

Deploying AI Applications to Multiple Environments: Coping with Environmental, Data, and Predictive Variety

Short Paper

Michael Weber

Technical University of Munich
Boltzmannstrasse 3
85748 Garching, Germany
mic.weber@tum.de

Maximilian Pfeiler

Technical University of Munich
Boltzmannstrasse 3
85748 Garching, Germany
maximilian.pfeiler@tum.de

Andreas Hein

Technical University of Munich
Boltzmannstrasse 3
85748 Garching, Germany
andreas.hein@tum.de

Jörg Weking

Queensland University of Technology
2 George St, Brisbane City
QLD 4000, Australia
joerg.weking@qut.edu.au

Helmut Krcmar

Technical University of Munich
Boltzmannstrasse 3
85748 Garching, Germany
helmut.krcmar@tum.de

Abstract

Deploying Artificial Intelligence (AI) proves to be challenging and resource-intensive in practice. To increase the economic value of AI deployments, organizations seek to deploy and reuse AI applications in multiple environments (e.g., different firm branches). This process involves generalizing an existing AI application to a new environment, which is typically not seamlessly possible. Despite its practical relevance, research lacks a thorough understanding of how organizations approach the deployment of AI applications to multiple environments. Therefore, we conduct an explorative multiple-case study with four computer vision projects as part of an ongoing research effort. Our preliminary findings suggest that new environments introduce variety, which is mirrored in the data produced in these environments and the required predictive capabilities. Organizations are found to cope with variety during AI deployment by 1) controlling variety in the environment, 2) capturing variety via data collection, and 3) adapting to variety by adjusting AI models.

Keywords: Artificial intelligence, machine learning, development, deployment

Introduction

Artificial Intelligence (AI) presents a vast opportunity for society and businesses alike (Stone et al. 2016). AI as a research field is concerned with making machines intelligent (Russell and Norvig 2021) and has developed a broad set of constantly evolving technologies, including machine learning, computer vision, and robotics (Berente et al. 2021; Stone et al. 2016). Organizations can apply AI technology in various ways, such as process automation, decision support, and human engagement (Davenport and Ronanki 2018). However, many organizations struggle with the implementation of AI technology (Benbya et al. 2020; Tarafdar et al. 2019). While many organizations manage to develop early pilot applications, there appears to be a large gap between those early pilots and actual productive deployment (Benbya et al. 2020). Productive deployments require considerable additional efforts, such as integrating the AI application in existing workflows and ensuring continuous monitoring (Baier et al. 2019; Paleyes et al. 2022). Furthermore, the productive environment of AI applications is typically more complex than the initial lab environment, which could lead to additional adjustments required (Paleyes et al. 2022).

To use AI applications at full scale and amortize development costs, larger organizations seek to deploy certain AI applications not only to one productive environment, but to multiple environments, such as different business units or firm branches. For example, an industrial firm might deploy an AI application for predictive maintenance to multiple machines at different factories. However, each of the factories could feature different setups and external influences that cause different maintenance patterns. Therefore, an AI application developed for one machine might not necessarily generalize to another machine. This typically forces organizations to conduct a costly and time-intensive retraining (e.g., Demlehner and Laumer 2020). Proven approaches from traditional software development to handle multiple deployment contexts, such as systematic code reuse or software variant management, require adaptation for AI and are yet to be broadly established in practice (Paleyes et al. 2022). This raises the question of how organizations approach the deployment of AI applications to multiple environments.

Assessing the current literature, Information Systems (IS) scholars have found important success factors, challenges, and emergent practices for AI deployment from an organizational perspective (e.g., Jöhnk et al. 2021; Pumplun et al. 2019) and project perspective (e.g., van den Broek et al. 2021; Zhang et al. 2020). While these studies greatly contribute to our understanding of AI deployment, prior literature has focused on projects that deploy AI applications to one environment instead of multiple environments. Therefore, we currently lack a detailed understanding of the context of multiple deployment environments, what unique challenges this context might incorporate, and how organizations approach such projects in the process of AI deployment. Considering the practical relevance of AI deployment to multiple deployment environments, especially for larger organizations, we aim to address this gap and ask the following research question: *How do organizations deploy AI applications to multiple environments?*

