processing or data flows in general) and technology stack, as users cannot see or verify the processes running on the cloud backend. As a result, interviewees suggest that organizations may still need to employ AIaaS specialists. Ideally, these AIaaS specialists should also have domain expertise in the industry their organization is active in, which can be a rare skill set.

4.4. Compatibility

A strong factor decreasing the compatibility of AIaaS are privacy concerns, mentioned in all interviews and in prior research [e.g., 32]. Since the cloud environment and the AIaaS system process data opaquely, it is difficult to audit and know what happens with the data. Such privacy concerns are complex to resolve and prevalent *"especially in the healthcare industry"* [i06]. In particular, AIaaS may not be compliant with corporate governance guidelines or regulations in some cases.

On the technology side, AIaaS may not work well with systems that need to infer decisions in real-time, due to a potentially high latency [6]. For example, a driving assistance system based on AI should be able to perform actions in real-time and, thus, AIaaS solutions that depend on a stable internet connection may not be suitable.

Also, organizations may not have enough data to develop own AI models using AI developer environments or infrastructures [37]. Interviewees were concerned that even if they had enough data, this data is often not organized and thus, either requires a strong pre-processing effort, or may not be usable at all. Extant literature, however, stresses that in these cases organizations may access pre-trained models offered by

prediction services or make use of data pre- and post-processing functions that are handled by the AI developer environments [14].

4.5. Trialability

In addition to the generally high simplicity of AI services, most AIaaS providers describe their offerings on their website and provide educational resources about their offerings, for example, through video tutorials. More importantly, most AIaaS providers offer a trial period which potential adopters can use to try out AIaaS. In some application scenarios, especially where the information system uses commoditized prediction APIs that are easy to integrate, adopters can choose one provider out of many possible providers and try out their solution while still having the flexibility to switch to another provider easily later on.

If, however, a user does not want to try out a commoditized prediction service but more complex offerings such as machine learning as a service (MLaaS), the process of trying the MLaaS may be more complex. The user may need to learn about the specific MLaaS offering, integrate it with a complex process into their information system, and potentially upload large amounts of data into the cloud storage. Thus, such integration efforts and costs can inhibit users from trying out AIaaS at all. In such cases, AIaaS offerings may even result in provider lock-in effects, where users cannot easily switch to another AIaaS.

5. Discussion

5.1 Principal findings

Given opposing views on the adoption of AIaaS, we identify this as a promising field of research. In our literature review, we identified 12 factors that drive the adoption of AIaaS, and 3 factors that inhibit the adoption of AIaaS. Based on our interviews, we can confirm these 15 factors, however, we identified additional 9 factors that inhibit the adoption of AIaaS, and relate to five DOI dimensions.

Our findings suggest that prior scholarly literature has drawn a too optimistic picture of the adoption of AIaaS, and that the adoption is more challenging in practice than theory suggests. One reason for this overly optimistic view may be that prior research mostly covers AI services and applications, and thus AIaaS literature solely relates to the conventional software as a service cloud model. On the contrary, cloud providers already began offering AI developer tools and AI infrastructure services, relating to the conventional platform- and infrastructure as a service cloud models that have been, however, neglected by prevalent research so far.

In general, we find that organizations are open to consider adopting AIaaS. This is true for both, small and medium-sized organizations as well as for large organizations, especially if their core field of business is not IT and, thus, in-house developed AI solutions are not viable. The most frequently discussed factors driving the adoption of AIaaS are the on-demand availability of computing resources, and factors that build on the automation potential of AIaaS. On the other hand, information privacy concerns, and the challenge to estimate the impact of AI were the most frequently mentioned factors inhibiting the adoption of AIaaS. Some factors are perceived as both, positive and negative. For example, the integration of AIaaS is typically much simpler than the integration of in-house AI systems. However, interviewees warned us that it can be complex because users still have to get familiar with the environment of the AIaaS provider.

