
**Business Models Powered by Machine Learning:
Exploring How Machine Learning Changes the Ways
Organizations Create, Deliver, and Capture Value**



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Department of Law and Economics
at the Technical University of Darmstadt

Dissertation

by

Oliver Andreas Vetter

submitted in fulfillment of the requirements for the degree of
Doctor rerum politicarum (Dr. rer. pol.)

First assessor: Prof. Dr. Peter Buxmann
Second assessor: Prof. Dr. Alexander Benlian
Darmstadt 2024

Vetter, Oliver Andreas

Business Models Powered by Machine Learning: Exploring How Machine Learning Changes the Ways Organizations Create, Deliver, and Capture Value

Darmstadt, Germany, Technical University of Darmstadt

Year thesis published in TUprints: 2024

Date of the viva voce: November 4, 2024

Published under CC BY-SA 4.0 International

<https://creativecommons.org/licenses/>

Declaration of Authorship

I hereby declare that the submitted thesis is my own work. All quotes, whether word by word or in my own words, have been marked as such.

The thesis has not been published anywhere else nor presented to any other examination board.

Ich erkläre hiermit ehrenwörtlich, dass ich die vorliegende Arbeit selbstständig angefertigt habe. Sämtliche aus fremden Quellen direkt oder indirekt übernommenen Gedanken sind als solche kenntlich gemacht.

Die Arbeit wurde bisher weder einer anderen Prüfungsbehörde vorgelegt noch veröffentlicht.

Oliver Andreas Vetter

Darmstadt, June 3, 2024

Abstract

In today's world, artificial intelligence (AI) permeates almost all areas of human life. Modern AI supports us both in our leisure time (e.g., built into applications on our smartphones that recommend music we may like, recognize people in pictures, or act as digital assistants) and in our work (e.g., by automating tasks or creating analyses, predictions, or almost perfectly formulated texts). For organizations, AI, particularly instances of AI built with the data-based learning approach of machine learning (ML), also unlocks entirely new possibilities. Such ML systems can, for instance, be integrated into organizational processes to achieve efficiency gains by automating (parts of) tasks or to elevate decision-making quality by provisioning information. Furthermore, as the examples above illustrate, ML also enables the creation of novel kinds of products and services. The associated entrepreneurial opportunities unveiled by the latest technological advancements in the field of ML are correspondingly diverse and numerous, offering enormous potential for exploitation through suitable business models. A business model is an activity system that illustrates the logic of how an organization conducts business, i.e., how it creates value through the creation of its products and services, how it delivers this value to its customers, and how it ultimately captures value for itself, e.g., in the form of profits.

However, at the same time, the challenges that accompany the integration of ML into the business models of organizations exhibit similar diversity and also differ from those posed by other digital technologies. Therefore, the existing literature underscores that organizations wishing to harness the power of ML to drive their business models must carefully consider the peculiarities of the technology to be able to benefit in the long term. Yet, the state of existing research on the actual implementation of the various facets of ML-driven business models is sparse and lacks insights into their alignment with the particularities of ML. To expand this understanding in both academia and practice, this dissertation incorporates five papers that successively investigate ML-driven business models along the three business model dimensions of value creation, value delivery, and value capture. It examines both the ML-induced challenges that arise in each of these dimensions and the opportunities unlocked by ML, elaborating on their influences on the business logic of organizations from the perspective of the respective dimension of the business model.

First, two studies address the dimension of ML-driven value creation. The creation of ML systems requires experts from various disciplines to collectively reflect on the organization's existing knowledge (e.g., when making sense of data), which can lead to the creation of additional knowledge (e.g., through insights into inefficiencies in routines). Moreover, their data-based learning enables ML systems to generate knowledge in a way that complements the strengths of humans and thus to uniquely contribute to knowledge creation and revision in organizations. Existing literature on organizational learning hence regards productive ML systems as a new type of learner alongside humans. Yet, the potential for learning during ML development efforts, which include interactions of interdisciplinary groups of experts and (prototypical) ML systems, has to date remained largely unexplored. The first associated study therefore illuminates the beneficial learning processes that the creation of ML systems can stimulate. It also highlights the resulting

human knowledge as a valuable additional by-product that can contribute to the knowledge base of the organization and thus to its long-term success. The second study examines a downside of the data-based learning approach of ML: the need for extensive experimentation during ML development. This runs counter to the demand of conventional business processes for efficiency and exploiting existing strengths, and organizations must allocate their limited resources between the two approaches, creating tensions during ML development that can take various forms in different structural approaches. Building on the theoretical foundation of organizational ambidexterity, the study identifies these tensions and corresponding tactics that organizations can employ to alleviate the tensions, depending on their manifestation, and facilitate ML-driven value creation.

Next, the dissertation discusses the second dimension, ML-driven value delivery. A particular problem in this context is that ML systems are often highly complex, making them and their outputs incomprehensible to humans. If they cannot use them due to a lack of understanding, customers of ML-driven business models may thus fail to benefit from the value the business model intends to deliver (i.e., the ML system or its outputs). Therefore, the literature on explainable AI contains approaches that can provide users of ML systems with explanations that disclose their inner workings and the reasoning behind their decisions. Yet, thus far, these approaches have lacked a focus on distinct user groups and their specific requirements of the system. Especially lay users have often been neglected in previous studies. However, fostering the lay users' understanding is critical if they are to incorporate the output of ML systems in their decision-making to benefit from the products and services of ML-driven business models. Hence, the third study in this dissertation follows a design science research process and presents an approach to elaborate the requirements specific to the users of ML systems. On this basis, the study further derives design principles for designing ML systems that provide user-centric explanations and thereby enhance value delivery.

Finally, two more studies shed light on the third dimension of ML-driven value capture. In pursuit of their own goals, organizations must align all components of their business model to enable the capture of value, e.g., the reaping of profits from their business model in the long term. Only creating valuable solutions and supplying them to customers does not guarantee value capture for the organization, as the decade-long search of Twitter (now X) for a suitable way to profit from its unique offering and massive user base illustrates. With the current literature yielding little clarity on the nature of ML-driven business models, the fourth study in this dissertation aims to create a fundamental understanding of the business model components that organizations must align for successful value capture. Specifically, the resulting taxonomy offers insights into the components of ML-driven business models and is supplemented by archetypes that represent structural compositions of ML-driven business models commonly found in practice. Building on these findings, the fifth study investigates the question of how organizations seeking to profit from ML-driven business models can successfully realize them, which is under-researched in today's scientific literature as well. Realizing business models is an inherently dynamic and iterative process. In the case of ML-driven business models, the particularities of ML systems further complicate the effort, due (for instance) to the additional uncertainty stemming from the experimental character of ML development. Therefore, the study shows that organizations must build dynamic capabilities to be able to successfully realize ML-driven business models in the long term. Moreover, the study develops microfoundations (e.g., practices or processes) that empower the creation of the necessary dynamic capabilities, consequently contributing to the understanding of how organizations can successfully capture value sustainably from their ML-driven business models.

The studies within this dissertation illustrate that organizations must consider the unique characteristics of ML when designing and implementing their ML-driven business models to achieve sustainable success. Specifically, they show that the effects of ML particularities, such as the need for extensive experimentation in ML development, can manifest themselves in all three dimensions of the business model and can be both inhibiting (e.g., through additional uncertainty in the realization of the business model) and value-adding (e.g., through stimulated learning processes). The studies further delineate how organizations can take these influences into account through appropriate responses. This dissertation thus represents an important step toward a holistic understanding of ML-driven business models, emphasizes the value of the business model perspective for investigating the influence of ML on the business logic of organizations, and yields contributions to strategic management, entrepreneurship, and information systems literature. Thereby, it provides fertile ground for future examinations of ML-driven value creation, value delivery, and value capture against the backdrop of the high-level technological and entrepreneurial dynamism in the field of ML.

Abstract (German version)

Nahezu alle Bereiche des menschlichen Lebens sind inzwischen von künstlicher Intelligenz (KI) durchdrungen. Moderne KI unterstützt uns sowohl in unserer Freizeit (z. B. in Form von Anwendungen auf unseren Smartphones, die uns Musik empfehlen, die uns gefallen könnte, die Personen auf Bildern erkennen, oder als digitale Assistenten fungieren), als auch bei unserer Arbeit (z. B. durch die Automatisierung von Aufgaben oder die Erstellung von Analysen, Vorhersagen, oder nahezu perfekt formulierten Texten). Für Organisationen eröffnet KI, insbesondere KI, die durch den datenbasierten Lernansatz des maschinellen Lernens (ML) geschaffen wird, ebenfalls gänzlich neue Möglichkeiten. So können solche ML-Systeme in Organisationsprozesse integriert werden, um Effizienzgewinne durch die autonome Erledigung von (Teilschritten von) Aufgaben zu realisieren oder die Qualität von Entscheidungen durch die Bereitstellung von Informationen zu verbessern. Darüber hinaus ermöglicht das ML jedoch auch die Erstellung völlig neuartiger Produkte und Dienstleistungen, wie die obigen Beispiele zeigen. Die unternehmerischen Chancen, die die jüngsten technologischen Entwicklungen im Bereich des MLs ermöglichen, sind dementsprechend vielfältig und zahlreich, und bieten ein enormes Potenzial für die Erschließung durch geeignete Geschäftsmodelle. Ein Geschäftsmodell wird dabei als ein System von Aktivitäten verstanden, welches die Logik veranschaulicht, wie eine Organisation ihre Geschäftstätigkeit ausübt, d. h. wie sie durch die Herstellung ihrer Produkte und Dienstleistungen Wert schöpft, wie sie diesen Wert an ihre Kunden liefert und wie sie letztlich Wert für sich selbst erfasst, z. B. in Form von Gewinnen.

Gleichzeitig weisen aber die Herausforderungen, die mit der Integration von ML in die Geschäftsmodelle von Organisationen einhergehen, eine ähnliche Vielfalt auf und unterscheiden sich zudem von denen, die andere digitale Technologien mit sich bringen. Die existierende Literatur unterstreicht daher, dass Organisationen, die die Kraft des MLs als Antrieb für ihre Geschäftsmodelle nutzen möchten, die Eigenheiten der Technologie sorgsam berücksichtigen müssen, um langfristig profitieren zu können. Der Stand bestehender Forschung zur tatsächlichen Umsetzung der verschiedenen Facetten ML-getriebener Geschäftsmodelle ist jedoch spärlich und mangelt an Erkenntnissen zu deren Ausrichtung auf die Besonderheiten von ML. Um dieses Verständnis sowohl in der Wissenschaft als auch in der Praxis auszubauen, werden in dieser Dissertation in fünf Beiträgen ML-getriebene Geschäftsmodelle sukzessive entlang der drei Geschäftsmodell-Dimensionen Wertschöpfung, Wertlieferung, und Werterfassung untersucht. Dabei werden sowohl die ML-induzierten Herausforderungen, die in jeder dieser drei Dimensionen auftreten, als auch die von ML eröffneten Möglichkeiten betrachtet und deren Einfluss auf die Geschäftslogik von Organisationen aus der Perspektive der jeweiligen Geschäftsmodell-Dimension herausgearbeitet.

Die vorliegende Dissertation befasst sich in zwei Studien zunächst mit der Dimension der ML-getriebenen Wertschöpfung. Die Erstellung von ML-Systemen erfordert, dass Experten aus verschiedenen Disziplinen gemeinsam über vorhandenes Wissen der Organisation reflektieren (z. B. bei der Auswertung von Daten), was zur Erzeugung von zusätzlichem Wissen führen kann

(z. B. durch Erkenntnisse über Ineffizienzen in Routinen). Darüber hinaus können ML-Systeme durch ihr datenbasiertes Lernen auch Wissen auf eine Art und Weise generieren, die komplementär zu den Stärken von Menschen ist, und so zur Schaffung und Revision von Wissen in Organisationen beitragen. Während in der bestehenden Literatur über organisationales Lernen produktive ML-Systeme deshalb als neue Art Lerner neben dem Menschen angesehen werden, ist potenzielles Lernen während der ML-Entwicklung, welche Interaktionen von interdisziplinären Expertengruppen und (prototypischen) ML-Systemen umfasst, bislang weitgehend unerforscht geblieben. Die erste zugehörige Studie betrachtet daher die vorteilhaften Lernprozesse, die während der Erstellung von ML-Systemen angeregt werden können. Sie hebt weiterhin das daraus entstehende menschliche Wissen als wertvolles, zusätzliches Nebenprodukt hervor, welches zur Wissensbasis der Organisation und damit zu ihrem langfristigen Erfolg beitragen kann. Die zweite Studie beschäftigt sich mit einem Nachteil des datenbasierten Lernansatzes von ML: dem Bedarf nach umfangreichem Experimentieren während der ML-Entwicklung. Da dieser dem Bedarf konventioneller Geschäftsprozesse nach Effizienz und Ausnutzung vorhandener Stärken entgegensteht und Organisationen ihre limitierten Ressourcen zwischen den beiden Ansätzen aufteilen müssen, erzeugt er Spannungen während der ML-Entwicklung, die sich in verschiedenen strukturellen Ansätzen unterschiedlich ausbilden können. Die auf der theoretischen Grundlage der organisationalen Ambidextrie aufbauende Studie identifiziert diese Spannungen und führt zugehörige Taktiken an, die Organisationen anwenden können, um die Spannungen je nach Ausprägung zu mildern und die ML-getriebene Wertschöpfung zu erleichtern.

Im Anschluss wird die Dimension der ML-getriebenen Wertlieferung diskutiert. Besonders problematisch ist in diesem Kontext, dass ML-Systeme häufig hochkomplex und somit die Systeme und ihre Ausgaben für Menschen unverständlich sind. Kunden von ML-getriebenen Geschäftsmodellen können deshalb Schwierigkeiten haben, einen Vorteil aus dem Wert (d. h. dem ML-System oder seinen Ausgaben) zu ziehen, den das Geschäftsmodell zu liefern beabsichtigt, wenn sie diesen aus mangelndem Verständnis nicht verwenden können. In der Literatur finden sich daher Ansätze der erklärbaren KI, die verwendet werden können, um Nutzende von ML-Systemen mit Erklärungen auszustatten, die deren Funktionsweise und Begründungen für ihre Entscheidungen offenlegen. Allerdings mangelt es diesen Ansätzen bislang an Fokus auf verschiedene Gruppen an Nutzenden und deren individuelle Anforderungen an das System. Insbesondere Laien wurden in bisherigen Studien häufig vernachlässigt. Die Förderung des Verständnisses von Laien ist jedoch von entscheidender Bedeutung, wenn diese die Ausgabe von ML-Systemen in ihre Entscheidungsfindung einbeziehen und so von den Produkten und Dienstleistungen ML-getriebener Geschäftsmodelle profitieren sollen. Die dritte Studie in dieser Dissertation folgt daher einem designwissenschaftlichen Forschungsprozess und stellt einen Ansatz vor, die Anforderungen der Nutzenden von ML-Systemen gezielt herauszuarbeiten. Darauf basierend werden zudem Designprinzipien für die Gestaltung von ML-Systemen abgeleitet, die den Nutzenden maßgeschneiderte Erklärungen bereitstellen und somit die Wertlieferung fördern.

Schließlich beleuchten zwei weitere Studien die Dimension der ML-getriebenen Werterfassung. Im Streben nach ihren eigenen Zielen müssen Organisationen alle Komponenten ihres Geschäftsmodells so ausrichten, dass sie die Erfassung von Wert, wie z. B. das langfristige Erzielen von Gewinnen aus ihrem Geschäftsmodell, ermöglichen. Die Entwicklung wertvoller Lösungen und deren Bereitstellung für Kunden garantiert keine Werterfassung für die Organisation selbst, wie die jahrzehntelange Suche von Twitter (jetzt X) nach einem geeigneten Weg, Gewinne aus seinem einzigartigen Angebot und seiner riesigen Basis an Nutzenden zu erzielen, verdeutlicht.

Da in der derzeitigen Literatur jedoch noch wenig Klarheit über die Beschaffenheit von ML-getriebenen Geschäftsmodellen herrscht, soll in der vierten Studie dieser Dissertation ein grundlegendes Verständnis für die Geschäftsmodellkomponenten geschaffen werden, welche Organisationen für die erfolgreiche Werterfassung ausrichten müssen. Konkret bietet die resultierende Taxonomie Einblicke in die Komponenten von ML-getriebenen Geschäftsmodellen und wird durch Archetypen ergänzt, die häufig in der Praxis anzutreffende strukturelle Kompositionen von ML-getriebenen Geschäftsmodellen darstellen. Darauf aufbauend befasst sich die fünfte Studie mit der ebenfalls in der aktuellen wissenschaftlichen Literatur noch unzureichend erforschten Frage, wie Organisationen, die von ML-getriebenen Geschäftsmodellen profitieren möchten, diese erfolgreich realisieren können. Die Realisierung von Geschäftsmodellen ist ein inhärent dynamischer und iterativer Prozess. Im Fall von ML-getriebenen Geschäftsmodellen wird dieser durch die Besonderheiten von ML-Systemen weiter erschwert, beispielsweise durch zusätzliche Ungewissheit, die sich aus dem experimentellen Charakter der ML-Entwicklung ergibt. Die Untersuchung zeigt, dass Organisationen daher dynamische Fähigkeiten aufbauen müssen, um ML-getriebene Geschäftsmodelle dennoch langfristig erfolgreich realisieren zu können. Darüber hinaus werden Mikrofundamente (z. B. Praktiken oder Prozesse) erarbeitet, welche den Aufbau der notwendigen dynamischen Fähigkeiten ermöglichen, und folglich zum Verständnis beigetragen, wie Organisationen nachhaltig erfolgreich Wert aus ihren ML-getriebenen Geschäftsmodellen erfassen können.

Die in dieser Dissertation enthaltenen Studien verdeutlichen, dass Organisationen bei der Konzeption und Umsetzung ihrer ML-getriebenen Geschäftsmodelle für deren nachhaltigen Erfolg die einzigartigen Besonderheiten von ML berücksichtigen müssen. Sie zeigen dabei insbesondere auf, dass sich die Auswirkungen von ML-Besonderheiten, wie z. B. die Notwendigkeit umfangreicher Experimente in der ML-Entwicklung, in allen drei Dimensionen des Geschäftsmodells manifestieren, und dabei sowohl hemmend (z. B. durch zusätzliche Ungewissheit bei Realisierung des Geschäftsmodells), als auch wertstiftend auftreten können (z. B. durch angestoßene Lernprozesse). Die Studien erläutern außerdem, wie Organisationen diesen Einflüssen durch entsprechende Maßnahmen Rechnung tragen können. Diese Dissertation stellt somit einen wichtigen Schritt in Richtung eines ganzheitlichen Verständnisses von ML-getriebenen Geschäftsmodellen dar, unterstreicht den Wert der Geschäftsmodellperspektive für Untersuchungen von ML-Einflüssen auf die Geschäftslogik von Organisationen, und liefert Beiträge zur Literatur in den Bereichen Strategic Management, Entrepreneurship, und Information Systems. Sie bietet damit einen fruchtbaren Boden für zukünftige Untersuchungen der ML-getriebenen Wertschöpfung, Wertlieferung, und Werterfassung vor dem Hintergrund der hohen technologischen und unternehmerischen Dynamik im Feld des MLs.