To answer this question, we employed an explorative multiple-case study of AI deployment projects in the context of a large industrial firm (Eisenhardt 1989; Yin 2018). So far, we have investigated four computer vision projects that deploy AI applications to multiple environments as part of an ongoing research endeavor. We rely on semi-structured interviews with key informants as primary data for data collection. In addition, we triangulated this data with secondary data (Yin 2018), which includes project reports, website information, and scientific papers on the respective computer vision techniques. For preliminary data analysis, we engaged in an iterative process of within- and cross-case analysis and the constant comparison of emergent hypotheses, data, and prior research (Eisenhardt 1989).

Our preliminary findings indicate that the deployment of AI applications to multiple environments is strongly influenced by environmental changes (e.g., background information, camera positions). This environmental variety is mirrored in the data produced in these environments and the AI application's required predictive capabilities. In addition, in the process of AI deployment, we found that organizations employ at least three coping mechanisms to address variety: 1) controlling variety in the environment, 2) capturing variety via data collection, and 3) adapting to variety by adjusting AI models.

This study is part of an ongoing research endeavor that aims to contribute to our understanding of the specifics of AI technology and its implications for IS deployment (Berente et al. 2021). Specifically, we complement literature on AI deployment (e.g., Paleyes et al. 2022; Zhang et al. 2020) by investigating the context of multiple deployment environments. In the future, we plan to advance this study by continuing data collection and deriving a comprehensive process model (Langley 1999) of how organizations deploy AI

applications to multiple environments. We aim to provide practitioners interested in AI deployment with a better understanding of the challenges and approaches associated with multiple deployment environments.

Background

Artificial Intelligence and Computer Vision

The term AI can be defined in two ways: First, AI refers to a broad and long-established research field that is concerned with making machines intelligent (Russell and Norvig 2021; Stone et al. 2016). The goal is to enable machines to perform tasks that “*we associate with human minds, such as perceiving, reasoning, learning, interacting with the environment, problem solving, decision-making, and even demonstrating creativity*” (Rai et al. 2019, p. iii). Second, AI serves as a collective term for multiple technologies, such as machine learning, deep learning, computer vision, natural language processing, knowledge-based reasoning, and robotics (Benbya et al. 2020; Stone et al. 2016). One or multiple of these technologies are applied when developing AI applications. In recent years, however, machine learning has dominated the discourse in research and practice (Russell and Norvig 2021), and contemporary AI applications typically build on machine learning (Berente et al. 2021). Hence, we use the term AI application to refer to applications that apply machine learning.

Machine learning enables machines to learn from prior experiences, for example, by generalizing from historical datasets (Russell and Norvig 2021). The result is typically a model (i.e., AI model) that produces predictions based on new input data. Deep learning presents an important subfield of machine learning, which uses multiple representation layers for generalization (Goodfellow et al. 2016). Advances in deep learning have caused major breakthroughs in many areas of AI, including natural language processing and computer vision (Stone et al. 2016). In computer vision, the goal is to make machines perceive based on visual input, such as images or videos (Russell and Norvig 2021). Typical tasks include image segmentation, activity recognition, and object recognition (e.g., identifying a cat).

Related Work on AI Deployment

Contemporary AI applications based on machine learning differ from conventional software through their 1) learning capabilities, 2) increasing autonomy, and 3) incomprehensiveness to certain audiences (Berente et al. 2021). Therefore, AI deployment requires a different approach to IS deployment. The typical AI deployment process is iterative and consists of the following steps: requirements definition, data management, AI model training and validation, and deployment (Ashmore et al. 2021; Paleyes et al. 2022).