Interestingly, some of these frequently discussed factors map well with ongoing research on ML. For example, the field of automated machine learning, also known as AutoML, aims to automate the deployment of ML systems [39]. Thus, advancements in this field could be worth integrating into AIaaS offerings and are likely to be perceived well by potential AIaaS users. Another example is the field of privacy-preserving ML, which aims to preserve information privacy when training or inferring ML models [40]. One emerging approach to reach this, for example, is by using hardware-assisted trusted execution environments, which provide an isolated environment for the confidential processing of information [41]. Some cloud providers recently started offering these for general-purpose computing (e.g., Microsoft's Azure Confidential Computing). Though AIaaS is a fast-moving field, there is a lag between cutting-edge research and AIaaS offerings in practice and between the release of new, innovative AIaaS offerings and the awareness of users. In line with DOI theory [17], adoption intentions might change over time once the maturity level of AIaaS increases, and prevalent challenges are solved.

The factors impacting the adoption of AIaaS typically strongly depend on the service type, industry, the use case, and the specific AIaaS in question. For example, information privacy concerns are prevalent across the healthcare industry, and in certain use cases (e.g., analyzing medical images) practitioners have doubts concerning the significant impact of deploying AIaaS or AI-based solutions in general. Another example are low latency requirements, which can be important for some cyber-physical systems that require information processing in real-time, such as robots.

5.2 Implications for research and practice

With our focus on the adoption of AIaaS, we study a novel field of research that is still characterized by opposing views. Our principal findings impose implications on both, researchers and practitioners. By comparing findings from extant research with findings from our interviews, we were able to validate many factors identified in prior research. However, the deployment of AIaaS into an existing information system is often more complex in practice than research and theoretical expectations suggest. As such, we identified 9 new factors inhibiting the adoption of AIaaS in organizations that were not discussed in extant research (cf. Table 3). Thus, future research on AIaaS could focus on overcoming the manifold inhibiting factors we identified. In particular, our results could help to guide future research on developing new, innovative AIaaS that could overcome these factors.

For practitioners, the implications of our research depend on the different stakeholders in the context of AIaaS, such as providers, consultants, or adaptors. Providers could use our results to develop AIaaS offerings that can be well-adopted. For this, providers need to follow and integrate the latest research results in several subfields of AI, such as AutoML or privacy-preserving machine learning. Since AIaaS offerings can be complex, providers need to communicate and explain their innovations well to enable potential adoption. As we identified AIaaS as a very dynamic field that comes with the risk of a provider lock-in, consultants should know the AIaaS offerings of several providers to propose the most well-suited AIaaS offering for their client. Otherwise, a client may be locked-in into an offering that is not ideal for them.

Organizations that consider using AIaaS need to critically evaluate whether the application of AI is generally beneficial for the specific use case. As AIaaS is a dynamic field and the latest developments may strongly influence its perception, organizations should be open to the latest innovations. For example, new innovations in the field of AIaaS promise confidential data processing and could combat privacy concerns.

5.3 Limitations and future research

Our study has several limitations. There are different stages of technology adoption and diffusion. Based on our research, we are in the beginning of AIaaS adoption. Not only established cloud providers offer AIaaS, but also start-ups and small and medium-sized enterprises are following the trend and provide unique services tailored to the needs of various industries. As more and more providers offer AIaaS, more organizations from every industry will be able to find solutions

that fit their specific use-cases making AI adoption more global and AIaaS even more compelling. Consequently, identified factors that impact an organization's decision to adopt AIaaS might change over time as AIaaS diffuses. For example, Microsoft recently integrated trusted execution environments in their cloud offerings, which could impact the privacy risk perception of AIaaS offerings based on these.

Furthermore, our study has limitations concerning the number and depth of interviews we conducted to gather necessary data. While we conducted eight expert interviews, future research might focus on gathering more information on specific findings to increase understanding. In addition, we chose a key informant method for data collection and focused on informants on a managerial level. This approach, while having advantages for our exploratory work, has the limitation that the data reflects perceptions of one person per company. Likewise, interviewees may have found it difficult to verbalize drivers and inhibitors of AIaaS adoption given its novelty. Moreover, we analyzed data based on how we interpreted it. Nevertheless, we confidently believe that we reduced potential interpretation bias by constantly discussing our interim findings with the research team and comparing the literature and interview findings.