Table of Contents

Declaration of Authorship	III
Abstract	IV
Abstract (German version)	VII
Table of Contents	X
List of Figures	XII
List of Tables	XIII
List of Abbreviations	XIV
1 Introduction	1
1.1 Overarching Motivation.....	1
1.2 Research Questions	2
1.3 Structure of this Dissertation	4
2 Research Context and Positioning of This Dissertation	8
2.1 Artificial Intelligence and Machine Learning	8
2.2 Business Models.....	10
2.3 The Need to Revisit Value Creation, Delivery, and Capture with Machine Learning	13
3 Research Paper A: Potential for Learning During Value Creation With Machine Learning	16
Abstract.....	16
3.1 Introduction.....	17
3.2 Theoretical Background	18
3.3 Qualitative Research Methodology	22
3.4 Results.....	24
3.5 Discussion.....	33
Acknowledgment	36
4 Research Paper B: Creating Value Despite Machine-Learning-Induced Tensions	37
Abstract.....	37
4.1 Introduction.....	38
4.2 Theoretical Background	39
4.3 Methodology.....	42
4.4 Results.....	46

4.5 Discussion.....	51
4.6 Conclusion.....	53
Acknowledgment.....	53
5 Research Paper C: Delivering Value Through Explainable Artificial Intelligence.....	54
Abstract.....	54
5.1 Introduction.....	55
5.2 Theoretical Background	56
5.3 Methodology.....	59
5.4 Results.....	62
5.5 Discussion.....	70
5.6 Conclusion.....	72
Acknowledgment.....	72
6 Research Paper D: Composition of Machine-Learning-Driven Business Models Enabling Value Capture	73
Abstract.....	73
6.1 Introduction.....	74
6.2 Theoretical Background	75
6.3 Methodology.....	77
6.4 Results.....	79
6.5 Conclusion, Limitations & Outlook.....	87
Acknowledgment.....	88
7 Research Paper E: Capturing Value by Realizing Machine-Learning-driven Business Models.....	89
Abstract.....	89
7.1 Introduction.....	90
7.2 Theoretical Background	91
7.3 Methodology.....	94
7.4 Results.....	96
7.5 Discussion.....	102
7.6 Conclusion.....	104
Acknowledgment.....	104
8 Overarching Discussion and Conclusion	105
8.1 Theoretical Contributions.....	105
8.2 Practical Contributions	108
8.3 Concluding Remarks	110
References.....	111
Appendix.....	134

List of Figures

Figure 1. The Business Model Canvas as an Exemplary Conceptualization of Business Models and Their Components, With Added Categorization Into Value Creation, Delivery, and Capture (Adapted From Osterwalder and Pigneur, 2010, Strategyzer.com, CC BY-SA 3.0)..... 11

Figure 2. Ambidexterity in ML Development and Operation..... 47

Figure 3. Mapping of XE to XR and DP..... 64

Figure 4. Shown Mock-ups for XE 6..... 66

Figure 5. Shown Mock-ups for XE 7..... 66

Figure 6. Shown Mock-ups for XE 8..... 67

Figure 7. Microfoundations of Dynamic Capabilities for Realizing ML-Driven Business Models Emergent From the Interview Data..... 98

List of Tables

Table 1. List of Publications Included in This Dissertation.	5
Table 2. Overview of Research Papers.	6
Table 3. Core Knowledge Types and Core Learning and Problem-Solving Processes.	19
Table 4. Overview of the Interviewed Experts.	23
Table 5. Identified Learning Sub-Processes in ML Developments.	25
Table 6. Study Participants Overview.	44
Table 7. Overview of Interviewees for Mock-up Evaluation.	61
Table 8. Evaluation Results from the Interviews for XE 7.	69
Table 9. Classification of Existing Taxonomies Adapted From Dehnert et al. (2021).	77
Table 10. Final Taxonomy, Visualized as a Morphological Box.	80
Table 11. Distribution of the Start-Ups' Characteristics Within Each Archetype.	85
Table 12. List of Interviewed Experts.	95
Appendix 1. Excerpt of the Examined Sample Set of B2B ML-Driven Start-Ups (Paper D).	134

List of Abbreviations

AGI	Artificial general intelligence
AI	Artificial intelligence
AM	Actionability measure
API	Application programming interface
BM	Business model
BMBF	German Federal Ministry of Education and Research
BMI	Business model innovation
B2B	Business-to-business
B2C	Business-to-consumer
CPM	Corporate performance management
DP	Design principle
GQ	General quality
IS	Information systems
IT	Information technology
KM	Knowledge measure
MAE	Mean absolute error
ML	Machine learning
PTKA	Project Management Agency Karlsruhe
RQ	Research question
UDP	User experience / user interface design principle
UI	User interface
UX	User experience
VHB	German Academic Association for Business Research
XAI	Explainable artificial intelligence
XE	Explainability element
XR	Explainability requirement
XQ	Explainable-artificial-intelligence-specific quality

1 Introduction

The creation of modern marvels such as self-driving vehicles, smartphones recognizing their owners' faces, radiology tools that detect brain tumors in medical imaging, and digital assistants that can converse with humans in natural language have recently been made possible through the use of contemporary artificial intelligence (AI; e.g., Berente et al., 2021; Lebovitz et al., 2021; Stone et al., 2016). Machine learning (ML) in particular has enabled many of the radical innovations that increasingly find their way into both our work and personal lives (e.g., Benbya et al., 2020b; Berente et al., 2021; Lyytinen et al., 2021; Yoo, 2010). ML denotes a set of techniques for creating instances of AI by letting algorithms autonomously abstract patterns from relationships in data, and then use these patterns to make predictions for new data (e.g., Brynjolfsson & Mitchell, 2017; Mitchell, 1997; Russell & Norvig, 2021). We interact with systems created with ML when we use navigation software that helps us circumvent a currently forming traffic jam, when we analyze ML-generated weather forecasts or sales predictions for the upcoming month, or when we ask ChatGPT to create a few images of conversing robots for a work presentation. With these examples only scratching the surface of the ways in which ML systems can assist in value-adding activities or even automate processes (Jordan & Mitchell, 2015), organizations developing and exploiting the unique capabilities of ML systems may thus stand to profit immensely.

1.1 Overarching Motivation

ML systems not only possess unique capabilities for use in a variety of scenarios to improve an organization's performance, e.g., by automating (parts of) processes or helping human experts make better decisions, but also unlock a plethora of entrepreneurial opportunities to offer novel products and services (e.g., Benbya et al., 2021; Chalmers et al., 2021; Davenport et al., 2020; Obschonka & Audretsch, 2020). In this regard, ML holds enormous potential for organizations to capitalize on, by either integrating ML into their existing business models to ideally exploit the strengths of the technology, or developing novel, unprecedented business models with ML at their core (e.g., Chalmers et al., 2021; Davenport et al., 2020; Townsend & Hunt, 2019). Some organizations do this with great success, whether by deeply embedding the technology into their internal processes, such as Uber in its software for managing drivers autonomously (Uber, 2024), by ubiquitously using the technology in almost all parts of their organization, such as Google (Burr, 2023), or by offering ML-based products and services, such as OpenAI with ChatGPT (OpenAI, 2022) or the German ML-based translation start-up DeepL, which, through its latest funding round, has recently reached a company valuation of over €1 billion and thus start-up unicorn status (DeepL, 2023). However, not all organizations see their ML endeavors succeed. On the contrary, many organizations fail to develop and deploy suitable ML systems or achieve business gains from their efforts (Gartner, 2024; Ransbotham et al., 2019). Yet, literature investigating the underlying reasons for stark contrasts in the success of organizations trying to realize business models employing ML systems remains scarce (e.g., Lange et al., 2021; Weber et al., 2022). To address this lack of research, this dissertation examines both the factors that make

ML endeavors unique as well as the business logic explicating the success of ML-infused organizations, providing guidance for scholars and practitioners alike.

To this end, I argue that the business model view poses a suitable perspective from which to dissect the numerous influences that ML systems exert on organizations seeking to integrate them into their core business, as well as the organizations' responses to these influences. In short, the business model concept describes the business logic of an organization, ranging from how the organization creates value through its products and services, through how it delivers value to its clients, to how it captures value for itself, e.g., in the form of profits, to pursue its goals and ensure its long-term survival (e.g., Osterwalder & Pigneur, 2010; Teece, 2010). Business model research is interdisciplinary in nature, spanning the strategic management, entrepreneurship, and information systems (IS) literatures (e.g., Steininger, 2019). In IS research in particular, the business model concept is seen as a "missing link" between strategy, processes, and information technology (IT; Veit et al., 2014). As such, it is often used to understand IT-induced changes in how organizations create, deliver, and capture value (e.g., Hartmann et al., 2016; Remane et al., 2016). Regarding the topic at hand, when ML systems play a vital role in an organization's business model, this dissertation conceptualizes it as an ML-driven business model (see Sections 2.3, 6.2.1, and 7.2.1). Chapter 2 further elaborates on the business model perspective and its suitability for investigating the influence of ML on the business logic of organizations.

The holistic approach of the business model perspective in regarding business value is particularly advantageous, as ML systems come with challenges that can permeate almost all parts of organizations implementing ML-driven business models and which differ from those of other digital technologies (e.g., Benbya et al., 2021; Berente et al., 2021). For instance, while the data-based learning approach of ML systems empowers them with the capabilities to realize the aforementioned use cases and many more, it poses novel requirements for their successful development that organizations must meet (e.g., Amershi et al., 2019; Choudhury et al., 2021). Such unique challenges for ML-driven business models emerge in a variety of facets both within and external to the organization and pertain to such aspects as ethics, trust, security, labor, and business strategy (e.g., Faraj et al., 2018; Kellogg et al., 2020; Martin, 2019a; Stone et al., 2016). For example, as the patterns saved within ML systems are often incomprehensible to humans, relying on the system's output could lead to organizations unwittingly targeting their advertising on the basis of a consumer's medical condition, or excluding a candidate from employment decisions due to marital status (e.g., Asatiani et al., 2021; Martin, 2019a). Moreover, such issues can be exacerbated when the degree of agency of ML systems is increased, for instance, when shifting the locus of managerial control to ML, as in the case of Uber letting their ML-powered app manage and direct its enormous fleet of human drivers (e.g., Kellogg et al., 2020; Möhlmann et al., 2021; Wiener et al., 2023). The next section delineates how the combination of ML-induced challenges and novel opportunities can affect different aspects of ML-driven business models, motivating the studies within this dissertation. Chapter 2 further expounds on the particularities of ML systems in the organizational context, and particularly on the need to revisit the business logic underlying ML-driven value creation, delivery, and capture.

1.2 Research Questions

For the purposes of this dissertation, the various investigated beneficial and obstructive effects of ML systems on the business logic of ML-infused organizations are categorized into three dimensions, namely, value creation, value delivery, and value capture. Notably, these three dimensions of the business model (see Section 2.2) are not neatly distinguishable from each other.

On the contrary, in a well-functioning business model, all three dimensions are inherently and inseparably interlinked, forming a coherent storyline explaining how the organization creates value, delivers it to customers, and ultimately captures value from its efforts (e.g., Al-Debei & Avison, 2010; Burström et al., 2021; Massa et al., 2017; Zott et al., 2011). These interdependencies become especially transparent in the examination of value capture, as an organization's possibilities for generating revenue and profits largely depend on how it has configured the other components of its business model (e.g., Teece, 2010; Tidhar & Eisenhardt, 2020). Therefore, the three dimensions are not to be viewed as strictly separable, but instead as three interconnected focal sides from which to understand the logic underneath ML-driven value creation, value delivery, and value capture, respectively.

First, regarding value creation, while responsible for many of ML's strengths, the data-based learning approach of ML systems also complicates their creation. In particular, building ML systems requires not only ML-specific resources and expertise, but also new working practices that are specifically adapted to accommodate data-based learning, for instance, through fostering fast and iterative experimentation (e.g., Amershi et al., 2019; Baier et al., 2019; Choudhury et al., 2021). Organizations must therefore take these difficulties into account when seeking to create value by developing ML systems. Yet, with many organizations currently ill-equipped to meet these demands, they may decide to reduce their investments into developing ML systems, for fear of wasting their valuable resources (e.g., Pumplun et al., 2019; Ransbotham et al., 2020; Sculley et al., 2015). This may prove to be a fatal fallacy, however, as recent research indicates that the learning performed by ML systems may also act as source of knowledge, stimulating learning processes in the organization (e.g., Argote et al., 2021; Balasubramanian et al., 2022; Sturm et al., 2021b). ML development further requires experts from various disciplines to engage with their existing knowledge (e.g., Studer et al., 2021), e.g., when evaluating the collected data, who may thereby unearth new knowledge, e.g., by finding inefficiencies in their routines. Therefore, organizations evaluating their ML development efforts may be well advised to consider learning processes generating knowledge as a valuable by-product of ML-driven value creation that can contribute to their organizational learning and thus their long-term performance (e.g., Argote et al., 2021; March, 1991). To scrutinize these ML-induced influences, I ask:

Research question 1 (RQ1): *How do organizations create value, despite ML-induced difficulties, while exploiting the capabilities of ML in ML-driven business models?*

Second, concerning value delivery, a major obstacle to the successful provision of value generated with ML to the customer is the aforementioned difficulty of humans to understand the ML system and their outcomes. This incomprehensibility stems largely from the complexity of modern ML systems and can diminish or entirely negate the benefit they can convey to customers who do not want to use systems they do not understand (e.g., Asatiani et al., 2021; Barredo Arrieta et al., 2020; Faraj et al., 2018). Coupled with the fact that ML systems can sometimes create unexpected results that may lead to undesired outcomes (e.g., Benbya et al., 2020b; Benbya et al., 2021), organizations must find ways to ensure that their clients can leverage the potential value of their ML systems. To this end, organizations can utilize explainable AI (XAI) approaches to foster their ML system's ability to disclose its inner workings and the reasoning leading to its decisions (e.g., Ågerfalk, 2020; Rai, 2020; Rudin, 2019). To examine how organizations can incorporate XAI to ensure that their clients can benefit from the ML systems they provide, I ask:

Research question 2 (RQ2): *How do organizations deliver value by providing explanations for inherently incomprehensible ML systems and their output in ML-driven business models?*

Third, with regard to value capture, organizations seeking to create and maintain competitive advantage and ultimately profit from ML systems must align all components of their ML-driven business model, to allow for their intended avenue for capturing value (e.g., Burström et al., 2021; Teece, 2010; Zott et al., 2011). However, the design of these components and their alignment proves difficult, with limited research available to offer guidance (Burström et al., 2021; Weber et al., 2022). As ML comes with capabilities and challenges unique to the technology (e.g., Benbya et al., 2021; Berente et al., 2021), the composition of ML-driven business models may consequently differ from other IT-related business models as well, thus meriting re-examination. Furthermore, the realization of successful business models has always been a difficult, multi-faceted, and iterative process (e.g., Schoemaker et al., 2018; Teece, 2018), and the particularities of ML may further complicate such endeavors, for instance, by adding another experimental component inherent in ML development (e.g., Choudhury et al., 2021; Studer et al., 2021). To investigate how organizations can nevertheless profit from ML-driven business models, I ask:

Research question 3 (RQ3): *How do organizations capture value by realizing ML-driven business models, and what constitutes these business models?*

1.3 Structure of this Dissertation

To answer the RQs introduced above, this dissertation encompasses five research papers that peer-reviewed conference proceedings have published. This section summarizes the papers' research approaches and contributions and illustrates the overall structure of the dissertation.

Table 1 displays all five papers, which are referred to as papers A, B, C, D, and E throughout the dissertation. Paper A contributes to answering RQ1 by scrutinizing the learning processes that can occur when organizations develop ML systems. It thereby highlights the additional value that the resulting knowledge gains of the involved domain experts, acquired through learning processes stimulated during ML development, can bring to the organization. Paper B complements the contribution to answering RQ1 by illuminating how the requirements of ML development for extensive experimentation uniquely complicate ML-driven value creation. It further provides suitable tactics that enable organizations to satisfy the demands of ML for experimentation under the consideration of strategic choices, particularly regarding the organization's structural set-up for ML development. Paper C proceeds to the topic of ML-driven value delivery (RQ2) with an examination of how organizations ensure that their customers can benefit from the potential value they aim to provide through their ML systems. Specifically, the paper studies how organizations can design their ML systems for this purpose by equipping them with explanations for the system's behavior and its output that are tailored specifically to the users of the system. Next, paper D contributes to RQ3 by first dissecting the structural composition of ML-driven business models through which organizations can seek to profit in the long term and then highlighting the ML-driven business models commonly employed in practice. To further explicate ML-driven value capture, paper E elaborates on the particularities of ML systems hindering the realization of business models and explores the capabilities that empower organizations to nevertheless successfully realize and ultimately profit from their ML-driven business model.

Table 1. List of Publications Included in This Dissertation.