Prior work has identified various important success factors, challenges, and emergent practices for AI deployment, both on organizational and project level: First, organizations need to familiarize themselves with AI technology (Keller et al. 2019) and formulate adequate requirements (Zhang et al. 2020). Often, AI is hard to grasp, and stakeholders develop unrealistic expectations of what AI can do (e.g., May et al. 2020; Reis et al. 2020; Weber et al. 2022). In addition, it is hard to predict beforehand whether a proposed AI use case will work in production (Dietz et al. 2021; Paleyes et al. 2022). Second, historical data needs to be collected, prepared, and managed for AI model training, which can be a tedious task (e.g., Jöhnk et al. 2021; Pumplun et al. 2019). A frequently observed problem is that historical datasets are unbalanced (Zhang et al. 2020) or deviate from the productive environment (Paleyes et al. 2022). Third, different skillsets need to be available and integrated for AI deployment, including skills in AI, data science, software, and the domain (Jöhnk et al. 2021; Weber et al. 2022). Of great importance is integrating users, who can actively shape the AI application by providing valuable domain expertise and feedback (e.g., May et al. 2020; van den Broek et al. 2021). Fourth, AI model training and validation typically follows a highly iterative procedure (Dietz et al. 2021) and demands much computing resources, which can be costly in practice (Baier et al. 2019; Paleyes et al. 2022). Fifth, organizations need to cope with additional efforts to integrate AI applications into existing workflows and systems (Baier et al. 2019; Paleyes et al. 2022; Tarafdar et al. 2019). For example, AI applications need to be constantly monitored to identify degrading performance or unexpected issues in production (Baier et al. 2019; Paleyes et al. 2022). Furthermore, users need to be properly instructed on the functioning of the AI application (Dietz et al. 2021; Reis et al. 2020).

In conclusion, AI’s unique characteristics pose new challenges to IS deployment that require investigation from a socio-technical perspective (Berente et al. 2021). While extant research has started to investigate

important success factors, challenges, and emergent practices, current studies predominately address the context of a single deployment environment. Therefore, we lack a thorough understanding of how organizations deploy AI applications to multiple deployment environments.

Research Approach

Case Overview and Data Collection

We conducted an explorative multiple-case study on projects that deploy AI applications to multiple environments (Eisenhardt 1989; Yin 2018). Given our limited understanding of the topic, a case study approach is particularly suited to explore *how* organizations deploy AI applications to multiple environments. Furthermore, studying multiple cases allows to distill potential contextual factors.

As part of an ongoing study, we investigated four computer vision projects (see Table 1 for case descriptions). The projects are being conducted within a large, internationally operating firm specializing in engineering, electronics, and industry-specific devices. Given the size and breadth of this firm, we were confident of finding a wealth of different projects that include AI deployment to multiple environments (e.g., to different factories). We decided to investigate projects that use computer vision technology for two reasons: First, this allows us to control for the variance that might stem from different AI technologies. Second, given advances in deep learning, computer vision has been one of the most important fields in AI in recent years (Stone et al. 2016), ensuring that our findings are directly relevant to a broad set of use cases.

The selection of the four projects followed theoretical considerations (Eisenhardt 1989). Initially, we selected two polar cases that either successfully managed to deploy or failed to deploy an AI application to a new environment (ERGO and SAFETY1). During early analysis of the cases, we realized that these two cases followed two approaches: 1) building one generalized AI model for all environments, and 2) adapting an AI model to fit a new environment. Therefore, we followed a replication logic and selected two more AI projects to cover each approach twice (SAFETY2 and TRAFFIC) to learn more about these two approaches.