With this study, we tap into a new field of behavioral aspects of AIaaS research, leaving plenty of opportunities for further ongoing research (e.g., conceptualizing our findings differentiated by industry or region, or trust in AIaaS [42]). Given that AIaaS is a broad field, future research could study the adoption of specific technologies in the context of AIaaS technologies, such as explainable ML, transfer learning, or privacy-preserving learning, or the convergence of AIaaS with further emerging technologies such as decentralized marketplaces empowered by distributed ledger technology and trusted hardware [43].

6. Conclusion

In this study, we synthesize prior research on AIaaS adoption and interview AI experts to factors that drive or inhibit the adoption of AIaaS in organizations. We identified 12 factors driving the adoption, and another 12 factors inhibiting the adoption. These factors relate to the AIaaS's relative advantage, observability, compatibility, complexity, and trialability. Comparing extant literature with our interview findings, we reveal that prior literature might have a too optimistic view on the adoption of AIaaS. With our study, we show that the adoption of AIaaS is a promising avenue for future research to support organizations and to inform research on AIaaS to further improve driving adoption factors and to combat inhibiting adoption factors.

7. References

- [1] LeCun, Y., Bengio, Y., and Hinton, G., "Deep learning", *Nature*, 521(7553), 2015, pp. 436–444.
- [2] Mahler, J., Matl, M., Liu, X., Li, A., Gealy, D., and Goldberg, K., "Dex-Net 3.0: Computing Robust Vacuum Suction Grasp Targets in Point Clouds Using a New Analytic Model and Deep Learning", in *Proc. of IEEE ICRA (Brisbane, QLD, 2018)*.
- [3] Madani, A., Arnaut, R., Mofrad, M., and Arnaut, R., "Fast and accurate view classification of echocardiograms using deep learning", *NPJ digital medicine*, 1(6), 2018.
- [4] Bughin, J., Seong, J., Manyika, J., Chiu, M., and Joshi, R., "Notes from the AI frontier: Modeling the impact of AI on the world economy". <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy> (retrieved 2020/07/13).
- [5] Chui, M. and Malhotra, S., "AI adoption advances, but foundational barriers remain". <https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain> (retrieved 2020/07/13).
- [6] Romero, F., Li, Q., Yadwadkar, N.J., and Kozyrakis, C., "INFaaS: Managed & Model-less Inference Serving". <https://arxiv.org/pdf/1905.13348.pdf> (retrieved 2020/03/31).
- [7] Davenport, T.H. and Bean, R., "Big Data and AI Executive Survey 2019". <http://newvantage.com/wp-content/uploads/2018/12/Big-Data-Executive-Survey-2019-Findings-Updated-010219-1.pdf> (retrieved 2020/07/13).
- [8] Ransbotham, S., Khodabandeh, S., Fehling, R., LaFountain, B., and Kiron, D., "Winning with AI", *MIT Sloan Manag. Rev.*, 61180, 2019.
- [9] Rai, A., Constantinides, P., and Sarker, S., "Next-Generation Digital Platforms: Toward Human-AI Hybrids", *MIS Q*, 43(1), 2019, pp. iii–ix.
- [10] Boag, S., Dube, P., El Maghraoui, K., Herta, B., Humer, W., Jayaram, K.R., Khalaf, R., Muthusamy, V., Kalantar, M., and Verma, A., "Dependability in a Multi-tenant Multi-framework Deep Learning as-a-Service Platform", in *Proc. of IEEE DSN-W (Luxembourg, Luxembourg, 2018)*.
- [11] Parsaeefard, S., Tabrizian, I., and Leon-Garcia, A., "Artificial Intelligence as a Service on Software-Defined Infrastructure", in *Proc. of IEEE CSCN (Granada, Spain, 2019)*.
- [12] Tramèr, F., Zhang, F., Juels, A., Reiter, M.K., and Ristenpart, T., "Stealing machine learning models via Prediction APIs", in *Proc. of USENIX Security (Austin, TX, USA, 2016)*.
- [13] Infiniti Research Ltd, "Artificial Intelligence-as-a-Service (AIaaS) Market by End-user and Geography - Forecast and Analysis 2020-2024". <https://www.technavio.com/report/artificial-intelligence-as-a-service-market-industry-analysis> (retrieved 2020/07/13).