RQ1: Value Creation	Paper A	Vetter, Oliver A., Sturm, Timo, Fecho, Mariska, & Buxmann, Peter (2023). Machine learning developments as stimuli for organizational learning. <i>Proceedings of the 44th International Conference on Information Systems (ICIS)</i> . VHB Rating ¹ : A
	Paper B	Vetter, Oliver A., Pumplun, Luisa, & Koppe, Timo (2023). An ambidextrous perspective on machine learning development and operation: The nexus of tensions, organizational structure, and tactics. <i>Proceedings of the 27th Pacific Asia Conference on Information Systems (PACIS)</i> . VHB Rating: C
RQ2: Value Delivery	Paper C	Vetter, Oliver A., & Efremov, Alexander (2023). A user-centric approach to explainable AI in corporate performance management. <i>Proceedings of the 27th Pacific Asia Conference on Information Systems (PACIS)</i> . VHB Rating: C; Honored with Best Complete Paper Runner-Up Award
RQ3: Value Capture	Paper D	Vetter, Oliver A., Hoffmann, Felix, Pumplun, Luisa, & Buxmann, Peter (2022). What constitutes a machine-learning-driven business model? A taxonomy of B2B start-ups with machine learning at their core. <i>Proceedings of the 30th European Conference on Information Systems (ECIS)</i> . VHB Rating: A
	Paper E	Vetter, Oliver A., Mehler, Maren F., & Buxmann, Peter (2023). As much art as science – examining the realization of business models driven by machine learning through a dynamic capabilities perspective. <i>Proceedings of the 31st European Conference on Information Systems (ECIS)</i> . VHB Rating: A

These five studies are grounded on various theoretical foundations and employ several methodological approaches, for which Table 2 presents an overview. As little evidence exists on the topic of organizations realizing ML-driven business models to date, this dissertation relies heavily on explorative research approaches, which are particularly suitable for uncovering phenomena in areas of research largely uncharted thus far (e.g., Corbin & Strauss, 2015; Myers & Newman, 2007). More specifically, papers A, B, and E, are grounded in data from a plethora of interviews with industry experts knowledgeable on the respective focal subjects. In particular, semi-structured interview guidelines were employed for data collection, as they provide both guidance for the interviews and the freedom to improvise with spontaneous relevant questions, depending on the course the interview takes (Myers & Newman, 2007). In papers A and E, the

¹ The VHB Rating, published by the German Academic Association for Business Research (VHB), served as the preferred means of assessing the quality of research papers in my doctoral studies. Specifically, this dissertation refers to the VHB Publication Media Rating 2024, which constitutes the latest version of the rating at the time of writing.

collected data forms the basis for sense-making through qualitative content analysis (Hsieh & Shannon, 2005; Saldaña, 2015). For paper B, an ML-based analysis of the interview data complements a manual analysis, drawing methodological inspiration from computational grounded theory (e.g., Nelson, 2020). Paper C presents a design science study, well aligned with the goal of the investigation, to generate guidance on how organizations can design ML systems that address the need for explainability in practice (e.g., Gregor & Hevner, 2013). Finally, Paper D rounds out the insights on ML-driven business models with a taxonomy development approach, building on existing business model literature before integrating empirical data on ML-driven business models found in practice to specify and extend the taxonomy (Nickerson et al., 2013). Regarding their theoretical foundation, all papers discuss topics connected to ML-driven business models and are thus linked through the business model view, with paper E drawing on it and paper D adopting it as theoretical basis. Therefore, Section 2.2 describes the business model view in detail. Moreover, with business models encompassing a multitude of facets, and each paper illuminating a different one, they build on different theoretical foundations relevant to the research questions they are investigating. In particular, paper A is based on organizational learning, paper B on organizational ambidexterity, and paper E on dynamic capabilities. Section 2.3 addresses the relevance of each of these theoretical foundations to the respective overarching research question, while elaborations on the theoretical foundations themselves appear in the respective papers.

Table 2. Overview of Research Papers.

Research Paper	Research Approach	Theoretical Foundation
Paper A	Qualitative content analysis	Organizational learning
Paper B	ML-assisted content analysis	Organizational ambidexterity
Paper C	Design science research	-
Paper D	Taxonomy development	Business model view
Paper E	Qualitative content analysis	Dynamic capabilities

Beyond the publications constituting this cumulative dissertation (see Table 1), I co-authored the following peer-reviewed papers during my doctoral studies at the Technical University of Darmstadt:

- Fecho, Mariska, Wahl, Nihal, von Ahsen, Anette, & Vetter, Oliver A. (2024). Paving the way to a green future with artificial intelligence – exploring organizational adoption factors. *Proceedings of the 57th Hawaii International Conference on System Sciences (HICSS)*, 821–830. VHB Rating: B
- Gräf, Miriam, Zöll, Anne, Wahl, Nihal, Ellenrieder, Sara, Hager, Florentina, Sturm, Timo, & Vetter, Oliver A. (2023). Designing the organizational metaverse for effective socialization. *Proceedings of the 27th Pacific Asia Conference on Information Systems (PACIS)*. VHB Rating: C
- Mehler, Maren F., & Vetter, Oliver A. (2023). How much are machine assistants worth? Willingness to pay for machine learning-based software testing. *Proceedings of the 31st European Conference on Information Systems (ECIS)*. VHB Rating: A

All publications² included in this dissertation appear in Chapters 3 through 7. Furthermore, Chapter 2 defines important concepts and delineates the overarching theoretical background relevant to the dissertation as a whole. Finally, Chapter 8 contains the overarching contributions of this dissertation, avenues for future research, and concluding remarks.

² The papers have been slightly adapted from their originally published version to have a consistent layout throughout this dissertation. They are further written from the first-person plural (i.e., “we”) perspective, since co-authors contributed to each publication.

2 Research Context and Positioning of This Dissertation

This chapter establishes the concepts that are fundamental to the five research papers that make up this cumulative dissertation. In particular, it details the terms “artificial intelligence,” specifically “machine learning,” and “business models,” and then elaborates on the need to revisit the value creation, value delivery, and value capture of organizations infused by ML.

2.1 Artificial Intelligence and Machine Learning

Many consider the origin of scholarly attention on *artificial intelligence* to be the 1955 proposal for the Dartmouth Summer Research Project on Artificial Intelligence, by McCarthy et al. (1955). The seminal conference on the topic one year later sparked a plethora of studies, as well as technological innovations in various fields such as education, healthcare, and transportation (e.g., Stone et al., 2016). However, despite the enormous amount of scientific literature to date, no universally agreed-upon definition of the term “artificial intelligence” exists, largely due to the challenge of defining what intelligence is (e.g., Russell & Norvig, 2021). Therefore, various research fields have brought forth definitions illuminating AI through different lenses. For instance, they may focus on the ability of AI to perform cognitive tasks similarly to humans (e.g., Benbya et al., 2020a; Brynjolfsson & Mitchell, 2017), or the ability to make sense of its environment by learning from data (e.g., Jordan & Mitchell, 2015; Castelvechi, 2016). Russell and Norvig (2021) summarize AI as the study and creation of rational agents, i.e., agents aiming to attain the best expected outcome for a particular objective. Today, these agents may perform such tasks as perceiving, problem-solving, innovation, decision-making, or automating physical processes (e.g., Benbya et al., 2020a). AI’s performance in many of these tasks has been improving rapidly in recent years due to increasing availability of data and processing power, even overtaking human performance in some respects (e.g., Ågerfalk, 2020; Brynjolfsson & McAfee, 2017). For instance, machines have become both better and faster at recognizing objects in images, and have also beaten the best human players of Go or poker (e.g., Brynjolfsson & McAfee, 2017).

Despite these outstanding achievements, today’s AI remains bound to the specific problem area it was developed for. Research distinguishes between this form of *artificial narrow intelligence*, which encompasses all instances of AI currently on the market, on the one hand, and *artificial general intelligence* (AGI) and *artificial superintelligence* on the other (e.g., Benbya et al., 2021). AGI denotes machines matching humans in general intelligence, including common sense and the ability to learn, reason, and plan. Scholars believe that once AGI exists, it might lead to the development of artificial superintelligence that far exceeds human performance and continually learns and improves, thus raising far-reaching ethical and philosophical arguments (Bostrom, 2017). However, as the latter two categories remain theoretical, this dissertation focuses solely on the creation, delivery, and capture of value through narrow AI. The most popular technique for implementing instances of narrow AI is *machine learning* (e.g., Brynjolfsson & Mitchell, 2017; Jordan & Mitchell, 2015). While the early approaches to creating rule-based narrow AI, such as

expert systems, required developers to hand-code desired behaviors into the machine, ML approaches allow systems to learn autonomously from given examples in the form of data (e.g., Russell & Norvig, 2021). ML algorithms achieve this capability by extracting patterns from datasets in a process called *training*, then storing the patterns in models for use in solving future instances of the given problem (e.g., Mitchell, 1997; Russell & Norvig, 2021). In this way, an ML algorithm might, for example, derive a rationale for predicting loan defaults from a dataset of past lending decisions, and save it in a model, e.g., a neural network. ML approaches that utilize models with multiple layers of computing elements (e.g., convolutional neural networks) are referred to as *deep learning* and often demand large amounts of data and processing power (e.g., Russell & Norvig, 2021). This category also encompasses *generative AI*, such as ChatGPT, GitHub Copilot, or DALL-E, which can be prompted to create new text, images, audio, or other media, and consists of deep learning models with billions of parameters (e.g., Benbya et al., 2024; Raisch & Fomina, 2024). While the studies presented in this dissertation do not specifically consider generative AI, due to its recent emergence, I abstract their findings from specific types of ML to provide contributions across all forms of machine learning, which thus yield insights for generative AI as well. Moreover, as most contemporary instances of AI are based on ML (e.g., Brynjolfsson & Mitchell, 2017; Jordan & Mitchell, 2015), the term ML system denotes instances of AI implemented through the use of ML within this dissertation.

The literature distinguishes three major types of ML (e.g., Mitchell, 1997; Russell & Norvig, 2021): supervised learning, unsupervised learning, and reinforcement learning. In *supervised learning*, the ML algorithm is supplied with training data in the form of input-output pairs (i.e., labeled data) and derives a function that allows it to predict outputs from inputs. *Unsupervised learning* algorithms receive no information on the desired output (i.e., the labels), and derive patterns from the input without feedback. For example, while a supervised learning system might process a large labeled image dataset of pets and pastries to extract the patterns that allow it to recognize a puppy in an image, an unsupervised learning system could identify clusters of similar images within the same dataset without requiring the labels but would need a human to make sense of the clusters (e.g., Brynjolfsson & McAfee, 2017; Russell & Norvig, 2021). A *reinforcement learning* system gathers its own experience by interacting with a (real or simulated) environment and receiving rewards or punishments as outcomes of its actions. Then, aiming to maximize its rewards, it learns from this experience to inform future actions, for instance, learning which stock trades lead to gains, or which chess moves lead to victory (e.g., Dempster & Leemans, 2006; Russell & Norvig, 2021).

ML enables the creation of IT systems with distinctive capabilities that can be of enormous value to individuals or organizations (e.g., Benbya et al., 2021; Brynjolfsson & McAfee, 2017; Jordan & Mitchell, 2015). Examples include the aforementioned image recognition systems for use in self-driving cars, forecasting systems that predict the amount of electricity generated from renewable energies or future states of the stock market, as well as multi-purpose agents that can be interacted with in natural language like the modern marvels ChatGPT from OpenAI or Gemini from Google. Conversely, ML systems also exhibit some characteristics that differ from those of traditional, non-ML IT systems, which raise unique challenges for ML projects in organizations, affecting both the development of the systems and their implementation in the organization (e.g., Ågerfalk, 2020; Berente et al., 2021; Choudhury et al., 2021). As training ML systems involves them autonomously learning from the data they receive, the success of ML projects is often uncertain *ex ante*, and the development process entails a high degree of experimentation (e.g., in preparing suitable input data or adjusting model parameters; e.g., Amershi et al., 2019; Choudhury et al., 2021). While many types of non-ML IT systems benefit from agile development practices,

the extensive experimentation that ML development requires exacerbates the need for fast, iterative processes (e.g., Amershi et al., 2019; Studer et al., 2021). Additionally, ML systems in operation necessitate continuous maintenance to ensure their performance with current and future input data, causing the boundaries between ML development and operation to blur (e.g., Studer et al., 2021).

Moreover, ML systems continue to pose challenges beyond their development that differ from those of traditional IT systems (e.g., Benbya et al., 2021; Berente et al., 2021). Berente et al. (2021) delineate the distinct characteristics of ML systems causing these unique challenges in three facets: autonomy, learning, and inscrutability. *Autonomy* denotes the capacity of ML to make decisions without requiring the input of human operators (e.g., Ågerfalk, 2020; Baird & Maruping, 2021). ML systems can thus take over tasks from humans, even tasks previously held firmly in human grasp like processing natural speech (e.g., Schuetz & Venkatesh, 2020), and can intentionally complement, constrain, or entirely substitute for humans at work (Murray et al., 2021). This previously unseen level of agency from IT systems blurs the line between human and machine capabilities and challenges assumptions about how humans and IT systems interact and produce outcomes (e.g., Schuetz & Venkatesh, 2020), promising both great potential and novel difficulties for the organizational processes of the future. *Learning* is the previously described ability of ML systems to automatically improve themselves by deriving patterns from data (e.g., Mitchell, 1997). However, with contemporary ML systems still belonging to the category of narrow AI, they require human experts to guide their learning process by framing the (business) problem to solve and interpreting the results (e.g., Seidel et al., 2020; Salovaara et al., 2019). As ML systems cannot react to scenarios for which they did not receive examples in their training process (in the form of data; e.g., Dennett, 2006), humans must engage with the learning aspect of ML, both when solving a new business problem and to ensure productive ML systems are continuously re-trained with all relevant scenarios. *Inscrutability* means that the statistical algorithms used in the creation of ML systems, such as neural networks, have exploded in complexity, due to ever-growing data and processing power availability, making them and their behavior incomprehensible to humans (e.g., Asatiani et al., 2021; Faraj et al., 2018). Therefore, decision-makers must take great care when choosing the use case and degree of autonomy of the ML system, especially as the self-learning systems can produce unexpected results (e.g., Benbya et al., 2020b). With the decisions of the ML system inscrutable to human experts, such unexpected results can lead to unanticipated consequences for organizations (e.g., Asatiani et al., 2021; Benbya et al., 2021), for instance, when an ML system considers a candidate's marital status in making employment decisions (e.g., Martin, 2019a). Organizations must thus carefully consider the impact of these particularities of ML systems when seeking to benefit from their unique capabilities.

2.2 Business Models

The concept of business models has been at the center of a rapidly increasing amount of scientific work in recent years (DaSilva & Trkman, 2014; Massa et al., 2017). In essence, a *business model* delineates the business logic of an organization, illuminating how the organization creates value through its products or services, how it delivers said value to its customers, and how its revenue, costs, and profits are structured to allow for profitability and long-term survival (Teece, 2010). More briefly, "a business model describes the rationale of how an organization creates, delivers, and captures value" (Osterwalder & Pigneur, 2010, p. 14). To this end, it further explains how an organization "chooses to connect with factor and product markets" (Zott & Amit, 2008, p. 3).

Therefore, the business model explains how an organization not only achieves its goals through the creation of its products or services from its inputs but also how it integrates itself into a supply chain or value network to create economic value (Ritter & Lettl, 2018). Conceptualizations of the business model often illuminate the overarching business logic through the components making up an organization's business model and the resulting relations between them (e.g., Al-Debei & Avison, 2010; Teece, 2010; Zott & Amit, 2010). As an example, Figure 1 depicts one such conceptualization that is regularly adopted in research (e.g., Weber et al., 2022), the *Business Model Canvas* by Osterwalder and Pigneur (2010), also immensely popular among practitioners (e.g., Bigelow & Barney, 2021). While the exact components differ between researchers, with Shafer et al. (2005) finding 42 different components within just twelve conceptualizations, most fall into four categories (Burkhart et al., 2011; Morris et al., 2005). *Offering components* express the nature of the product or service being created as well as how the organization delivers said offering to the customers. *Market components* describe the customers to whom the organization delivers said value, their characteristics and interaction requirements, as well as the organization's position in the value chain. *Internal components* are the organization's core competencies, delineating the capabilities or activities that make up the source of its competitive advantage, i.e., that enable its value creation. *Economic components* explain the organization's logic for earning profits, i.e., for capturing value, and include the organization's revenue model, cost structure, and ability to achieve margins. Moreover, regarding the relation between components, the literature exclusively considers business model components to be interdependent, meaning that changes in one component can induce changes in other components and vice versa (Burkhart et al., 2011).

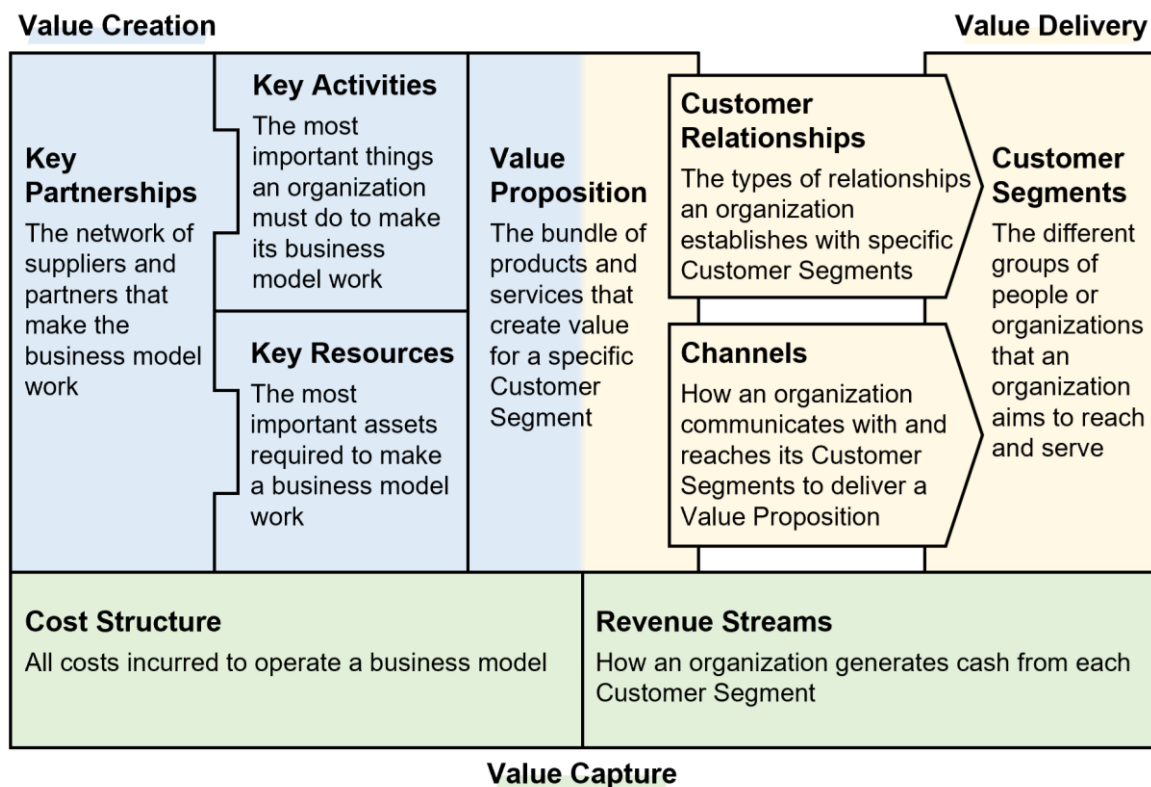


Figure 1. The Business Model Canvas as an Exemplary Conceptualization of Business Models and Their Components, With Added Categorization Into Value Creation, Delivery, and Capture (Adapted From Osterwalder and Pigneur, 2010, Strategyzer.com, CC BY-SA 3.0).