Case	Case Description
ERGO	The goal of project ERGO is to apply computer vision for the ergonomic risk assessment of industrial workers. For this, a smartphone application is provided that has an AI model embedded. The smartphone is used to record a worker, for example, when lifting objects. The AI model is responsible for detecting the joints of the human body, upon which the ergonomic risk assessment is made. The project collaborates with a startup that develops and offers the smartphone application to multiple customers. Seventeen test users across five countries use the application. The project is currently being evaluated and will potentially be rolled out on a large scale.
SAFETY1	The goal of project SAFETY1 is to apply computer vision for the safety monitoring of industrial workers at heights at outdoor construction sites (2 meters or above). Therefore, a camera monitors the loading and unloading area. The AI model detects whether the workers follow the safety protocols, for example, wearing a helmet or having a lifeline attached. The project collaborates with an external technology provider that implemented the application. The application was initially tested within a factory and then piloted for several weeks at an outdoor construction site. The project terminated after the pilot phase as it was too costly.
SAFETY2	The goal of project SAFETY2 is to apply computer vision for the safety monitoring of industrial workers at factories. For this, multiple cameras are installed at different locations at a factory. The AI model analyzes images to detect whether workers follow the safety protocols, for example, wearing a face mask, minding speed limits, and non-pedestrian zones. The internal technology department develops the application. Currently, 20-25 cameras have been installed at a pilot factory at multiple locations. From there, a large-scale roll-out to other factories is planned. In addition, new use cases, such as detecting workers' helmets, are planned.
TRAFFIC	The goal of project TRAFFIC is to apply computer vision to monitor and analyze public traffic at intersections. For this, pre-installed cameras at intersections are used. The AI model recognizes vehicles and classifies them according to their type (e.g., car, bus, motorcycle). The internal technology department develops the application. The application was piloted at multiple intersections across three countries (Southeast Asia and Europe) and is now offered as a service.

Table 1. Case overview

So far, we have collected data from August 2021 to April 2022. As primary data, we conducted semi-structured interviews with key informants of the projects (Yin 2018). For ERGO, we interviewed the project manager (45 min), the co-founder of the external startup (45 min), and two users (30 min). For SAFETY1, we interviewed the project manager (45 min) and received text and voice messages as responses from the CEO of the external firm. For SAFETY2, we interviewed the project manager (60 min) and the computer vision research group manager (60 min). For TRAFFIC, we interviewed the computer vision research group manager (75 min). In addition, we interviewed a computer vision scientist on his/her experience across different computer vision projects (70 min). Our interview guidelines included questions on 1) the purpose and status of the project, 2) general success factors and challenges, and 3) experiences on the deployment of the AI application to new environments. The variety between environments emerged as a central theme during data collection. Therefore, we updated our guidelines to include specific questions on the differences between environments. To reduce bias from one source (Yin 2018), we also collected secondary data from three different sources: 1) internal documents (e.g., project reports and wiki entries), 2) website information in case of external collaborations, and 3) scientific papers on the applied computer vision techniques (e.g., Seo et al. 2015; Sivaraman and Trivedi 2013).