- [14] Elshawi, R., Sakr, S., Talia, D., and Trunfio, P., “Big Data Systems Meet Machine Learning Challenges: Towards Big Data Science as a Service”, *Big Data Research*, 14, 2018, pp. 1–11.
- [15] Wang, W., Gao, J., Zhang, M., Wang, S., Chen, G., Ng, T.K., Ooi, B.C., Shao, J., and Reyad, M., “Rafiki: Machine Learning As an Analytics Service System”, *Proc. of VLDB Endowment*, 12(2), 2018, pp. 128–140.
- [16] Rouhani, B.D., Hussain, S.U., Lauter, K., and Koushanfar, F., “ReDCrypt: Real-Time Privacy-Preserving Deep Learning Inference in Clouds Using FPGAs”, *ACM TRET*, 11(3), 2018, 1–21.
- [17] Rogers, E.M., *Diffusion of innovations*, 4th edn., Free Press, New York, 1995.
- [18] Janakiram, M.S.V., “AI Stack”.
<https://www.forbes.com/sites/janakiramsv/2018/02/22/the-rise-of-artificial-intelligence-as-a-service-in-the-public-cloud/> (retrieved 2020/03/31).
- [19] Xu, D., Wu, D., Xu, X., Zhu, L., and Bass, L., “Making Real Time Data Analytics Available as a Service”, in *Proc. of ACM SIGSOFT QoSA* (Montreal, Canada, 2015).
- [20] Meyer, G., Adomavicius, G., Johnson, P.E., Elidrissi, M., Rush, W.A., Sperl-Hillen, J.M., and O’Connor, P.J., “A Machine Learning Approach to Improving Dynamic Decision Making”, *INF SYST RES*, 25(2), 2014, pp. 239–263.
- [21] Qiu, L. and Kumar, S., “Understanding Voluntary Knowledge Provision and Content Contribution Through a Social-Media-Based Prediction Market: A Field Experiment”, *INF SYST RES*, 28(3), 2017, pp. 529–546.
- [22] Sulaiman Abdallah Alsheibani, Yen Cheung, and Chris Messom, “Factors Inhibiting the Adoption of Artificial Intelligence at organizational-level: A Preliminary Investigation”, in *Proc. of AMCIS* (Cancun, Mexico, 2019).
- [23] Schneider, S. and Sunyaev, A., “Determinant factors of cloud-sourcing decisions: reflecting on the IT outsourcing literature in the era of cloud computing”, *Journal of Information Technology*, 31(1), 2016, pp. 1–31.
- [24] Tornatzky, L.G. and Klein, K.J., “Innovation characteristics and innovation adoption-implementation: A meta-analysis of findings”, *IEEE Trans. Eng. Manag.*, EM-29(1), 1982, pp. 28–45.
- [25] Teigeler, H., Lins, S., and Sunyaev, A., “Drivers vs. Inhibitors - What Clinches Continuous Service Certification Adoption by Cloud Service Providers?”, in *51st HICSS* (Honolulu, HI, USA, 2018).
- [26] Jeyaraj, A., Rottman, J.W., and Lacity, M.C., “A Review of the Predictors, Linkages, and Biases in IT Innovation Adoption Research”, *J Inf Technol*, 21(1), 2006, pp. 1–23.
- [27] Lacity, M.C., Khan, S., Yan, A., and Willcocks, L.P., “A review of the IT outsourcing empirical literature and future research directions”, *J Inf Technol*, 25(4), 2010, pp. 395–433.
- [28] Corbin, J.M. and Strauss, A.L., *Basics of qualitative research: Techniques and procedures for developing grounded theory*, SAGE, Los Angeles, 2015.