The extensive literature on the topic has thus explored what a business model is and brought forth various definitions (e.g., Zott et al., 2011; Al-Debei & Avison, 2010; Casadesus-Masanell & Ricart, 2010), has characterized constituting components and their interconnections (e.g., Osterwalder & Pigneur, 2010; Teece, 2010; Zott & Amit, 2010; Burkhart et al., 2011; Remane et al., 2016), and has further examined methods for developing novel innovative business models (e.g., Spieth et al., 2014; Foss & Saebi, 2015), as well as many more topics. Yet, thus far, the concept still lacks consensus on its definition as well as conceptual clarity (Ritter & Lettl, 2018; Casadesus-Masanell & Ricart, 2010). The diversity of definitions is at least partly attributable to the fact that scholars from different literatures have examined business models from different subject-specific perspectives, resulting in the emergence of various interpretations of the term (Massa et al., 2017; Zott et al., 2011). For instance, the concept has gained traction in not only strategy (e.g., Casadesus-Masanell & Ricart, 2010; Teece, 2010), technology and innovation management (e.g., Leppänen et al., 2023; Massa & Tucci, 2014), entrepreneurship (e.g., Chalmers et al., 2021; George & Bock, 2011), and IS (e.g., Aubert & Chan, 2024; Böttcher et al., 2022; Hartmann et al., 2016) research but also further fields such as social entrepreneurship (e.g., Seelos & Mair, 2007) and environmental sustainability (e.g., Schaltegger et al., 2012). George and Bock (2011) note that “the academic literature on this topic is fragmented and confounded by inconsistent definitions and construct boundaries” (p. 83). While such definitional and conceptual ambiguity is not uncommon when scholars are developing novel concepts with wide-ranging applicability (see Zott et al., 2011), the literature on business models has developed separately from strategy literature, despite their overlapping intent to explicate variations in firm performance (Lanzolla & Markides, 2021; Porter, 2001; Zott et al., 2011). A debate exists in the strategic management literature on the value of the business model, with strategy scholars expressing concerns, e.g., that “the concept of a business model has no established theoretical grounding in economics or in business studies” (Teece, 2010, p. 174) and may merely be a re-interpretation of concepts already well-established in strategy literature (e.g., Arend, 2013; Bigelow & Barney, 2021). For instance, the internal resources and capabilities that organizations must manage and leverage for value creation in the business model perspective (e.g., Osterwalder & Pigneur, 2010) already appear at the core of the resource-based view (Barney, 1991; Helfat & Peteraf, 2003). Therefore, whether the business model is meaningfully distinct from existing strategy concepts remains an issue of debate (e.g., Prescott & Filatotchev, 2021), for which Bigelow and Barney (2021) provide a discussion of potential distinctions. While the authors doubt significant differences, they concede that the business model view may be particularly suitable for conveying strategy fundamentals to decision-makers engaged in complex and fast-changing environments, due to its overwhelming popularity among managers and entrepreneurs. Moreover, they argue that the concept may offer further advantages in highlighting the interdependencies within organizations as relevant factors explicating firm performance in empirical analyses (Bigelow & Barney, 2021).

Other scholars are more optimistic about the value that the business model view can contribute to the strategy literature. Snihur and Eisenhardt (2022) even suggest that the business model might dethrone strategy as most important source of competitive advantage due to its strategic complexity and dynamism in comparison to the simpler, more static views of traditional strategy such as the resource-based view. Lanzolla and Markides (2021) posit that the attempt to establish the business model as a new concept may indeed have been futile but has also distracted scholars from more fruitful endeavors. They propose focusing instead on exactly the interdependencies between a business model’s activities, internal and external to the organization, that can provide a novel lens through which to examine strategy and contribute new theoretical insights (Lanzolla & Markides, 2021). Thus, while the business model perspective may or may not count as an

entirely new field, it relaxes some of the assumptions of traditional strategic management theories (e.g., the resource-based view or the positioning school), allowing the development of novel insights as an extension of strategy (Massa et al., 2017). In this regard, the business model can serve as the “webbing” between theories, allowing ideas from one theory to inspire aspects of other theories by providing a logical connection between them and their elements (e.g., modeling the value of resources as demand, thereby connecting the resource-based view and the demand-side perspective; Ritter & Lettl, 2018; Burström et al., 2021). Additionally, in IS literature, the business model concept is regarded as the “missing link” connecting strategy, processes, and IT (Veit et al., 2014). The business model perspective is therefore well-suited to investigating and comprehending technological change from the business value side (Steininger, 2019). Consequently, it often serves as a lens in IS research to examine how novel IT affects existing business models or drives the creation of new ones (e.g., Bock & Wiener, 2017; Hartmann et al., 2016; Spiegel, 2016). Thereby, IT can take different roles in business models, ranging from facilitating operations internally, mediating at the customer interface, or being the product itself, to a ubiquitous presence in all components of the business model (Steininger, 2019). In turn, the role that IT plays in the business model significantly influences the organization’s business logic (Steininger, 2019; Veit et al., 2014). The business model view can thus support examining not only how IT can be integrated into an organization and its departments but also, in particular, IT’s impact on how organizations create, deliver, and capture value. This dissertation therefore adopts the business model perspective due to both its suitability for investigating the interplay between IT and business logic and its interconnectivity with strategy theories. Moreover, the dissertation follows the definition of the term *business model* as an activity system with an organization at its core, which encompasses the focal organization’s internal and external activities and, through their interdependencies, links the organization’s value creation and delivery with value capture (e.g., Afuah, 2003; Hedman & Kalling, 2003; Lanzolla & Markides, 2021; Ritter & Lettl, 2018; Snihur & Eisenhardt, 2022; Teece, 2010; Zott & Amit, 2010).

2.3 The Need to Revisit Value Creation, Delivery, and Capture with Machine Learning

In the context of ongoing digitalization, advancements in IT, such as the ever-increasing availability of data or progress in data analytics, can profoundly change the business models of organizations, or enable the development of entirely new types of business models (e.g., Hartmann et al., 2016; Steininger, 2019; Veit et al., 2014). In this regard, ML in particular is unlocking an enormous variety of entrepreneurial opportunities for organizations (e.g., Chalmers et al., 2021; Obschonka & Audretsch, 2020). Consequently, their high level of disruptive potential is pressuring organizations to develop new or redesign their existing business models, or risk losing their competitive advantages to the growing number of ML-infused organizations, such as Alibaba, Amazon, Google, or Uber (e.g., Lee et al., 2019; Wamba-Taguimdje et al., 2020). In turn, organizations that integrate ML into their business models can expect improved firm performance through the abilities of ML systems, e.g., in automating and optimizing processes, informing and interacting with humans, or enabling the creation of entirely novel products and services (e.g., Benbya et al., 2020b; Benbya et al., 2021; Burström et al., 2021; Davenport et al., 2020; Wamba-Taguimdje et al., 2020). For example, ML systems can enhance the productivity of legal professionals by assisting them with non-routine tasks (i.e., client-facing or creative work) and automating such routine tasks as analyzing contracts and conducting corresponding legal research (e.g., Armour & Sako, 2020). Different ML systems can aid engineers in developing novel materials (e.g., Correa-Baena et al., 2018) or support radiologists by detecting malignant nodules in computer tomography scans (e.g., Pumplun et al., 2023). Thereby, the ML system can take

various roles described in Section 2.2 (i.e., facilitator, mediator, outcome, or ubiquity; see Steininger, 2019) within or outside the organization's structural boundaries, with some organizations utilizing the technology ubiquitously throughout their business model. For instance, Uber manages the entirety of its more than six-million-driver remote workforce through an ML-powered app that decides which routes to take and which passengers to pick up (Uber, 2024). While the ML system is therefore essential to the success of the organization's business model, it also begets distinct challenges. For instance, Uber must address its algorithmic management approach's potentially negative effects on the drivers (e.g., incomprehensibility of the ML system and its decisions or lack of socialization at work) or risk adversarial action or increased driver turnover (Möhlmann et al., 2021; Wiener et al., 2023). When ML systems play an integral part in the business models of organizations, their unique characteristics (detailed in Section 2.1) significantly influence how organizations create, deliver, and capture value (e.g., Benbya et al., 2020a; Burström et al., 2021; Steininger et al., 2022). This dissertation conceptualizes such business models as ML-driven business models (see Sections 6.2.1 and 7.2.1). However, exactly how these ML influences manifest themselves is still largely uncertain, with many business models incorporating ML being nascent to date. Important current issues of organizations, regarding value creation, value delivery, and value capture with ML appear below, and respective studies in the following chapters investigate them in detail.

Regarding an organization's **value creation** with ML, the characteristics of ML systems pose requirements that must be taken into account during their development processes. For instance, the data-based learning approach of ML systems encourages fast development cycles with quick feedback mechanisms to identify and focus on suitable solutions (e.g., Amershi et al., 2019; Shearer, 2000). Moreover, ML systems often continue learning after their development (to the organization's benefit) and therefore require continuous maintenance even after deployment, blurring the lines between their development and operation (e.g., Amershi et al., 2019; Studer et al., 2021). However, the aforementioned need of ML development for quick exploratory processes conflicts with the needs of conventional business processes to exploit the organization's existing strengths (e.g., Gerbert et al., 2020; Pumplun et al., 2019), leading to tensions when creating ML systems as organizations must allocate their limited resources. Chapter 4 presents an in-depth examination of these tensions and potential alleviating tactics under the consideration of different structural set-ups for ML development. Beyond such challenges, the learning capability of ML systems can also prove uniquely valuable to organizations. Scholars currently assume that this capability merits seeing ML systems as a new type of learner alongside humans (e.g., Argote et al., 2021; Sturm et al., 2021b) that can help organizations make sense of data from their environment (e.g., regarding competitors or customer responses) more easily and rationally than through humans alone (e.g., Benbya et al., 2021). Simultaneously, the new ML learners can also adversely affect learning processes occurring in organizations, e.g., by reducing the diversity in routines and background knowledge flowing into the organizational knowledge base (e.g., Balasubramanian et al., 2022). Organizations are well advised to let humans and ML systems learn in tandem, to complement and amplify their respective capabilities (e.g., Lyytinen et al., 2021). However, the literature on such joint learning has thus far focused on productive ML systems, leaving potential learning during ML development unexplored. Yet, as the creation of ML systems is a collaborative effort of experts from various disciplines who engage with their existing knowledge in collective sense-making and evaluation processes while interacting with data and prototypical ML systems (see, e.g., Studer et al., 2021), ML development may act as a unique field of interaction for potential learning driven by the experts' encounters (Nonaka, 1994). The study that Chapter 3 details thus illuminates the various learning processes that are possible during the creation of ML systems.

With respect to **value delivery**, organizations must ensure that their customers can receive the value they create with the help of ML systems. In particular, the inscrutability of ML can constitute a significant hurdle in this endeavor, as for users to consider the output of ML systems in their decision-making and thus benefit from them, they must be able to understand both the systems and the reasoning for their output (e.g., Barredo Arrieta et al., 2020; Martin, 2019b). Otherwise, organizations may risk unexpected outcomes, such as problematic ethical implications (e.g., considering a candidate's marital status for employment decisions or integrating race into decision-making on allocating city police forces; e.g., Benbya et al., 2021; Martin, 2019a). To counteract their inherent inscrutability, researchers have brought forth XAI approaches to reveal the inner workings of ML systems as well as the reasoning behind their output (e.g., Ågerfalk, 2020; Rai, 2020; Rudin, 2019). However, to date, XAI research has largely focused on offering explanations to ML experts, while neglecting the user group that many ML-driven business models intend to deliver value to, namely lay users, who seek to benefit from the ML systems, e.g., by utilizing them in their decision-making (Ellenrieder et al., 2023; Gaube et al., 2023). To address this issue, Chapter 5 presents a design science study that develops an approach to designing ML systems in a way that fulfills the different needs for explainability of each of their user groups and specifically takes lay users, such as managers, into account.

Last, concerning **value capture**, organizations must take all aforementioned particularities of ML into account when looking to create and maintain competitive advantage, and ultimately profit from the development of business models infused by the technology. As previously mentioned, an organization's possibilities for value capture are strongly interleaved with the surrounding components of its business model (e.g., Teece, 2010; Tidhar & Eisenhardt, 2020), which can be seen in the struggle of many digital start-ups to capture value from their inherently difficult-to-monetize value proposition (e.g., Steininger et al., 2022). It took Twitter (now X), for instance, over a decade to identify a suitable revenue model to sustainably profit from its massive user base (Mangalindan, 2010; Steininger et al., 2022). The structural composition in which all business model components are arranged and interconnected thereby determines an organization's options to enhance its position in the market in relation to its competitors, to respond to disruption, and to enter new or defend existing markets (e.g., Snihur & Markman, 2023). Therefore, the taxonomic study in Chapter 6 scrutinizes the structural composition of ML-driven business models to lay the groundwork for a holistic investigation of how organizations can sustainably capture value with the help of ML. It sheds light on the components constituting ML-driven business models (including, e.g., ML-specific value propositions and revenue models) and delineates business model archetypes commonly found in practice. Additionally, the realization of business models in general is a difficult and inherently uncertain process that requires entrepreneurs to experiment with elements of their business model (e.g., Burnell et al., 2023; Snihur & Eisenhardt, 2022), which might be further complicated by the uncertain and experimental nature of ML development (Amershi et al., 2019; Choudhury et al., 2021). Organizations seeking to establish and capture value from ML-driven business models in the long term may thus need to build new capabilities, e.g., due to the aforementioned ML-induced uncertainty, the heightened importance of ecosystems for ML-driven business models, or the high dynamism of their environment (e.g., Burström et al., 2021; Chalmers et al., 2021; Lange et al., 2021; Steininger et al., 2022). In particular, dynamic capabilities play an integral part in developing business models in uncertain and dynamic environments (see Ricciardi et al., 2016; Schoemaker et al., 2018; Steininger, 2019; Teece, 2018). Therefore, an investigation of how organizations can build the necessary dynamic capabilities to realize ML-driven business models with (economic) success appears in Chapter 7 of this dissertation.

3 Research Paper A: Potential for Learning During Value Creation With Machine Learning

Title: Machine Learning Developments as Stimuli for Organizational Learning
Authors: Vetter, Oliver A.³; Sturm, Timo³; Fecho, Mariska; Buxmann, Peter
Published in: Proceedings of the 44th International Conference on Information Systems (ICIS 2023)

Abstract

Organizational learning is a fundamental process that defines organizational behavior and thereby strongly influences organizational performance. As organizations increasingly adopt machine learning (ML) systems in their routines, the need to illuminate the impact of learning machines on organizational learning processes becomes increasingly urgent. In particular, due to their highly interdisciplinary and collaborative nature, ML developments—as organizations' activities aimed at creating productively usable ML systems—may hereby represent an important but not yet well understood mechanism for fostering organizational learning. To explore how ML developments affect organizational learning processes, we interviewed 42 experts who are frequently involved in ML developments. Our findings suggest that ML developments can enhance organizational learning by stimulating a variety of organizational learning processes that generate a wealth of explicit and tacit knowledge in organizations.

Keywords: Organizational learning, machine learning, development, knowledge, tacit, explicit

³ Note: Both authors contributed equally to this research.

3.1 Introduction

Organizational learning lies at the heart of organizational behavior because it represents the process by which organizations continually (re)define their norms, innovations, and routines (e.g., Argote & Miron-Spektor, 2011; Levitt & March, 1988). Learning thereby enables organizations to effectively adapt to their environment, which is critical to their success and, if done wrong, can even jeopardize the organizations' survival (e.g., Argote et al., 2021; Huber, 1991; March, 1991). Since organizations cannot learn on their own, they must rely on the learning of their members (e.g., Fang et al., 2010; March, 1991). Therefore, facilitating and effectively coordinating the individual learning processes and interactions of its members represent the key issue behind effective organizational learning, which has already spurred decades of research (great overviews exist such as, e.g., Argote et al., 2021; Argote & Miron-Spektor, 2011; Fiol & Lyles, 1985; Huber, 1991). However, as human members have traditionally been the ones who learn in organizations, the limitations of human cognition have complicated organizational learning since its very beginnings (e.g., Levinthal & March, 1993; March, 2006, 2010; Simon, 1991).

Hoping to overcome the limitations of human cognition, organizations are increasingly recognizing the great potential of information systems (IS) based on machine learning (ML) algorithms (e.g., Berente et al., 2021; Benbya et al., 2021). Such ML systems can learn autonomously by inferring patterns from data to create models for guiding behavior (e.g., Mitchell, 1997; Russell & Norvig, 2021). Because of their ability to learn, recent research now also considers ML systems as a new type of organizational learner besides humans; that is, not just as another tool that only supports human learning (like traditional non-ML IS; e.g., Argote et al., 2021; Alavi & Leidner, 2001; Kane & Alavi, 2007), but rather as active learners that can contribute their own knowledge to organizational learning (e.g., Argote et al., 2021; Balasubramanian et al., 2022; Sturm et al., 2021b). Indeed, the great computational power of ML systems, which have already beaten the best human Go player, surpassed human capabilities in object recognition, and recently begun to revolutionize the way we write text, makes them increasingly appear as a panacea for improving organizational efficiency and effectiveness (e.g., Jordan & Mitchell, 2015; Lindebaum et al., 2020).

Despite this great potential, many organizations struggle to create productive ML systems. As a result, many organizations are afraid to waste their scarce resources and reduce their funding for ML developments (e.g., Ransbotham et al., 2020; Sculley et al., 2015), which are typically projected efforts focused on developing ML systems for productive use (e.g., Amershi et al., 2019; Studer et al., 2021). While this may seem like an appropriate risk mitigation strategy at first glance, it may also become a fatal fallacy, as such ML developments may act as a powerful form of organizational learning: from data selection to ML model evaluation to knowledge sharing with ML systems (e.g., Amershi et al., 2019; Wirth & Hipp, 2000), the development and use of ML systems involves an interdisciplinary learning process among an organization's domain experts, data scientists, and/or ML algorithms that would not otherwise occur (e.g., knowledge sharing sessions between data scientists and engineers to develop a predictive maintenance solution). Since ML developments require reflection on existing knowledge (e.g., through collective data exploration) and can enable new knowledge creation beyond the mere development of ML systems (e.g., insights into flaws in routines), they may have a significant impact on the use, extension, and retention of an organization's knowledge. ML developments may therefore play a crucial role in facilitating and stimulating organizational learning. The discontinuity of ML developments may thus lead to a critical competitive disadvantage in the long run—even if an organization's ML developments frequently fail to yield productive ML systems.