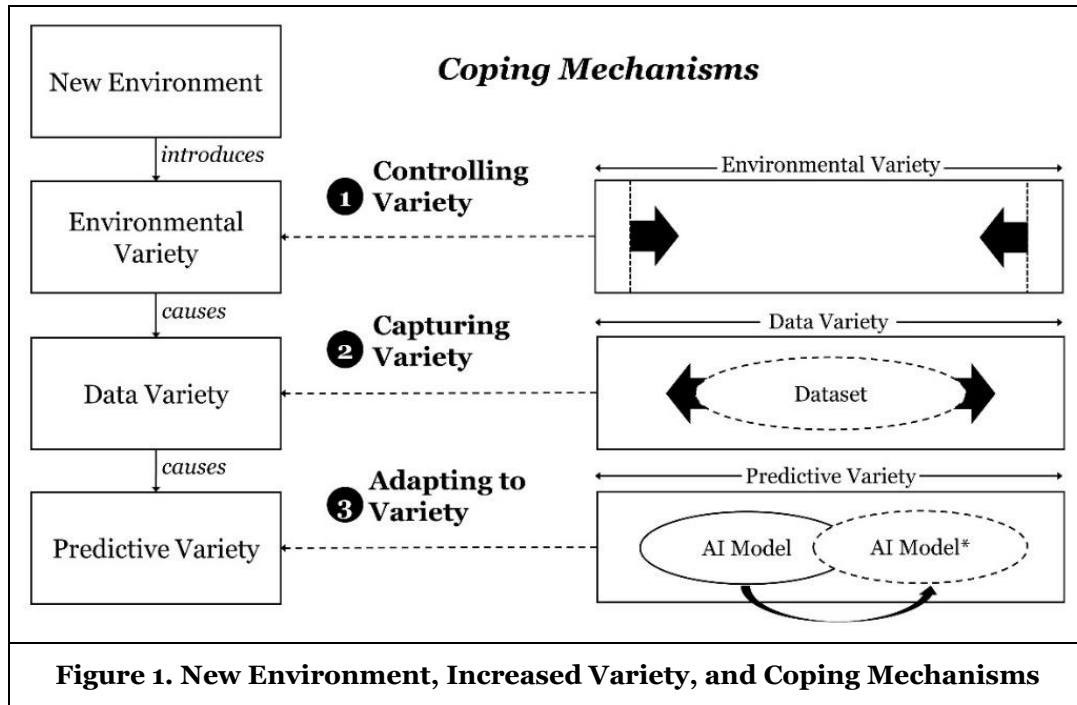
Preliminary Data Analysis

We engaged in a highly iterative process of within-case and cross-case analyses and the constant comparison of emergent hypotheses, data, and prior research (Eisenhardt 1989). First, two researchers independently analyzed all data as part of a within-case analysis. We coded the data using open, axial, and selective coding, which helps to reduce the data and support sense-making activities (Corbin and Strauss 1990). We then discussed each case within the team and constructed initial write-ups for each case. Second, we conducted a cross-case analysis by juxtaposing cases to identify similarities and differences (Eisenhardt 1989). Based on these comparisons, we developed early hypotheses of how organizations approach AI deployment to multiple environments and what factors influence its outcome.

We then evaluated and refined these hypotheses in multiple iterations by going back into the data. This procedure led us to adjust our data collection, as described in the previous section. For example, we added two cases to gain additional insights on two approaches that had emerged. Accompanying this hypotheses development, we constantly challenged the novelty of our findings regarding prior works on AI deployment. This helped us to distill potentially unique factors in the context of multiple deployment environments. For example, we decided not to deep-dive into the importance of data quality or collaboration between stakeholders, as they present more generic factors in AI deployment. After multiple iterations, we arrived at a preliminary explanatory model on how organizations deploy AI applications to multiple environments.

Preliminary Findings

We propose a preliminary explanatory model for how organizations deploy AI applications to multiple environments (see Figure 1). The model aims to explain how new environments increase variety 1) between environments, 2) in the data, and 3) in the required predictive capabilities, which complicates the deployment of AI applications to multiple environments. In addition, the model aims to explain how organizations can cope with variety by 1) controlling variety in the environment, 2) capturing variety via data collection, and 3) adapting to variety by adjusting AI models. In the following, we depict these findings in more detail and provide empirical evidence from the case study.



New Environments and Increased Variety

All four cases provide vivid evidence of how new environments increase the variety 1) between environments, 2) in the data, and 3) in the required predictive capabilities. This variety emerged as an important theme complicating the deployments. A research group manager highlights its importance:

“I am seeing the biggest problem comes with the change in the background. The objects might look differently. The lighting conditions might be different. The background clutter could be different. [...] There is a truck here, and there are people. So, this changes from one side to the other, and this is the major problem in transferring from one location to another.”

– Computer Vision Research Group Manager (SAFETY2)

Each new environment might be different: First, the camera that records input data might be of a different quality than the original training data. Second, the camera might be positioned slightly differently, producing unseen angles or distances on the scene. Third, external conditions, such as the lighting or weather effects, can alter the input data. Fourth, the object of interest might be hidden behind an interfering object. For example, in ERGO, the posture and movements of workers can disappear behind the objects that workers are lifting. Fifth, the object of interest may appear in a different shape. For example, in TRAFFIC, buses and cars in one country look different than in another country.

The environmental variety also impacts the required training data, as the training data is typically produced in these environments and should represent the environments. Hence, a greater variety in the environment causes a greater variety in the required training data. Furthermore, an AI model should provide appropriate predictions given the input data in an environment. Hence, a greater variety of the data in the environment causes a greater variety in the required predictive capabilities.