- [29] Urquhart, C., Lehmann, H., and Myers, M.D., “Putting the ‘theory’ back into grounded theory: guidelines for grounded theory studies in information systems”, *Inform Syst J*, 20(4), 2010, pp. 357–381.
- [30] Reif, M., Shafait, F., Goldstein, M., Breuel, T., and Dengel, A., “Automatic classifier selection for non-experts”, *Pattern Analysis and Applications*, 17(1), 2014, pp. 83–96.
- [31] Yao, Y., Xiao, Z., Wang, B., Viswanath, B., Zheng, H., and Zhao, B.Y., “Complexity vs. Performance: Empirical Analysis of Machine Learning As a Service”, in *Proc. of IMC* (London, UK, 2017).
- [32] Truex, S., Liu, L., Gursoy, M.E., Yu, L., and Wei, W., “Demystifying Membership Inference Attacks in Machine Learning as a Service”, *IEEE T SERV COMPUT*, forthcoming, 2019.
- [33] Bao, B., Xiang, Y., Li, Y., Lyu, S., Munshi, H., and Zhu, H., “Scalable Cloud Service For Multimedia Analysis Based on Deep Learning”, in *Proc. of IEEE ICMEW* (San Diego, CA, USA, 2018).
- [34] Shaukat, Z., Fang, J., Azeem, M., Akhtar, F., and Ali, S., “Cloud Based Face Recognition for Google Glass”, in *Proc. of ICCAI* (Chengdu, China, 2018).
- [35] Dorard, L., Reid, M.D., and Martin, F.J., “AzureML Anatomy of a machine learning service”, in *Proc. of PAPIs* (Sydney, Australia, 2016).
- [36] Kaplunovich, A. and Yesha, Y., “Cloud big data decision support system for machine learning on AWS: Analytics of analytics”, in *Proc. of IEEE Big Data* (Boston, MA, USA, 2017).
- [37] Arnaldo, I., Veeramachaneni, K., Song, A., and O’Reilly, U., “Bring Your Own Learner: A Cloud-Based, Data-Parallel Commons for Machine Learning”, *IEEE COMPUT INTELL M*, 10(1), 2015, pp. 20–32.
- [38] Gujarati, A., Elnikety, S., He, Y., McKinley, K.S., and Brandenburg, B.B., “Swayam: Distributed Autoscaling to Meet SLAs of Machine Learning Inference Services with Resource Efficiency”, in *Proc. of ACM/IFIP/USENIX Middleware* (Las Vegas, NV, USA, 2017).
- [39] Hutter, F., Kotthoff, L., and Vanschoren, J., *Automated machine learning: methods, systems, challenges*, Springer Nature, 2019.
- [40] Lu, W.-j. and Sakuma, J., “More Practical Privacy-Preserving Machine Learning As A Service via Efficient Secure Matrix Multiplication”, in *Proc. of WAHC* (Toronto, Canada, 2018).
- [41] Jauernig, P., Sadeghi, A.-R., and Stapf, E., “Trusted Execution Environments: Properties, Applications, and Challenges”, *IEEE Security & Privacy*, 18(2), 2020, pp. 56–60.
- [42] Thiebes, S., Lins, S., and Sunyaev, A., “Trustworthy artificial intelligence”, *Electronic Markets*, 2020, pp. 1–18.
- [43] Pandl, K.D., Thiebes, S., Schmidt-Kraepelin, M., and Sunyaev, A., “On the Convergence of Artificial Intelligence and Distributed Ledger Technology: A Scoping Review and Future Research Agenda”, *IEEE Access*, 2020.