So far, however, the role of ML developments for organizational learning has been largely neglected by research. Unfortunately, research on the general impact of IS on organizational learning can only be of limited help, as existing studies primarily focus on the impact of IS use rather than IS development (e.g., Alavi & Leidner, 2001; Argote et al., 2021; Argote & Miron-Spektor, 2011; Kane & Alavi, 2007) and naturally neglect the particularities of ML systems (e.g., involvement of data scientists and system functionality being defined primarily by data analysis rather than human-defined rules; Amershi et al., 2019; Brynjolfsson & Mitchell, 2017; Sturm et al., 2021a). Only recently have a few scholars begun to explore the impact of productive ML systems as a new type of organizational learner on organizational learning (e.g., Argote et al., 2021; Balasubramanian et al., 2022; Lyytinen et al., 2021; Ransbotham et al., 2020; Sturm et al., 2021a, 2021b). Yet, research that takes a broader view to include the impact of *human* learning processes that occur within ML developments on organizational learning is still lacking, leaving organizations without clear guidance. To help organizations better manage and grasp the value of their ML developments, we aim to explore the impact of ML developments on different organizational learning processes. We thus ask the following research question (RQ): *How do ML developments affect organizational learning processes?*

To answer our RQ, we adopt a qualitative research approach that allows us to explore and synthesize the experiences of 42 experts who have been frequently involved in ML developments. Our findings suggest that ML developments can indeed contribute value to organizational learning. We further find that these contributions involve different learning processes depending on the ML development phase. Our result is a framework for how ML developments can stimulate different types of organizational learning processes.

3.2 Theoretical Background

We first introduce organizational learning, with a particular focus on Nonaka's (1994) seminal work on the knowledge creation spiral. We then turn to the foundations of ML and the typical processes within ML developments that we revisit as a form of problem solving. Table 3 summarizes the core concepts that we introduce in both subsections and which will act as a theoretical structure for our subsequent analysis. We conclude by crystallizing the need to revisit organizational learning in the context of ML developments.

Table 3. Core Knowledge Types and Core Learning and Problem-Solving Processes.

Source	Concept	Definition	Examples
Nonaka (1994)	Knowledge types	Explicit knowledge: Knowledge that is codified in some form and easy to articulate.	<i>Manuals, standard operating procedures, and patents</i>
		Tacit knowledge: Knowledge that is deeply rooted in personal experience and difficult to articulate.	<i>Riding a bike, playing an instrument, and dancing a choreography</i>
	Learning processes	Socialization (tacit to tacit): Integrating tacit knowledge by creating shared experiences.	<i>Apprenticeship, mentoring, and immersion in a community of practice</i>
		Externalization (tacit to explicit): Codifying tacit knowledge into explicit knowledge.	<i>Capture personal experiences as a metaphor, a concept, or a story</i>
		Combination (explicit to explicit): Integrating different explicit knowledge by combining documented concepts.	<i>Integrate rule sets, documents, and/or frameworks</i>
	Internalization (explicit to tacit): Creating tacit knowledge by gaining personal experience with explicit knowledge.	<i>Apply textbook knowledge about skating, singing, and cooking to gain personal experience with it</i>	
Basadur et al. (1982)	Problem-solving processes	Problem finding: Recognize and construct relevant problems.	<i>Identify and describe bottlenecks and decision-making flaws</i>
		Problem solving: Search for adequate solutions by exploring potential solutions for given problems.	<i>Define quality requirements, test and compare solution candidates</i>
		Solution implementation: Integrate selected solutions into organizational processes.	<i>Provide additional tools to accelerate a time-consuming process</i>

3.2.1 Organizational Learning

Organizational learning is the process of gathering experience in some organizational context and deriving knowledge from that experience to guide future actions (e.g., Argote & Miron-Spektor, 2011; Levitt & March, 1988; March, 2010; Nonaka, 1994). Experience thereby denotes one's own or others' unit of task performance (e.g., making decisions, observing others performing routines). Learners' experiences are thus assumed to primarily contain recollections of chosen actions and their consequences, enriched with information about the context in which the actions took place (e.g., Argote et al., 2021; Argote & Miron-Spektor, 2011). Learners then learn from the available experience by inferring conclusions from it, attempting to generalize and reconcile the collected knowledge (e.g., decision rules, approaches to performing routines) with existing knowledge (e.g., Argote & Miron-Spektor, 2011; March, 2010). Thereby, learners can learn from others' experiences either directly (e.g., by drawing conclusions from observing others' actions) or indirectly by integrating others' knowledge with their own knowledge (e.g., by integrating their own and others' decision rules; e.g., Argote & Miron-Spektor, 2011; Nonaka, 1994).

To further nuance organizational learning processes, Nonaka (1994) famously introduced the now seminal theory of the knowledge creation spiral. In his theory, Nonaka (1994) first introduces two core knowledge types (i.e., *explicit* and *tacit* knowledge) and then outlines four core learning processes that focus on integrating and converting between the two knowledge types (see also

Table 3 for an overview). First, **explicit knowledge** refers to knowledge that is codified or documented in some form, such as written or electronic documents, databases, or other tangible formats. Explicit knowledge can be easily articulated, shared, and accessed by others. Examples include textbooks, technical manuals, standard operating procedures, patents, and scientific articles. Second, **tacit knowledge** refers to knowledge that is difficult to articulate, codify, or transfer to others. This type of knowledge is often deeply rooted in an individual's personal experience, intuition, and insights, and it may be difficult to formalize or explain in words. Tacit knowledge is typically acquired through personal experience, observation, and informal learning. For instance, tacit knowledge are skills like riding a bike, playing an instrument, and dancing a choreography as they involve a complex interplay of physical and mental abilities that are difficult to describe in words and thus hard to teach explicitly through written or spoken instructions.

To optimize knowledge creation, Nonaka (1994) argues that organizations must enable a continuous conversion between both types of knowledge. To do so, Nonaka (1994) identified four iteratively connected processes: First, **socialization** involves the sharing and integration of tacit knowledge by creating shared experiences through social interaction, observation, and emulation. For instance, socialization includes learning from others through apprenticeship, mentoring, and immersion in a community of practice. By developing a shared understanding of a particular domain within a specific context, humans gather personal experience that helps them approximate others' tacit knowledge. Second, **externalization** focuses on the transformation of tacit knowledge into explicit knowledge. By doing so, externalization can help facilitate the communication of tacit knowledge to some extent. For instance, this includes capturing personal experiences as a metaphor, a concept, or a story. Third, **combination** focuses on creating new explicit knowledge by integrating different explicit knowledge that already exists. Here, individuals synthesize different forms of explicit knowledge, such as rule sets, documents, and frameworks. Thereby, combination allows individuals to build on established knowledge to create new insights and perspectives. Fourth, **internalization** involves the transformation of explicit knowledge into tacit knowledge. By applying explicit knowledge to the particularities of specific contexts, individuals can gain a deeper understanding of the knowledge and its applications, and thereby develop new tacit knowledge based on the gathered experiences. For instance, a skateboarder can internalize textbook knowledge of a certain trick, by enriching the knowledge with their own related experience by trying the trick themselves.

Nonaka (1994) theorizes the spiral of knowledge creation as an iterative process that continuously traverses all four learning processes, shifting between explicit and tacit knowledge. Because of the spiral's iterative nature, effective organizational learning requires an organization to facilitate each of the four learning processes to enable well-functioning, mutually stimulating knowledge creation.

3.2.2 Machine Learning Development

The approach behind modern artificial intelligence (AI) that has driven recent breakthroughs (e.g., ChatGPT or AlphaGo) is the use of ML algorithms (e.g., Berente et al., 2021; Brynjolfsson & Mitchell, 2017). ML algorithms allow IS to derive patterns from data to create ML models that are then used to solve given problems (e.g., deriving a rationale for how to grant a loan; e.g., Mitchell, 1997; Russell & Norvig, 2021). In doing so, ML resembles organizational learning: ML systems use experience (i.e., data capturing units of task performance) to infer conclusions from it (i.e., models of how to perform a task), thereby attempting to generalize and reconcile the contained knowledge (e.g., Sturm et al., 2021b). In contrast to traditional non-ML IS that were only able to

support human learning (e.g., emails, Zoom, or repositories supporting human knowledge exchange; e.g., Alavi & Leidner, 2001; Kane & Alavi, 2007), this enables ML systems to learn autonomously and contribute their own knowledge to organizations' stock of knowledge (e.g., Argote et al., 2021; Sturm et al., 2021b). As a result, ML systems can be viewed as a new type of organizational learner besides humans (e.g., Balasubramanian et al., 2022; Ransbotham et al., 2020; Sturm et al., 2021b).

ML systems are realized through ML developments, which are project-led efforts to develop and deploy ML systems productively (e.g., Amershi et al., 2019; Studer et al., 2021). Such ML developments are typically a highly iterative and interdisciplinary practice involving a diverse set of collaborating domain and technical experts (e.g., Amershi et al., 2019; Wirth & Hipp, 2000). For example, developing an ML system for predicting wind turbine failures requires domain experts with knowledge of wind turbine behavior to help define the problem, and data scientists to perform analyses and translate them into a suitable solution, which must then be evaluated again by the domain experts. To coordinate this process, numerous process frameworks have been proposed for ML developments (e.g., CRISP-DM, SEMMA, and KDD) that are widely used in practice (e.g., Amershi et al., 2019; Azevedo & Santos, 2008; Martinez et al., 2021; Martínez-Plumed et al., 2021). While these frameworks differ in their specific process steps, they share a common ground of key phases that resemble established phases of problem-solving processes as famously conceptualized by Basadur et al. (1982). Following this rationale, we can revisit these key phases for ML development contexts:

ML developments typically begin with a **problem finding** phase, which focuses on articulating a relevant problem with data that a planned ML system is intended to solve. First, this requires experts to identify problems that are organizationally meaningful and adequate for ML systems to solve. This includes activities such as exploring business domains and available data to identify bottlenecks and flaws in organizational processes or practices. Second, once problems have been identified, experts need to create a concise and realistic representation of the identified problems. This includes activities such as exploring the availability of data, selecting a data sample that is a representative and comprehensive collection of exemplary problem instances, and preparing data to ensure high data quality (e.g., accuracy, completeness, timeliness, and consistency) and to develop additional variables to describe problem instances more holistically. The second phase of ML developments typically focuses on **problem solving**, which involves the search for ML models that adequately suit the identified problems. This first requires the organization to develop a shared understanding of what an appropriate ML model should entail. This includes activities such as defining baselines for evaluation metrics for when an ML model is good enough to be used, and clarifying transparency requirements for ML models (e.g., requiring highly transparent models such as decision trees). Second, once the requirements for appropriate ML models have been defined, experts search for possible ML models by creating and evaluating a variety of ML models, and identify the seemingly best ML model. This includes activities such as selecting ML algorithms, parameterizing ML algorithms, and training, testing, and comparing prototypical ML models. Lastly, the third phase focuses on **solution implementation**, which aims to integrate ML systems into organizational processes. This typically involves redesigning established processes to include ML systems and to design and sustain their emerging interplay with humans. For instance, this requires organizations to rethink how inputs and outputs of ML systems can be integrated to ensure effective workflows and how to adapt the role of involved humans.

As ML developments are therefore a highly interdisciplinary process that involves collective reflection on given problems, potential solutions, and the integration of knowledge between

multiple human experts (and ML systems), organizations' ML endeavors may not only impact organizational learning through productively used ML systems that now learn side-by-side with humans but also by the additional human learning processes that are stimulated by the processes entailed in ML developments.

3.2.3 *The Need to Revisit Organizational Learning in the Context of ML Developments*

For decades, researchers have analyzed the processes of *human-driven* organizational learning and how organizations can effectively coordinate these processes to improve organizational performance (e.g., Argote et al., 2021; Argote & Miron-Spektor, 2011; Fiol & Lyles, 1985; Huber, 1991). However, with their ability to learn, ML systems increasingly participate as a new type of learner in organizational learning alongside humans (e.g., Ransbotham et al., 2020; Sturm et al., 2021b). While research has recently recognized the importance of understanding ML's impact on organizational learning (e.g., Argote et al., 2021; Balasubramanian et al., 2022; Berente et al., 2021; Sturm et al., 2021b), the few studies that exist focus only on the impact of ML systems that are already in *productive use*. While these studies provide crucial insights on how to effectively manage the resulting learning dynamics between humans and ML systems, this focus neglects the potential impact of the underlying ML development activities required to enable such productive ML systems in the first place. This is problematic because research that overlooks ML developments may thereby neglect a novel context of analytically-driven interactions between domain and technical human experts that are likely to offer new opportunities for stimulating organizational learning processes (e.g., human experts sharing their domain knowledge to prepare data for training and evaluating a planned ML system). Moreover, if organizations frequently fail to produce productively usable ML systems, they run the risk of misjudging the impact of reducing their ML developments, which may help them save scarce resources in the short term, but may become a fatal long-term fallacy if they thereby inhibit valuable organizational learning processes. While existing research emphasizes the need for further analysis of the impact of ML systems on organizational learning due to such consequences (e.g., Argote et al., 2021; Sturm et al., 2021b), research that also considers the impact of the preceding ML developments remains non-existent. As with existing ML research, research on the general impact of IS on organizational learning focuses mainly on IS use rather than IS development (e.g., Alavi & Leidner, 2001; Argote et al., 2021; Argote & Miron-Spektor, 2011; Kane & Alavi, 2007), and inherently neglects the particularities of ML developments (e.g., Amershi et al., 2019; Brynjolfsson & Mitchell, 2017; Sturm et al., 2021b)—also leaving us with limited help in unpacking the impact of ML developments. As a result, the current discussion runs the risk of being too narrowly focused, which can lead to ill-informed decisions for organizations. Hoping to contribute to broadening the perspective of the current discussion, we now turn to our study to analyze how ML development activities contain opportunities to stimulate different learning processes and thus can serve as important mechanisms to improve the long-term performance of organizations.

3.3 **Qualitative Research Methodology**

Due to the lack of research on organizational learning in ML developments, we pursued a qualitative research approach through interviews with professionals from various industries. Expert interviews are one of the most important data collection tools in research contexts that lack sufficient empirical evidence (Myers & Newman, 2007), allowing us to examine ML developments' impacts on organizational learning in a wide variety of contexts. Our goal is to

develop a theoretical foundation for organizational learning processes involving humans and ML systems that occur during the different ML development phases.

We conducted 42 in-depth semi-structured interviews with experts who are frequently involved in ML developments. Each interview was conducted online and lasted 55 minutes on average. The experts were recruited from our personal and professional networks, primarily through LinkedIn (Butts et al., 2015). To broaden the scope of our analyses and develop a more general theory, we interviewed experts from a variety of different industries between March 2021 and July 2022 (Davison & Martinsons, 2016). In total, our experts' experiences span 15 industries, with software (23.8%), manufacturing (11.9%), and telecommunication (11.9%) being the most represented. Most experts work for inter- or multinational organizations. To capture diverse perspectives on ML developments, we included both technology- and domain-oriented experts. Having worked on various ML developments over several years, the experts can offer a deep and diverse set of experiences. Table 4 provides an overview of the interviewed experts.

Table 4. Overview of the Interviewed Experts.

ID	Position	Sex	Exp.	Industry	ID	Position	Sex	Exp.	Industry
I1	Data scientist	m	3 yr.	Aviation	I22	Data scientist	m	12 yr.	Trade
I2	Data scientist	w	6 yr.	Software	I23	Data scientist	m	6 yr.	Tourism
I3	Data scientist	m	8 yr.	Software	I24	Data scientist	m	3 yr.	Technical Testing
I4	Domain expert	m	5 yr.	Software	I25	Data scientist	m	4 yr.	Automotive
I5	Data scientist	w	6 yr.	Healthcare	I26	Data scientist	w	8 yr.	Software
I6	Domain expert	m	9 yr.	Telecomm.	I27	Domain expert	w	6 yr.	IT Consulting
I7	Domain expert	m	12 yr.	Software	I28	Manager	m	6 yr.	Telecomm.
I8	Data scientist	m	8 yr.	E-Commerce	I29	Manager	m	7 yr.	IT Consulting
I9	Domain expert	w	5 yr.	Energy	I30	Data scientist	w	3 yr.	Telecomm.
I10	Domain expert	m	6 yr.	Automotive	I31	Data scientist	m	7 yr.	Manufacturing
I11	Data scientist	m	7 yr.	Software	I32	Manager	m	3 yr.	Manufacturing
I12	Data scientist	m	5 yr.	Automotive	I33	Data scientist	m	6 yr.	Research
I13	Domain expert	m	4 yr.	Automotive	I34	Domain expert	m	3 yr.	Software
I14	Data scientist	m	3 yr.	Healthcare	I35	Data scientist	m	3 yr.	IT Consulting
I15	Domain expert	m	5 yr.	Manufacturing	I36	Manager	w	5 yr.	Infrastructure
I16	Data scientist	m	9 yr.	Market Research	I37	Data scientist	m	2 yr.	Manufacturing
I17	Data scientist	w	2 yr.	Healthcare	I38	Manager	m	6 yr.	Software
I18	Domain expert	m	4 yr.	IT Consulting	I39	Manager	w	7 yr.	Aviation
I19	Data scientist	m	7 yr.	Manufacturing	I40	Domain expert	w	7 yr.	Healthcare
I20	Data scientist	m	8 yr.	Telecomm.	I41	Manager	m	5 yr.	Software
I21	Data scientist	w	11 yr.	Software	I42	Domain expert	m	10 yr.	Telecomm.

Note. Experience has been shortened to Exp. and Telecommunication to Telecomm.

At the beginning of each interview, the interviewees were introduced to the topic and our RQ. The interview guide consisted of a series of open-ended questions and was divided into four sections: The first section aimed to familiarize the interviewees with the interview situation and to develop an understanding of the interviewees' general expertise (i.e., descriptions of participants' current position, past ML development experience, and current ML development involvement). The

subsequent three sections then focused on the three ML development key phases: Problem finding (i.e., how the interviewees identify appropriate problems and prepare related data), problem solving (i.e., how the interviewees train and evaluate ML models), and solution implementation (i.e., how ML systems are used productively in their organization). We used the guide as a basis for structuring our interviews according to the prepared questions, while also allowing for improvisation and spontaneous questions based on the particular course of the interview (Myers & Newman, 2007). After mutual consent, the interviews were recorded and transcribed to analyze the participants' responses. For interview I13, we were only able to take notes instead of making a recording.