Controlling Variety

We found that organizations can control variety in the environment in some cases. Reducing the variety between environments limits the variety of possible data inputs for the AI application. One aspect that can be controlled is the positioning of the camera. For example, in ERGO, users can install the AI application on their smartphone. The users are instructed to properly record a worker from a particular angle, which reduces variety from different angles and distances to the scene. In contrast, using pre-installed, static security cameras complicate ergonomic risk assessment:

“One of the limitations currently is you need a certain vantage point to get meaningful results from the body position. So, with security cameras, it is oftentimes too far away and at the wrong vantage point to meaningfully collect information to perform these types of assessments.”

– Co-Founder External Startup (ERGO)

Other aspects that can be controlled concern background objects and external conditions. For example, monitoring worker safety can be applied at external project sites (SAFETY1) or more controlled environments like factories (SAFETY2). The project manager from SAFETY1 explained that “a factory is a controlled environment” where “you have buildings and production lines” and “everything is very defined.” In contrast, “at a project site, it is more a chaotic scenario” with “a lot of work going on in parallel”. For example, firms might store excess material at a place covering the scene. Given the challenges that were experienced at external project sites in SAFETY1, SAFETY2 decided to only deploy its safety monitoring at factories where the environment was in control of the organization. SAFETY1 was cancelled after the pilot, because fine-tuning the AI application for each environment was too costly.

Capturing Variety

We found that organizations can capture variety in some cases via data collection. By collecting a diverse dataset, they can build an AI model that generalizes well and accounts for edge cases that occur in different environments. One possible way is to collect new data for each new environment. For example, in TRAFFIC, the team collects new data for each environment to arrive at a generalized AI model:

“We had to make sure that we collect a good amount of data from our deployment sites to get the model to understand the environmental factors, the background information, and how a vehicle looks like. Then, of course, the human mind generalizes it very easily, but for the model, it essentially relies on the geometry of the bus, as well as the color, so different regions have different colors, and that made it really hard for the model to generalize.”

– Computer Vision Research Group Manager (TRAFFIC)

Besides collecting new data for each new environment, another way is to explicitly collect data for certain edge cases. For example, in ERGO, the startup collected data for situations where the human body was partially not visible to still assess the posture and movements in these situations. In doing that, the startup can serve one generalized AI model to all its customers:

“Imagine our system trying to evaluate a new video where someone is interacting with some kind of machinery where different parts of their body may be not visible. We've captured similar situations in our datasets. [...] So, when our model is looking at that video that we didn't train on, it is actually able to generalize where those joint positions are of the human body.”

– Co-Founder External Startup (ERGO)

Adapting to Variety

We found that organizations can adapt to variety in some cases by adjusting AI models. By adapting the AI application to a new environment, it learns to handle new aspects of the environment. For example, although SAFETY2 features a more controlled environment, there is still considerable variety between different factories. For example, the production lines could look different, and the pre-installed cameras could have different angles and distances to the scenes. Therefore, the team collects new data at each factory to retrain an existing set of AI models:

“This is based on a thing called transfer learning. [...] We just deploy the same set of models in the new location and retrain them. When you retrain [the model], you're essentially achieving that transferability, [...] you are transferring the weights and then you are fine tuning [the model] to adapt to the new conditions.”

– Computer Vision Research Group Manager (SAFETY2)

Similarly, SAFETY1 used existing AI models and fine-tuned them to adapt to differences in the distances to the scene or lighting conditions. However, in SAFETY1, this approach took several weeks to reach a desirable performance and was ultimately deemed too expensive.

Preliminary Discussion and Next Steps

This ongoing study aims to contribute to IS research on AI deployment by investigating the context of multiple deployment environments. So far, we started to investigate four computer vision projects within a large industrial firm. Our preliminary findings highlight how new environments introduce variety between environments. This variety is mirrored in the data produced in these environments and the AI application's required predictive capabilities. Our findings suggest that variety presents a major challenge for AI deployment to multiple environments. In essence, variety affects the generalization of AI models (Demlehner and Laumer 2020) and thereby preempts a seamless deployment.

We found that organizations apply at least three coping mechanisms to cope with variety during AI deployment. First, organizations can control variety between environments, for example, by standardizing camera placements. This stands in line with prior research highlighting the importance of existing organizational workflows and systems in AI deployment (e.g., Paleyes et al. 2022; Tarafdar et al. 2019). Our findings suggest that organizations might benefit from (re-)designing the existing workflows and systems (Tarafdar et al. 2019) in a way that provides more standardized environments for the AI application. Second, organizations can capture variety via data collection, for example, by collecting data at each new environment. Thereby, organizations aim to appropriately cover the true data distribution and relevant edge cases, a critical factor for AI deployment (Ashmore et al. 2021; Zhang et al. 2020). However, repeatedly collecting and preparing data from multiple environments enhances the complexity and effort associated with data management. This supports prior findings that organizations should build a sophisticated data management to support AI deployments (e.g., Jöhnk et al. 2021; Weber et al. 2022), for example, by systematically keeping track of what data exists and in which environment it was collected (Ashmore et al. 2021). Third, organizations can adapt to variety by adjusting AI models, for example, by fine-tuning the AI model to the new environment. Typically, it is hard to predict whether an AI application will work in the productive environment (Paleyes et al. 2022), and thus, the repeated training of AI models is a common practice in AI deployment (Dietz et al. 2021). However, our findings suggest that even if an AI application has been successfully validated in one environment, there is no guarantee that the AI application will also work in another environment. (Paleyes et al. 2022). Consequently, organizations need to again validate the AI application and potentially initiate a new AI training cycle for each new deployment environment, limiting the seamless reusability of AI applications.

In the future, we plan to advance this ongoing study as follows. First, we will continue to investigate the current projects, as some plan to deploy the AI applications to even more environments. This could reveal contextual factors of deploying AI applications at scale, such as the potential saturation of variety. Second, we consider adding additional cases to our multiple-case study to challenge and advance our emergent findings (Eisenhardt 1989). Specifically, we want to better understand the coping mechanisms' antecedents, boundaries, and outcomes. For example, we are interested in when organizations decide to control the environment and how exercising control potentially conflicts with other organizational goals. Third, we envision to derive a comprehensive process model (Langley 1999) of how organizations deploy AI applications to multiple environments. This process model could explain when organizations encounter specific challenges such as those associated with variety and under what circumstances organizations apply different coping mechanisms.

Acknowledgements

This work was in part funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project no. 464594907.

References

- Ashmore, R., Calinescu, R., and Paterson, C. 2021. "Assuring the Machine Learning Lifecycle: Desiderata, Methods, and Challenges," *ACM Computing Surveys (CSUR)* (54:5), pp. 1-39.
- Baier, L., Jöhren, F., and Seebacher, S. 2019. "Challenges in the Deployment and Operation of Machine Learning in Practice," *27th European Conference on Information Systems (ECIS)*, Stockholm & Uppsala, Sweden.