As mentioned above, our goal is to identify learning processes that occur within ML developments. We analyzed the interview data by following the directed content analysis, which is appropriate for validating and extending existing theories (Hsieh & Shannon, 2005). An iterative multi-cycle, multi-researcher coding process with the following coding strategies was performed (Saldaña, 2015): First, we employed attribute coding to select descriptive and essential information about participants' characteristics. Second, we then utilized hypothesis coding to create primary codes according to the related literature of organizational learning. We used these primary codes, based on previous findings about human learning in organizations, as a structural framework to categorize the codes specific to the ML context that emerged in the following steps, including the codes *externalization*, *combination*, *internalization*, and *socialization* along each problem-solving phase (i.e., *problem finding*, *problem solving*, and *solution implementation*) of ML development. Third, we applied descriptive coding to uncover novel insights into learning processes during ML developments, which produced numerous subcodes for each primary code. Fourth, to make sense of the vast amount of codes and to remove redundant and irrelevant codes, we used pattern coding to integrate the identified codes into emergent themes that represent major subprocesses of the learning processes in the problem-solving phases (e.g., "*tacit knowledge about relevant variables becoming more structured and externalized when reflecting on possible data for a given problem*" was a code categorized into the theme "*domain experts externalize their problem understandings*" of the *externalization* processes in the *problem finding* phase). Each of us conducted these steps separately in each iteration, after which we intensively discussed our findings and synthesized commonalities into a shared code system to gradually develop a consensus on the codes. As part of the iterative process, we continued to collect new data to answer questions that arose during coding until no new insights emerged. As we observed no new insights in the final interviews, we noticed that we had reached theoretical saturation and stopped interviewing after the 42nd interview (Flick, 2004b). Finally, we formulated summaries (shown in Table 5) of the main themes and selected quotes from our experts that best represent the corresponding codes, which we describe below.

3.4 Results

Combining the four processes of Nonaka's (1994) knowledge creation spiral with Basadur et al.'s (1982) problem-solving phases provides us with a fruitful structure for exploring the learning processes that occur within ML developments. Table 5 summarizes the learning processes we have identified as occurring in the three phases of ML development, which we now describe below.

Table 5. Identified Learning Sub-Processes in ML Developments.

	Problem finding Main stimulus: Data scientists Main learning focus: Problems	Problem solving Main stimulus: Data scientists Main learning focus: Solutions	Solution implementation Main stimulus: ML systems Main learning focus: Solutions
Externalization (tacit → explicit)	Domain experts externalize their problem understandings.	Domain experts externalize their understanding of an adequate solution.	Domain experts' solutions get externalized by ML systems.
Combination (explicit → explicit)	Domain experts compare actual observed and expected problem behavior.	Domain experts integrate insights of candidate solutions with their solutions.	Domain experts integrate their solutions with those suggested by ML systems.
Internalization (explicit → tacit)	Domain experts experiment with newly discovered types of problem instances.	Domain experts experiment with solutions suggested by data scientists.	Domain experts experiment with solutions suggested by ML systems.
Socialization (tacit → tacit)	Domain experts share experiences and resolve conflicting problem understandings.	Domain experts exchange their solutions and resolve uncovered conflicts.	Domain experts share experiences and resolve conflicting solutions with ML systems.

3.4.1 Problem Finding

As detailed above, the problem finding phase of ML developments focuses on recognizing problems and constructing representations of those problems in data. During our interviews, it became apparent that data scientists serve as the main stimulus for the various learning processes that take place in this phase as they approach domain experts with derived data insights in different ways, which we describe below.

Externalization: Our experts emphasized that at the beginning of ML developments, domain experts (experts with deep knowledge of the focused problem domain but without deep data science knowledge; e.g., bankers, engineers) and data scientists (experts with deep data science knowledge but without deep domain knowledge) typically begin by participating in knowledge sharing sessions. While these sessions aim to help the data scientists gain a rough understanding of the problem domain, our interviews revealed that these sessions involve more than just transferring the domain experts' understanding of the problem to the data scientists. By having to articulate their often implicit understanding of problems to data scientists, who are typically laypeople in the given domain, domain experts are forced to reflect on their tacit knowledge of a problem and explain it in a simple and transparent way along explicit variables that can potentially be reflected in the data—an exercise that domain experts often only rarely have to perform in their day-to-day work in the domain, but which helps them better clarify their understanding to themselves and make it readily available to others (since their normal practice requires them only to act on the problem, but not to clearly explain their rationale for doing so):

“We often start with the data collection and then a bunch of us get together, look at some documents, including data scientists, product managers, and some other people, to see what’s in there. Then we talk to the [domain experts] to learn what they need, what they want. And then we come together and decide what the desired data collection looks like. [...] And through that process, we often learn a lot more about the problem itself.” (I34)

The experts further emphasized that these articulations are then carefully documented during the sessions because they are needed for subsequent data preparation:

“What exactly is the problem we focus on? What is the purpose of the project? What do we want to investigate in the first place? What do we want to predict somehow? That’s exactly what it means to initially try to understand the processes and to understand the problem correctly, and then to outline with the domain experts what kind of data we need.” (I17)

“You first have to find out what problem is behind the data, so creating something like a documentation or something like that is an issue.” (I6)

In this way, ML developments can help organizations clarify, express, and document critical tacit knowledge on the respective problem domain that has remained exclusive to individual experts, thereby making it accessible to others in the form of explicit knowledge.

Combination: Our experts stated that, based on the understanding gained in the initial knowledge sharing sessions, data scientists then typically begin to explore and prepare the available data to progressively understand the data and improve data quality. In doing so, our experts explained that data scientists engage in an ongoing exchange with domain experts to clarify their understanding of why certain data characteristics may occur (e.g., data distributions, correlations, outliers). To do this, data scientists typically confront domain experts with their objective data descriptions, and possibly their own “lay” hypotheses, and ask for an explanation. This requires domain experts to reflect on their expected and the actual explicit problem behavior observed in the data and formulate a plausible explanation. To do this, domain experts often first explicate their reasoning of why certain problem behavior (and thus data characteristics) can occur, and then use these explanations to challenge the actual observed problem behavior. To this end, this process can help domain experts evaluate their understanding of the problem behavior: If their expected and the actual observed problem behaviors match, then this practice reinforces their correct understanding. If they differ, however, domain experts and data scientists typically engage in a mutual process of sensemaking, in which they gradually combine the domain experts’ explanations and the explicit results of the data scientists’ data analyses to extend or revise the domain experts’ understanding of the problem domain. In particular, an important case that can stimulate this process is when data scientists identify outliers in the data, forcing domain experts to explain and reflect on (ab)normal problem behavior that may have remained hidden to them:

“When you’re monitoring an IoT device, you may realize at some point that some operations can be performed better or more efficiently and then you communicate this to [the domain experts], who have often not thought about it this way.” (I24)

“We would preprocess the data and sit with the domain experts to show them what we are able to see in the data, and then ask them for explaining something that seems off. In doing so, we were sometimes able to detect that some of the regular maintenance they did was not useful. So they would learn a lot of things about the practices, something that they do, the norms, unexpectedly.” (I5)

In this way, ML developments can help organizations evaluate, revise, and extend explicit knowledge by combining explicated domain knowledge with data insights.

Internalization: Our experts further described that after being confronted with explicit data insights, domain experts often engage in further investigation to better understand problem instances they observed in the data but were unfamiliar with why they occur and how they behave—especially when the insights reflect anomalies in their domain. In doing so, domain

experts use the available insights to guide their own experimentation with previously unfamiliar problems to improve their understanding of them, unpacking the conditions under which they may occur and the consequences of dealing with them in different ways. In the process, domain experts gain new experience with novel problems observed in others' experiences reflected in the data, which they can use to enrich their own experience and ultimately revise their established practices and develop new approaches to account for the novel problem instances (e.g., considering special cases in their day-to-day decision making that they may have overlooked so far):

"We actually found a few things from the data and some of it just by looking at it, so it was just the analysis—and maybe the visualization, that was also very important. [...] And then, on the [domain expert] side, people said: oh, that's interesting, we have to take a closer look at that." (I4)

"Through this data preparation, we can actually see where things are not working so well, where the data pipelines are incomplete, where business logic is being applied that should not actually be applied, where quality issues occur. That helps us data scientists, but on the other hand [...], of course, we feed this back to the responsible domain experts by reaching out to them like: 'In this step, the process is not working as described. You should investigate in the process why this is the case.'" (I40)

In this way, ML developments can help organizations foster the exploration of unusual experiences to diversify individuals' tacit knowledge.

Socialization: Finally, our experts also described how collective data analysis can foster socialization of experts into each other's experiences and underlying rationales for performing routines and making decisions. The experts stated that, to ground their assumptions well, data scientists often reach out to more than one domain expert to gather multiple opinions about certain ambiguities in the data, with the goal of assessing whether there exists a consensus among domain experts. In doing so, data scientists confront multiple domain experts with data insights at the same time, which often stimulates discussions among the domain experts about the correctness of actual practices that can be observed in the data. This often reveals existing ambiguities and, in particular, inconsistencies between the approaches of domain experts—in the form of so-called "aha moments"—which in turn enable valuable consensus-building processes in ongoing discussions and evaluations (e.g., agreeing on compromises or defining exceptions). Especially when the domain experts come from different departments and thus base their reasoning on different organizational contexts and experiences, the stimulated socialization can help domain experts "look outside their ordinary box" and thereby reflect, revise, and extend their accumulated expertise:

"Or what we also have quite often, which is always very, very exciting, when we have two case workers in a room and we look at an outlier, then one says, 'Yes, yes, it's clear, decision A was correct'. And the other one looks at the one with big eyes and says, 'Nah, I would definitely go for B.'" (I19)

In this way, ML developments can help organizations share and resolve complementary and conflicting tacit knowledge about the collectively faced problem domain.

3.4.2 Problem Solving

The problem solving phase aims to build the ML systems that can best solve the problems represented in the prepared data. To this end, data scientists typically perform the training of ML models to propose candidate solutions, while domain experts are heavily involved in their evaluation to identify the seemingly best candidate. As in the problem finding phase, our interviews suggest that data scientists act as the main stimulus of involved organizational learning processes: Data scientists repeatedly approach domain experts to let them evaluate their derived candidates of ML models, which stimulates different learning processes that help domain experts gradually improve their understanding of adequate solutions, as follows.

Externalization: Since data scientists lack a deep insight into the domain, our experts emphasize that domain experts are typically urged to articulate their perception of an adequate solution in an easy-to-understand manner to data scientists, allowing them to define and document the requirements that an ML system must meet to adequately solve a given problem. In this way, domain experts describe the requirements that seem most relevant to them, while data scientists act as a kind of translator between domain experts and ML algorithms (e.g., translating their understanding of described solution requirements into available data variables and quantifiable metrics that can be interpreted by ML algorithms). In addition, our experts highlight that data scientists expect domain experts to identify and correct anomalous learning by the ML systems when evaluating them, which typically requires them to justify and explain why they believe certain behavior of an ML model is inadequate for solving a problem:

“You start with a very rudimentary model, maybe basic distinctions it can do. For example, for some sort of vision model, we started off with some basic classes [...] which you then show to the business to get their buy-in. Then, how do you need to fine tune it? What kind of granularity do you eventually need? This is where you need to involve the domain experts.” (I22)

“Simply dealing with the situation or with the systems leads to the emergence of knowledge, i.e., a great deal of implicit knowledge is made explicit. [...] Often you have specialist departments that have been doing this and that all their lives and for them it is completely intuitive, so to speak, and they implicitly know how to do it. If you then build software from it or deal with the subject matter, then you have to do it all explicitly and it is discussed. By doing so, the knowledge becomes explicit and then usually ends up in some documentation.” (I29)

In this way, ML developments can help organizations externalize and document tacit knowledge about existing perceptions of adequate problem solutions.

Combination: Our experts described that, when analyzing data to build potential ML models, data scientists typically explore a variety of correlations that they hope may add a relevant piece to solving a given problem. In doing so, data scientists search for and evaluate potential (parts of) solutions that they aim to eventually capture in an ML model. For instance, data scientists of an online retailer might recognize drops in sales during world cup soccer matches and seek to incorporate according information into their ML model for sales prediction. Since these newly discovered correlations are not yet proven to represent causation, however, they must be thoroughly evaluated by domain experts attempting to integrate them with their existing solution descriptions. While this evaluation is often used for improving the input fed into the ML system, it can also yield novel explicit knowledge in the organization when domain experts deem the newly-discovered patterns to be promising (parts of) solutions and seek to combine them with their existing explicit knowledge to revise their standard practices. In particular, when data scientists experiment with ML approaches that may not be suitable for solving the problem (e.g.,

due to being too complex or computationally intensive), they may be able to draw previously-unnoticed connections from the iterative experimentation with the uncovered patterns (e.g., uncovering preferences of certain customer groups). Our experts emphasize that by reflecting on the adequacy of proposed patterns, domain experts often identify complementarities or discrepancies between their own rationales and the proposed solution patterns, allowing them to extend or detail their own solution approaches with additional aspects (e.g., adding new guidelines for handling exceptions or simplifications of their previous approach) or to identify and revise some of their potentially outdated or incorrect rationales:

“For the marketing team, it was also interesting because they knew, for example, that their marketing campaign had an impact on sales, but they didn’t know how much of an impact. They were able to form an expectation that 30% off on shoes has more impact than 10% on hats. But they didn’t know for sure that this was really the case and only anticipated it as one is 10%, the other is 30%, and they guessed that people buy more shoes than hats. But we had the data to try to correlate this.” (I21)

“Just because I’ve seen the three classes in the past doesn’t necessarily mean that a fourth one won’t come along. [...] So if there’s a fourth class that you can clearly see, hopefully you’ll notice that. And if there’s a class that you haven’t even had on your radar yet, you might notice that because something is classified differently than you think. So I think humans have to continue to keep that in mind if that’s relevant to their problem.” (I3)

In this way, ML developments can help organizations evaluate, revise, and extend explicit knowledge by combining existing explicit knowledge with explicated insights into alternative solution approaches.

Internalization: Our experts note that the problem solving phase can also encourage domain experts to adopt solution patterns proposed by the data scientists and experiment with them in their respective domain. For instance, when data scientists, while working on an ML system to predict machine breakdowns, notice that a specific combination of machine parameters significantly heightens the chance of a failure, the engineers (i.e., domain experts) may want to run suitable tests on the machine to better understand the cause of the breakdowns and the conditions of when they may appear. Such experimentation can guide domain experts towards forming new tacit knowledge they can use to refine and extend their own solution practices when handling future problem instances. Thereby, expanding on the insights gained through internalization in the problem finding phase, further domain knowledge may be unearthed through the continued occupation with the problem and corresponding solutions:

“So in the beginning one store is compared with itself. Then, we can look how a particular item behaves in different stores, and then we can check how the whole market of this particular item behaves. And in each of these stages we use an algorithm to highlight what is not inline, and that is then manually reviewed again and, if necessary, either changed or sent back completely.” (I16)

By exposing domain experts to potential data patterns, data scientists can also encourage them to think about the solution from a more pattern-driven perspective. Our experts report that domain experts often begin experimenting with the patterns to understand whether they can serve as (parts of) problem solutions, helping them to enrich their existing reasoning with unorthodox solution perspectives. For example, uncovered patterns can help to think about what other conditions might be considered to trigger a particular solution approach:

“Tinder is a good example. If you do not swipe—can you trap that as a signal? If you didn’t like him or her, how can you take that and feed it back to the system?” (I8)

In this way, ML developments can help organizations create, revise, and extend tacit knowledge by initiating and guiding experimentation with externalized data patterns in new directions.

Socialization: Our interviews reveal that in the process of finding appropriate solutions, data scientists often convene involved domain experts, ideally from different departments/areas, in knowledge sharing sessions, primarily to have domain experts collectively evaluate solution candidates proposed by data scientists. These sessions reflect opportunities for domain experts to meet with domain experts from the same or different areas to discuss how to appropriately solve a common problem. This allows domain experts to learn from each other through emergent discussions in which they share their perspectives on proposed solutions as well as their own best practices, fostering a culture free of silo thinking. For example, if data scientists are developing an ML system to provide sales reps with real-time advice about customers and are sharing emergent patterns that help close sales, the domain experts can share their own experiences of what works and what does not for them from different perspectives in the emergent discussion (e.g., sales reps and product engineers share their impressions of solutions). Our experts emphasize that these sessions can spark interactions among domain experts that can broaden each expert’s perspective by giving them insight into each other’s solution practices, help resolve conflicting approaches, and combine complementary approaches to gradually build consensus:

“In the discussion about what these results mean and what we can do differently with them—‘how can we rearrange our lines? How can we adapt our shift schedules? How can we better deploy our personnel?’—so the knowledge gain takes place in the discourse about the project results, so with the interpretation of the results or the intermediate results. As I said, we follow a phase model with different stop or go decisions along different iteration cycles. The discourse about the intermediate results along this development process is what represents the essential benefit for [the organization]—besides the fact that the final ML model improves something by X percentage points every day in operation afterwards.” (I42)

In this way, ML developments can help organizations share tacit knowledge about solutions across departments and areas of expertise, helping to resolve potential conflicts and make complementary knowledge more accessible.

3.4.3 Solution Implementation

Once adequate ML systems are built, they are implemented into existing processes for productive use. With the emerging interactions between domain experts and developed ML systems when both operate in the same domain, our experts emphasized that ML systems act as the main stimulus for the learning processes of domain experts that occur in the solution implementation phase, enabled by the frequent confrontation of domain experts with the solution approaches generated and applied by the ML systems.

Externalization: Our experts note that the developed ML systems can be used to continuously capture the tacit knowledge of domain experts: By tracking domain experts’ behavior in data (e.g., by recording their past decisions), ML systems can be enabled to continuously observe domain experts’ behavior, which allows the developed ML systems to mimic domain experts’ solution approaches and thereby approximate their applied tacit knowledge with patterns captured in ML models. Our experts underscore that this allows them to store in their ML systems knowledge that was previously held only in the hard-to-access mental models of individuals. Our experts further

emphasize that when organizations use transparent ML models (i.e., by using inherently transparent ML models (e.g., decision trees) or Explainable AI (XAI) methods) in their ML development, the use of the resulting ML systems allows them to externalize and disseminate the captured tacit knowledge of domain experts and thereby make it usable by other experts (e.g., who can now adopt other experts' decision rules that are explicitly described in a decision tree). As a result, our experts underline that ML systems can facilitate the dissemination of existing domain knowledge to other members of the organization and protect this knowledge from loss (e.g., due to personnel turnover):

“In the past, employees had a great deal of knowledge in their heads, which they then somehow managed via Excel lists and knew that they now had to calculate up/down again or something, that would flow into the models. And then, of course, it is no longer so dependent on the employees—that sounds bad now and probably also fuels the fears of one or the other employee—but it is definitely no longer so dependent on the employees.” (I10)

In this way, ML developments can help organizations improve their continuous capture and documentation of domain experts' tacit knowledge by approximating tacit knowledge with transparent ML systems.

Combination: According to our experts, domain experts often learn by comparing and integrating their solution approaches with the ML systems' approach and their provided explanations in their daily tasks (e.g., when faced with a problem instance, they may enrich their solution with the approach outlined and applied by a transparent ML system, such as an ML system that outlines effects of marketing campaign parameters it has observed and considered to derive its sales predictions). Yet, our experts stress that the intensity of this learning process depends on the relevance of the knowledge generated and the level of automation of the given use case. For example, while domain experts may see little value in learning about the rationale behind autonomously tagged images from an ML system, an ML system that predicts crash data for automotive product development may provide highly interesting information on how to stimulate secure new product design ideas. Our interviews show that when domain experts learn from the ML system in this way, they are trying to reconcile the new explicit knowledge created by the ML system with the explicit knowledge (e.g., descriptions of effective product designs) that already exists in their organization:

“When the model makes a decision that the administrator can't understand, there's a lot of skepticism at first. [...] And it is very, very important to discuss this and then, I would say, in 98% of the cases, you generally come to a common denominator where the specialist department also says: 'Yes, okay, you can see it that way, that makes sense.'” (I9)

However, our experts caution that domain experts must remain critical in their evaluation of such confrontations by ML systems. Since ML systems can make erroneous predictions, domain experts must mitigate the risk of spreading erroneous ML-generated knowledge throughout the organization, which could otherwise replace correct human knowledge. Our experts therefore highlight the need for domain experts to continuously verify that ML systems' proposed inputs are indeed based on causality:

“Anyone who calls himself an expert must, of course, be capable of criticism and question these findings. [...] If I have a correlation, and I see it often enough... the more often you see it, the more you can at least believe that it's stable. But that still doesn't say anything about causality, because you should still think about common causes and things like that.” (I3)

In this way, ML developments can help organizations evaluate, revise, and extend explicit knowledge by combining existing explicit domain knowledge with carefully evaluated insights from ML systems.

Internalization: Our experts emphasize that transparent ML systems (e.g., decision trees or applied XAI approaches) can provide domain experts with explicit rationales for how the ML systems would argue for performing certain routines, which can then trigger domain expert experimentation with these rationales. For example, when using ML systems to predict the effectiveness of security measures in a new product design, the ML system can supplement its estimated effectiveness of a given design with specific reasons for how it reached its decision. When the rationales differ from the domain experts' own solution approaches, our experts find that the domain experts do not just naively adopt the rationales as new approaches, but typically begin to experiment with the rationales, testing their actual performance in different contexts (e.g., trying rules that the ML system would use in certain scenarios instead of their own approaches). In this way, the domain experts can gain new experiences guided by the ML systems' approaches, which enables them to revise and extend their own solution approaches:

"This happens very often. Especially through the description of causal relationships [through XAI approaches], anomalies or peculiarities are revealed that you had not considered before. This is very, very helpful and is really my be-all and end-all: If you can learn something new from it, then this can already bring you a lot further, because you can try this knowledge as an engineer in future cases and projects." (I25)

Yet, our experts caution that this process can also become vicious if the ML systems' erroneous reasoning creeps slowly and unnoticed into the domain experts' solutions as they interact with the ML systems and develop confidence in the veracity of the ML systems' described rationales through repeated interactions rather than through reasoned justification based on comparison with their own domain knowledge:

"The interaction is actually always a problem for me, because it can happen that people suddenly rest on it and say: 'Yes, that thing is always right anyway, so I don't do my job anymore.' There's the situation where you kind of say: 'Okay, I just trust that thing and I always do what it says.' [...] So it also works the other way around, that the model influences the human being. That also works. And to find out this interaction is very challenging. This is very, very problematic." (I19)

In this way, organizations can acquire new tacit knowledge as human experts (un)intentionally collect new experiences guided by the rationales offered by transparent ML systems.

Socialization: Our experts underscore the importance of domain experts also learning about the tacit knowledge stored in non-transparent ML systems by observing the ML systems' behavior. By observing how ML systems make decisions, domain experts can, over time, approximate to some extent the rationale used by the ML systems, thereby creating their own tacit knowledge from the experienced task performances by the ML systems, and thus gradually become socialized into the ML systems' problem-solving behavior. For example, doctors who use a non-transparent ML system that suggests diagnoses for a while may learn to anticipate when and how the ML system is likely to reach certain conclusions and incorporate these heuristics into their own decision-making rationale. Our experts believe that this process is important, not only because it can help ensure the quality of the ML systems being used (as domain experts typically also judge the correctness of the approximated rationale of the ML system), but also because it requires domain experts to continually reflect on and question the correctness and currency of their own

rationales and to revise them when necessary. Once domain experts identify a mismatch between an ML system's rationale and their own, our experts emphasize that domain experts then often engage in a collective sensemaking process in which domain experts reach out to other domain experts and data scientists to understand why there are conflicts between the ML systems' rationale and their own and how to resolve them, with the goal of reaching a conclusion about which of the two rationales is better suited as an adequate solution or how to integrate both rationales to create a new superior solution, which typically involves intensive discussions between data scientists and domain experts and possible reconfigurations of affected ML systems:

"The old veterans sit there, so to speak, who know exactly what is important—or think they know what is important—and they have to say: 'Yes, it makes sense if the model says that if I now allow ten millimeters more forward displacement, it will have a positive effect on my bus acceleration'. So in the end, SHAP and LIME [i.e., XAI approaches] are used to discuss whether this prediction makes sense or not, whether the model itself makes sense or not. So that is then discussed. Together, locally." (I25)

"We do that [i.e., observing ML systems' behavior], but we need time to do it, because to find these patterns is not so easy. You have to use these models at least one or two years in production to know some patterns, so that you can say: 'Ok, this model works like this and we found this pattern.'" (I14)

"Before we fully changed the approach to a machine learning approach, there were two or three months where the [domain experts] would still predict the sales and the machine learning system would also predict the sales—and then there was a comparison. So, they would check, because they knew what they would come up with and they would maybe see if they are not right or that the model is not working well for this kind of campaign that we have and we would discuss it with them." (I21)

In this way, ML developments can help organizations share and resolve complementary and conflicting tacit knowledge between their human experts and ML systems.

In sum, our experts reported a variety of learning processes stimulated by the confrontation of domain experts with data scientists and ML systems. These processes can provide several opportunities to contribute to the revision of existing and the creation of new explicit and tacit domain knowledge in the organization. With the need for domain experts to reflect on their own understanding of problems and adequate solutions, and to engage in mutual learning with others with whom they may not typically interact, the three phases of ML developments seem to provide powerful stimuli for the four key organizational learning processes. In particular, by confronting domain experts with the need to explain their understanding to non-experts (i.e., data scientists), to express and evaluate it in terms of clear data variables and patterns, and to resolve conflicts and experiment with the understandings and observations of others (e.g., other domain experts, data scientists, and even ML systems), ML developments can provide unique encounters for stimulating organizational learning.

3.5 Discussion

Organizational learning is a crucial process that lies at the heart of organizational behavior and is known to fundamentally control an organization's performance (e.g., Argote et al., 2021; March, 1991). As organizations should therefore be careful to optimize their organizational learning processes, decades of research have analyzed how organizations can increase their organizational learning effectiveness (e.g., Argote & Miron-Spektor, 2011; Huber, 1991). More recently, research

has recognized that ML systems can participate as a new type of organizational learner alongside humans, which can strongly influence organizational learning (e.g., Argote et al., 2021; Sturm et al., 2021b). So far, a handful of studies have explored how productive ML systems can contribute knowledge to organizational learning through learning dynamics that may emerge between humans and ML systems (i.e., Balasubramanian et al., 2022; Lyytinen et al., 2021; Ransbotham et al., 2020; Seidel et al., 2019; Sturm et al., 2021b). With our study, we aimed to explore not only how this human-AI interplay can change organizational learning, but how the preceding and continuous ML developments can stimulate different valuable types of learning processes among the human members of organizations, even when ML systems are not yet being used productively.

To this end, our study offers several theoretical contributions. First, we add to the emerging discussion on the impact of ML on organizational learning (e.g., Argote et al., 2021; Seidel et al., 2019). While existing work primarily focuses on the influence of productively deployed ML systems (e.g., Balasubramanian et al., 2022; Sturm et al., 2021b), our study shows that understanding the impact of ML on organizational learning can benefit from a more holistic perspective that considers not only the direct impact of ML systems by observing their mutual influence with human learners, but also takes into account the interactions between humans that are required to realize ML systems in the first place. In this sense, our study provides initial empirical evidence on how human-centered organizational learning processes unfold within ML developments, which we hope can provide a theoretical foundation for further research.

Second, our study demonstrates that ML developments can, indeed, serve as an important mechanism for stimulating organizational learning. The sub-processes uncovered by our study show how ML development activities, even if they do not result in productive ML systems, can add value by initiating different learning processes for domain experts. Our interviews highlight that ML developments can thus act as a new kind of “field of interaction” (i.e., as proposed by Nonaka (1994) as an important learning mechanism driven by encounters) where domain experts, lay people (i.e., data scientists), and ML systems can meet to collectively share, reflect, and revise organizations’ domain knowledge. Since the stimulated learning processes do not only reside in ML developments, but domain experts are likely to disseminate their new knowledge to other parts of the organization when they return to their domain-specific departments (e.g., traders involved in developing a trading ML system applying their newly acquired knowledge in their trading activities and sharing it with other traders), our study suggests that the consequences of the identified learning processes may have a large impact on organizations’ knowledge stock in the long run. However, while our study confirms the existence of these processes, their actual consequences for organizational learning remain unknown and deserve further attention in future studies. Moreover, uncovering the conditions that benefit or harm the identified learning processes may also provide much-needed insights into how to effectively coordinate ML developments and their consequences for improving organizational learning. To this end, while our experts emphasized the positive side of ML developments as learning stimuli, it may be helpful to further understand if and how they can also turn vicious to (other) organizational learning processes, and act as an alternative, obstacle, or complement to other stimulating mechanisms (e.g., rotational programs, interorganizational partnerships; e.g., Argote et al., 2021; Nonaka, 1994).

Third, the uncovered processes highlight that ML developments involve two different types of stimuli depending on the ML development phase: the data scientists and the developed ML systems. By confronting domain experts with new insights, emerging ambiguities, needed explanations, and their generated hypotheses and models, the interaction of domain experts with

data scientists or ML systems provides opportunities for domain experts to substantially reflect on, revise, and articulate their knowledge. Our interviews show that these processes can help improve the availability and quality of organizations' domain knowledge (e.g., by forcing the articulation of tacit knowledge, triggering the revision of outdated knowledge). This finding reflects a valuable observation, because a well-known major problem in organizational learning research is the scarcity of activities that enable reflection and substantial revision of conventional knowledge (e.g., the tendency to favor exploitation of known over exploration of unknown solutions (Levinthal & March, 1993) and lack of questioning the adequacy of established learning conditions and goals through "double-loop learning" (Argyris, 1976)). ML developments may thus serve as a useful new mechanism that organizations can deliberately use to promote such reflective and revisionist activities, thereby helping to liberate organizations from suboptimal or outdated conventional knowledge (see, e.g., Argote et al., 2021; Argyris, 1976; Levinthal & March, 1993). While our study uncovers ML developments as a fruitful mechanism, research may greatly benefit from further studies on how organizations can effectively coordinate the interaction between domain experts, data scientists, and ML systems to integrate ML developments as a strategic element that fosters organization-wide domain knowledge.

Fourth, our study shows how ML systems can affect different types of knowledge (i.e., explicit and tacit knowledge). In particular, the role of ML systems as carriers of tacit knowledge has recently received increasing attention in research, stressing the potential role of ML systems as tacit knowledge repositories that can help prevent certain knowledge losses (e.g., Hadjimichael & Tsoukas, 2019; Lebovitz et al., 2021). In line with this research, our study confirms that ML systems are used as tacit knowledge repositories across different domains and industries. Interestingly, our study shows that ML systems can further serve as a powerful tool for externalizing tacit knowledge (e.g., a decision tree that mimics and describes experts' approaches) and as a stimulus for human learning of tacit knowledge. ML developments can thus provide an important mechanism to facilitate the management of tacit knowledge, which is known to be a difficult and tedious endeavor (e.g., Argote et al., 2021; Nonaka, 1994). Here, more research is needed to understand how organizations can effectively use ML systems to convert between tacit and explicit knowledge while avoiding losses and biases in the translation between humans' and/or ML systems' knowledge.

From a practical perspective, our study emphasizes that organizations should be cautious about reducing investment in ML development, even if they repeatedly fail to develop ML systems for productive use. Since most "unsuccessful" ML developments still involve the first two phases of ML development (i.e., problem finding and solving), organizations may otherwise miss out on significant benefits that can result from the learning processes involved—and thus miss a powerful driver that can improve their long-term knowledge needed to differentiate themselves from their competitors. In addition, our interviews emphasize that organizations should ensure that domain experts are well integrated into ML developments. Only when domain experts have the opportunity to interact with and be confronted by the insights of data scientists and ML systems can domain experts bring new insights back to their domains and infuse the new knowledge into further organizational learning processes, thereby spreading gained knowledge throughout the organization. Finally, organizations should be keen to allocate additional time and resources to ML developments to allow space for domain experts' analysis of insights and discussions that may not contribute to the development of the planned ML systems, but are focused on exploring potential process failures and resolving identified conflicts and ambiguities in existing knowledge. Otherwise, organizations may stifle valuable learning processes that may

ultimately prove to be more than just a nice-to-have byproduct, but an essential stimulus for vital performance gains in the long run.

Of course, our study has several limitations. First, although our respondents cover a wide range of experiences across industries, roles, and ML use cases, there may still be data biases that we could not completely eliminate. Further qualitative and quantitative studies in different contexts may therefore help to uncover such biases and validate the applicability of the identified processes. Second, because we based our analyses exclusively on participants' post-hoc descriptions, our findings can only consider the details that were remembered and deemed relevant by our participants. In particular, in-depth case studies of ML developments could help to provide complementary insights into the learning processes uncovered in our study that we were unable to observe in the interviews. Finally, while we attempted to capture experiences with different types of ML systems across different industries, domains, and roles of participants, the processes identified may vary depending on the ML system under development. Here, future research could help validate and contextualize the learning processes for ML developments of specific types of ML systems.

Our study is only a first step in understanding the impact of ML developments on the crucial processes of organizational learning. As more research is needed to better understand how ML developments affect organizational learning and how ML developments can be used as a strategic means to improve organizational performance in the long run, we hope that our study itself can serve as a fruitful theoretical "stimulus" for future research to help rethink organizational learning theory in the era of AI.

Acknowledgment

This research was funded by the German Federal Ministry of Education and Research (BMBF) within the "Innovations for Tomorrow's Production, Services, and Work" program (funding number 02L19C150) coordinated by the Project Management Agency Karlsruhe (PTKA).

4 Research Paper B: Creating Value Despite Machine-Learning-Induced Tensions

Title: An Ambidextrous Perspective on Machine Learning Development and Operation: The Nexus of Organizational Structure, Tensions, and Tactics

Authors: Vetter, Oliver A.; Pumplun, Luisa; Koppe, Timo

Published in: Proceedings of the 27th Pacific Asia Conference on Information Systems (PACIS 2023)

Abstract

Organizations from all industries have recently begun to develop and operate machine learning (ML) systems. While ML promises to improve an organization's effectiveness and efficiency, developing and operating ML systems remains challenging as these systems differ significantly from traditional software and require novel work practices that run counter to existing business processes. These conflicting demands create tension in the organization as resources to develop and operate ML systems are limited. Organizations thus seek to leverage scarce resources by employing a range of organizational structures and tailored tactics. To explore the interplay of organizational structures, tensions, and tactics, we conducted an explorative expert interview study informed by computational grounded theory methodology. We took an ambidextrous perspective to identify four central tensions and associated tactics employed within given organizational structures. Further, we found that organizations are moving from centralized and decentralized structures to hybrid ones to enable effective ML development and operation.

Keywords: Machine learning, organizational structure, ambidexterity

4.1 Introduction

Machine learning (ML) has been on the rise in recent years and has become a buzzword in almost every industry. By automatically monitoring machine systems, detecting cyber-attacks, or intelligently handling customer inquiries, among other things, it promises to make organizations more effective, efficient, and potentially provide them with a competitive edge (Brynjolfsson & McAfee, 2017). As a result, many organizations feel pressured to develop and subsequently operate systems that are based on ML (e.g., Baier et al., 2019; Pumplun et al., 2019). At the heart of such ML systems are applicable models that are trained based on incoming data and learning algorithms (Mitchell, 1997; Russell & Norvig, 2016). In this regard, the development and operation process of ML systems differs significantly from that of conventional software systems, as it requires new competencies, for example, in the field of data science or data engineering (e.g., competencies on data cleansing, learning algorithm selection, data pipeline building; Amershi et al., 2019). Furthermore, the data-based learning approach implies that the development and commissioning of ML systems often has a highly experimental character, for example, in preparing the data or fine-tuning the model's hyperparameters (Amershi et al., 2019; Choudhury et al., 2021).

At the same time, with the increasing relevance of ML systems, it is no longer just large software providers seeking to develop and operate such systems; organizations from a wide range of industries have begun to explore this new technology (Gerbert et al., 2020). Common to most of these organizations is that their experience and existing resources for developing and operating ML systems are scarce, and that previous working practices are inconsistent with the experimental nature of ML (e.g., Baier et al., 2019; Choudhury et al., 2021; Gerbert et al., 2020; Pumplun et al., 2019). In particular, the development and operation of ML systems require a shift toward more flexible, faster, and exploratory processes that may not align with existing conventional business processes (e.g., Gerbert et al., 2020; Pumplun et al., 2019). In other words, ML development requires organizations to engage in exploration activities (see March 1991). Conversely, conventional business processes inherently favor exploitation (e.g., Uotila et al., 2008), for example, because it yields positive returns in a shorter period of time than exploration (e.g., Levinthal & March, 1993). As a result, there may be strong tensions as not only financial, human, and technical resources need to be allocated between the old and new workflows, but also changes in the way of working have to be implemented. A difficult task, as the organization and its members are often driven by day-to-day business and need to meet operational goals as well.