- Benbya, H., Davenport, T. H., and Pachidi, S. 2020. "Artificial Intelligence in Organizations: Current State and Future Opportunities," *MIS Quarterly Executive* (19:4), pp. ix-xxi.
- Berente, N., Gu, B., Recker, J., and Santhanam, R. 2021. "Managing Artificial Intelligence," *MIS Quarterly* (45:3), pp. 1433-1450.
- Corbin, J. M., and Strauss, A. 1990. "Grounded Theory Research: Procedures, Canons, and Evaluative Criteria," *Qualitative Sociology* (13:1), pp. 3-21.
- Davenport, T., and Ronanki, R. 2018. "Artificial Intelligence for the Real World," *Harvard Business Review* (96:1), pp. 108-116.
- Demlehner, Q., and Laumer, S. 2020. "Shall We Use It or Not? Explaining the Adoption of Artificial Intelligence for Car Manufacturing Purposes," *28th European Conference on Information Systems (ECIS)*, Virtual Conference.
- Dietz, J., Glaser, K., and Hoehle, H. 2021. "Uncertainty Reducing and Handling Strategies in ML Development Projects," *42th International Conference on Information Systems (ICIS)*, Austin, TX.
- Eisenhardt, K. M. 1989. "Building Theories from Case Study Research," *Academy of Management Review* (14:4), pp. 532-550.
- Goodfellow, I., Bengio, Y., and Courville, A. 2016. *Deep Learning*. Cambridge, MA: The MIT Press.
- Jöhnk, J., Weißert, M., and Wyrтки, K. 2021. "Ready or Not, Ai Comes—an Interview Study of Organizational Ai Readiness Factors," *Business & Information Systems Engineering* (63:1), pp. 5-20.
- Keller, R., Stohr, A., Fridgen, G., Lockl, J., and Rieger, A. 2019. "Affordance-Experimentation-Actualization Theory in Artificial Intelligence Research—a Predictive Maintenance Story," *40th International Conference on Information Systems (ICIS)*, Munich, Germany.
- Langley, A. 1999. "Strategies for Theorizing from Process Data," *Academy of Management Review* (24:4), pp. 691-710.
- May, A., Sagodi, A., Dremel, C., and van Giffen, B. 2020. "Realizing Digital Innovation from Artificial Intelligence," *41th International Conference on Information Systems (ICIS)*, Virtual Conference.
- Paley, A., Urma, R.-G., and Lawrence, N. D. 2022. "Challenges in Deploying Machine Learning: A Survey of Case Studies," *ACM Computing Surveys (CSUR)*, pp. 1-26.
- Pumplun, L., Tauchert, C., and Heidt, M. 2019. "A New Organizational Chassis for Artificial Intelligence—Exploring Organizational Readiness Factors," *27th European Conference on Information Systems (ECIS)*, Stockholm & Uppsala, Sweden.
- Rai, A., Constantinides, P., and Sarker, S. 2019. "Editor's Comments: Next-Generation Digital Platforms: Toward Human–Ai Hybrids," *MIS Quarterly* (43:1), pp. iii-x.
- Reis, L., Maier, C., Mattke, J., Creutzenberg, M., and Weitzel, T. 2020. "Addressing User Resistance Would Have Prevented a Healthcare Ai Project Failure," *MIS Quarterly Executive* (19:4), pp. 279-296.
- Russell, S. J., and Norvig, P. 2021. *Artificial Intelligence: A Modern Approach*, (4th ed.). Boston, MA: Pearson.
- Seo, J., Han, S., Lee, S., and Kim, H. 2015. "Computer Vision Techniques for Construction Safety and Health Monitoring," *Advanced Engineering Informatics* (29:2), pp. 239-251.
- Sivaraman, S., and Trivedi, M. M. 2013. "Looking at Vehicles on the Road: A Survey of Vision-Based Vehicle Detection, Tracking, and Behavior Analysis," *IEEE Transactions on Intelligent Transportation Systems* (14:4), pp. 1773-1795.
- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A., Shah, J., Tambe, M., and Teller, A. 2016. "Artificial Intelligence and Life in 2030. One Hundred Year Study on Artificial Intelligence," Report of the 2015-2016 Study Panel, Stanford University, Stanford, USA.
- Tarafdar, M., Beath, C. M., and Ross, J. W. 2019. "Using Ai to Enhance Business Operations," *MIT Sloan Management Review* (60:4), pp. 37-44.
- van den Broek, E., Sergeeva, A., and Huysman, M. 2021. "When the Machine Meets the Expert: An Ethnography of Developing Ai for Hiring," *MIS Quarterly* (45:3), pp. 1557-1580.
- Weber, M., Engert, M., Schaffer, N., Weking, J., and Krcmar, H. 2022. "Organizational Capabilities for Ai Implementation—Coping with Inscrutability and Data Dependency in Ai," *Information Systems Frontiers*, pp. 1-21.
- Yin, R. K. 2018. *Case Study Research and Applications: Design and Methods*, (6th ed.). Los Angeles, CA: SAGE.
- Zhang, Z., Nandhakumar, J., Hummel, J., and Waardenburg, L. 2020. "Addressing the Key Challenges of Developing Machine Learning Ai Systems for Knowledge-Intensive Work," *MIS Quarterly Executive* (19:4), pp. 221-238.