As of now, most companies are still in the early stages of developing and operating ML systems and are struggling to resolve the tensions involved (van der Meulen & McCall, 2018; Ransbotham et al., 2019). Therefore, many organizations have been ineffective and inefficient in their ML development and operation to date, leading to significant delays in ML system deployment or even ML project failures (van der Meulen & McCall, 2018). Previous research on other technological innovations shows that organizations need to choose an organizational structure in the early stages of developing and operating an innovative technology (e.g., Haffke et al., 2017a; Sharma et al., 2014). According to Günther et al. (2017), an organizational structure is the system by which technical and human resources are organized into departments. In this context, there are two essential structural means to organize the organization and its units: The centralized structure, in which a competence center is entrusted with the development and commissioning of innovative technologies, and the decentralized structure, in which existing specialist departments are charged with these tasks (Haffke et al., 2017a; Sharma et al., 2014). Thus far, however, it is not known whether these structures are equally useful in the development and operation of ML

systems (Ransbotham et al., 2017). Moreover, the chosen structure can lead not only to a relaxation but also to an aggravation of tensions, as it can bring its own disadvantages. For example, employees in decentralized structures may be overwhelmed by the simultaneous management of the day-to-day business and new technological innovations (e.g., Davenport, 2018; Gerbert et al., 2020; Sharma et al., 2014). Organizations must therefore give additional thought to how they can resolve these potentially increased tensions in the respective organizational structures.

To date, little research in the field of ML has examined the tensions that arise in the development and operation of ML systems, how the choice of organizational structure affects such tensions, and what tactics organizations can pursue in their respective organizational structures to ensure successful ML projects. Few available studies deal specifically with the challenges encountered in developing and operating ML systems (e.g., Baier et al., 2019; Marabelli et al., 2021), but not the interplay of these challenges with organizational structures and possible tactics for addressing them. As a result, there is a lack of understanding of how organizational structures and tactics interact with each other and help resolve emerging tensions within the ML development and operation process. In order to provide organizations with helpful insights on how to ensure more successful ML projects, our study investigates the following research questions:

(1) What organizational tensions arise during the development and operation of ML systems and how do they differ within different organizational structures?

(2) How do organizations cope with these tensions effectively through tactics?

To answer these questions, we conducted an explorative interview study inspired by grounded theory methodology. To analyze the interviews, we adopted an ambidextrous perspective, whereby ambidexterity is defined as an organization's capability to manage opposing demands (Raisch & Birkinshaw, 2008). As the demand for fast and flexible exploratory processes in ML development stands opposed to the tendency of conventional business processes to focus on exploitative activities (e.g., Gerbert et al., 2020; Pumplun et al., 2019), our ambidextrous perspective provides helpful guidance to examine possible tensions and tactics within the ML development and operation process. This way, our study contributes to the success of future ML projects by providing insight into how tactics can be precisely aligned with tensions in organizational structures.

4.2 Theoretical Background

We first provide an overview of the current state of research concerning the development of ML systems from a technical and strategic perspective. Next, known organizational structures are described, and organizational ambidexterity is introduced as a meta-theoretical lens for our interview analysis.

4.2.1 ML Development and Operation

As a variety of organizations are developing and operating ML systems today (Gerbert et al., 2020), it is important to understand the unique characteristics of ML for organizations to understand the tensions that arise. ML systems constitute a possible instance of artificial intelligence; that is, a field of research concerned with the development of intelligent machines (McCarthy, 2007). ML systems enable processes to be fulfilled that were otherwise reserved for humans, such as diagnosing illnesses, answering customer inquiries, or monitoring the

functioning of machines (e.g., Brynjolfsson & McAfee, 2017). In doing so, ML systems exhibit some characteristics that distinguish them from other forms of digital technologies. Berente et al. (2021) subsume these differentiating characteristics under the facets of *learning*, *inscrutability*, and *autonomy*, each of which uniquely complicates the implementation of ML systems in organizations. Technically, ML systems differ from conventional systems in that they are based on applicable models derived from data (Mitchell, 1997; Russell & Norvig, 2016). *Learning*, therefore, describes that these models are not programmed using explicit if-then rules but are trained using data and statistical learning algorithms (Amershi et al., 2019). When using ML systems, the models are, in turn, applied to new data, e.g., to perform classification or clustering (Mitchell, 1997; Russell & Norvig, 2016). As ML technologies are capable of processing large amounts of data while considering complex interactions between variables, they may uncover novel knowledge in the data that has previously been missed by domain experts (Agarwal & Dhar, 2014; Müller et al., 2016). However, ML systems are incapable of reacting to environment states, which they have not been trained with (Dennett, 2006), and require human guidance for these interpretations (Seidel et al., 2020). Organizations must, therefore, carefully evaluate whether to let ML systems learn independently to glean novel insights in a truly exploratory manner or to involve domain experts for the ML system to remain more relevant to the domain in which it is embedded (van den Broek et al., 2021). The *inscrutability* of ML systems is caused by the increasingly complex statistical learning algorithms and ever-larger data sets that are being used to train models (e.g., neural networks), making them increasingly complicated and incomprehensible to humans (Meske et al., 2022; Rudin, 2019). This is particularly problematic since these systems can produce unexpected outcomes (Benbya et al., 2020a), which can lead to unanticipated consequences within organizations (Benbya et al., 2021). Furthermore, employees who use ML systems in their decision-making process may even experience an increase in uncertainty because they lack explanations of the underlying reasoning for the ML output (Lebovitz et al., 2022). Implementing ML systems into organizations must, therefore, not necessarily lead to increases in efficiency, and corresponding projects must be carefully evaluated and managed. Lastly, *autonomy* denotes the capability of ML systems to perform actions without human intervention or human knowledge (Baird & Maruping, 2021), which allows them to complement, constrain, or replace humans at work (Murray et al., 2021). As ML systems are implemented in organizations, this shift of agency toward the machines can give rise to fears among ML-developing employees that their occupation will become obsolete (Vaast & Pinsonneault, 2021).

The striking differences of these characteristics of ML systems and their impact on implementing organizations from other digital technologies suggest that developing and operating ML systems also poses unique managerial challenges. Indeed, to customize and reuse the models forming the core of ML systems, new skills need to be developed in the organization that differ from capabilities in prior software development (Amershi et al., 2019). In particular, integrating ML systems into the existing software landscape is complex due to their unpredictable, possibly erroneous outputs (Benbya et al., 2020a; Amershi et al., 2019). The development and operation of ML systems as a whole is strongly characterized by an experimental character. It is difficult to predict in advance how ML systems will work and what steps are required to train a high-performance model (Amershi et al., 2019; Choudhury et al., 2021). A general ML development process is as follows (e.g., Shearer, 2000; Amershi et al., 2019): First of all, the *business and data* must be thoroughly understood, followed by the *preparation of the data* used to train the model. The third step is the *creation of the model* itself and its subsequent *evaluation*. The actual development concludes with the *deployment* step. After deployment, the ML system is operated by observing and *maintaining* the system behavior, possibly triggering further development

cycles. Therefore, the operation of ML systems is part of the development process as well. Since all steps are iterative and typically performed multiple times, the clear boundaries between development and operations are becoming increasingly blurred (Amershi et al., 2019; Studer et al., 2021), again highlighting the need for flexible exploratory processes.

4.2.2 *Potential Organizational Structures*

ML systems are a strategic technological innovation that requires a transformation of organizational resources, their allocation, and ultimately that of the entire organizational structure (Pumplun et al., 2019). At a high level of abstraction, initial research is already addressing organizational alignment, strategy, and deployment of ML systems (e.g., Baier et al., 2019; Jöhnk et al., 2020; Pumplun et al., 2019). However, a detailed consideration of potential organizational structures, their perils, and promises has so far been lacking. Basically, when developing and operating ML, organizations can choose whether they prefer the synergistic benefits of centralized organizational ML structures or whether an organically growing, decentralized structure is advantageous to them (Tallon et al., 2013). Some insights into these organizational design decisions can be derived from the research field of bimodal information technology (IT), which deals with the organizational structures that can be implemented to conduct IT projects. Organizations can either set up a divisionally separated bimodal IT for new projects or might choose to entrust existing IT or business units with the IT development (Haffke et al., 2017a). This discussion also occurs in the development of data-driven technologies such as analytics and the dissemination of their results throughout the organization. Sharma et al. (2014), for instance, call to investigate how existing organizational structures affect organizational members' ability to generate data-based insights. Indeed, the optimal structure for developing analytics and disseminating its results is not obvious, as "multiple actors from different parts of the organi[z]ation" (Sharma et al., 2014, p. 435) need to collaborate intensely. Rigid organizational structures can constrain the required cross-disciplinary collaboration, which is why Günther et al. (2017) request for additional research examining how businesses can create flexible organizational structures to foster the required cooperation. Compared to analytics, when integrating ML systems, organizations face the additional challenge of not only having to distribute individual analytics results throughout the organization but systems that act and learn in an automated manner and must be integrated into the existing system landscape and maintained over time. For the development and operation of ML systems, it is thus still unclear which organizational structure is advantageous and how tensions within the chosen structure can be relieved through meaningful tactics. One perspective that can help better understand the tensions and approaches to resolving them in organizations is organizational ambidexterity, which is discussed in more detail below.

4.2.3 *Organizational Ambidexterity as a Meta-Theoretical Lens*

Demands within an organization constantly contradict each other to a certain extent, so tensions arise and trade-offs must be made (Gibson & Birkinshaw, 2004; Raisch & Birkinshaw, 2008). If one demand is fulfilled, the other may be neglected, complicating the simultaneous fulfillment of needs. This conflict is addressed under the theoretical concept of ambidexterity, which offers an approach to managing such tensions (Gibson & Birkinshaw, 2004; Raisch & Birkinshaw, 2008). Organizational ambidexterity is defined as "an organization's ability to be aligned and efficient in its management of today's business demands while simultaneously being adaptive to changes in the environment" (Raisch & Birkinshaw, 2008, p. 375). In this sense, only organizations that excel

in both areas can balance the tensions that arise between the demands of existing business and innovation (March, 1991). Although the tensions cannot be wholly eliminated, organizations can manage conflicting demands and reduce the tensions significantly by employing effective tactics to find a trade-off (Gibson & Birkinshaw, 2004; March, 1991; Raisch & Birkinshaw, 2008). As a result, organizations that embrace ambidexterity as a means can achieve better firm performance (Raisch & Birkinshaw, 2008).

Ambidexterity has been investigated and specified in multiple studies. For instance, the conflicting demands of alignment and adaptability (Gibson & Birkinshaw, 2004), rigor and agility (Lee et al., 2010), or exploration and exploitation (March, 1991) were examined. For our study, the concept of exploration versus exploitation is particularly relevant, as the need for exploration in ML development and operation (e.g., Gerbert et al., 2020; Pumplun et al., 2019) conflicts with the tendency of conventional business processes toward exploitation (e.g., Uotila et al., 2008). While exploitation deals with activities such as “refinement, choice, production, efficiency, selection, implementation, execution,” exploration refers to terms such as “search, variation, risk-taking, experimentation, play, flexibility, discovery, innovation” (March, 1991, p. 71). Exploration and exploitation thus form a bipolar relationship, however, they are not mutually exclusive but rather interdependent opposites in a duality (Putnam et al., 2016; Farjoun, 2010). Organizations looking to balance the trade-off may take various approaches, many of which can be subsumed under two ideal types of ambidexterity: structural and contextual (Raisch & Birkinshaw, 2008; O’Reilly & Tushman, 2013; Ossenbrink et al., 2019). Thereby, structural ambidexterity describes the simultaneous pursuit of both exploration and exploitation through separate structural units with different competencies that are held together by the common strategy of the organization (O’Reilly & Tushman, 2008). Contextual ambidexterity then denotes approaches that create a fostering environment for individual employees to decide for themselves when and how to allocate time and resources toward exploration or exploitation (Gibson & Birkinshaw, 2004). In practice, organizations choose which of the types to pursue based on a variety of reasons, for instance, the amount of identifiable new business opportunities or the required capabilities for seizing them (Ossenbrink et al., 2019).

In the case of ML development and operations, the demand of the existing business is focused on efficiency and thus exploitation, while data-based learning requires experimental, novel, and exploratory processes, which can create strong tensions in the organization (Davenport, 2018). Organizational ambidexterity has played a role in many different streams of literature, including technological innovation, organizational adaptation, or organizational design (Raisch & Birkinshaw, 2008). Moreover, it has been adopted as a theoretical concept in studies on applying data-based technologies (e.g., Kowalczyk & Buxmann, 2015). This led us to apply organizational ambidexterity as a meta-theoretical lens for our study, especially since it dovetails with our empirical observations: *“You have to choose a side. In that sense, of course, it’s a trade-off”* (E09).

4.3 Methodology

To date, the existing literature offers little insight into answering the research questions of what tensions arise in the development and operation of ML and how these can be resolved within different organizational structures. Therefore, we conducted an explorative expert interview study (Bogner et al., 2009), which is a suitable methodology to address under-researched problems (Myers & Newman, 2007; Corbin & Strauss, 2015). Furthermore, to ensure rigor and structure in making sense of our data, we draw from grounded theory methodology to guide our interview study and apply principles of the methodology for our data collection and analysis

(Urquhart et al., 2010; Seidel & Urquhart, 2013). The application of grounded theory has gained popularity in recent years, especially to explore the impact of technological change (Birks et al., 2013; Wiesche et al., 2017) as in the field of ML, e.g., to study ML adoption in healthcare (Lebovitz, 2019). Grounded theory is not designed to test existing theories but aims to derive new theories from research data (Birks et al., 2013; Urquhart & Fernández, 2013).

In general, grounded theory studies include several steps that enable a structured analysis of the data (Birks et al., 2013; Wiesche et al., 2017). After the subject and the field of research upon which the study is to be conducted are determined, data collection and analysis are performed. These follow an iterative process through which insights are gradually gained and expanded. In this regard, the selection of interview participants is not done in advance but is performed successively based on the preceding data and insights gained (i.e., theoretical sampling; Birks et al., 2013; Glaser & Strauss, 1967). The data collection is terminated as soon as theoretical saturation is reached (i.e., no further insights emerge from additional data). In order to establish the correspondence between reality, data, and the research field, the resulting grounded theory is integrated against the background of existing theoretical concepts (Glaser, 1978).

4.3.1 *Data Collection*

In-depth interviews are one of the most prevalent data sources in grounded theory studies (Charmaz & Belgrave, 2012). Since they offer insights into complex real-world contexts, we conducted interviews with experts, most of whom work at the management level and are responsible for integrating ML systems into the organization. All experts interviewed have profound experience in the development and operation of ML systems in a large number of projects and can thus report from a wealth of experience. We approached the experts through professional online networks and ongoing industrial collaborations. To counteract potential selection bias and in line with grounded theory methodology, we applied a theoretical sampling approach to decide on the interview participants successively based on prior data (Birks et al., 2013; Glaser & Strauss, 1967). With this in mind, we interviewed experts from various organizations of different sizes and industries that follow different organizational structures for developing and operating ML systems. We laid a focus on acquiring experts from established, larger organizations, as initial interviews showed that it is particularly these organizations that have the personnel resources and structural granularity to actively address and openly decide on the presented conflict between centralization and decentralization. This also includes experts from consulting firms who have insights into a wide range of client companies. We continually added more questions to our interview guide to cover topics that emerged during the research process. For instance, the first interviews focused particularly on the ML development and operation process and associated challenges, while later interviews also discussed solutions to the identified tensions. The sampling process was terminated once we reached theoretical saturation and no new themes or relationships emerged within the newly collected interview data. We conducted two rounds of interviews, one from the end of 2019 to the first quarter of 2020 and another in the first quarter of 2021. In total, we conducted sixteen in-depth interviews (see Table 6), which were recorded and transcribed to allow for detailed analysis. The interviews lasted an average of 55 minutes, resulting in 875 minutes of interview data.

Table 6. Study Participants Overview.

ID	Position	Experience	Organizational Size	Industry
E01	Senior Manager ML	8 yr.	Very large	Consulting
E02	Managing Director ML	23 yr.	Very large	Consulting
E03	Consultant ML	4 yr.	Very large	Software
E04	ML Adoption Lead	20 yr.	Very large	Software
E05	Product Manager ML	28 yr.	Very large	Software
E06	Global Head of ML	13 yr.	Very large	Automotive
E07	Product Manager ML	10 yr.	Very large	Automotive
E08	Head of ML	17 yr.	Very large	Transport
E09	Head of ML	6 yr.	Large	Software
E10	Product Manager ML	7 yr.	Large	Software
E11	Manager ML	11 yr.	Medium	Consulting
E12	Head of ML	6 yr.	Small	Consulting
E13	Head of ML	7 yr.	Very large	Aviation
E14	ML Specialist	3 yr.	Very large	Automotive
E15	Digital Transformation Manager	23 yr.	Very large	Insurance
E16	Head of ML	21 yr.	Very large	Insurance

Note. Very large: empl. > 1000/ rev. > 500M €, large: empl. > 250/ rev. > 50M €, medium: empl. > 50/ rev. > 10M €, small: empl. ≤ 50/ rev. ≤ 10M €.

In each interview, we used a semi-structured interview guide that included five sections: First, the purpose of the study and the interview process were explained to the participants to create a common understanding of the phenomena investigated. Second, we included questions about the ML development and operation process as well as the challenges encountered. Third, we asked about the organizational structure used in each case, whether there were alternatives to the chosen organizational structures, and what advantages and disadvantages they would have had. Fourth, we discussed possible alternative courses of action, solutions to the identified tensions, and asked about possible distinctions between tensions and tactics in different organizational structures. Lastly, we asked about success factors and potential competitive advantages from different approaches.

4.3.2 Manual and ML-Based Data Analysis

Throughout the whole analysis process of our explorative interview study, we followed the principles of grounded theory methodology to ensure a reproducible, structured analysis. In grounded theory studies, insights emerge from the codes that researchers derive from the data and are progressively synthesized and abstracted to a more conceptual level (Glaser, 1978). More recently, information systems (IS) research has called for supplementing this coding process with ML-based computational analysis to provide a highly exploratory perspective on the data that is largely independent of the experiences and attitudes of the researchers performing the coding (Baumer et al., 2017; Nelson, 2020). This endeavor is summarized under the term *computational grounded theory*, and allows to combine the high processing power of computer programs and the slower yet more comprehensive capabilities of humans (Baumer et al., 2017; Nelson, 2020). In

order to capture ideas about themes and relationships within the data that emerge during the coding process, memo writing is conducted in parallel (Glaser, 1978).

With this in mind, we based our coding process on the procedure suggested by Glaser (1978), iteratively jumping back and forth between the different coding stages. To perform the coding, we used the NVivo 12 software. Furthermore, we engaged in comprehensive memo writing throughout the whole process to capture ideas (e.g., in the form of various diagrams). To achieve rigor, we employed multi-researcher triangulation and discussed the ideas intensively among the authors and two further IS researchers (Carter et al., 2014).

In the first step, we followed a mixture of human-led open coding and a computational approach to open coding. We began by going through the interview data line by line, assigning codes at the level of whole sentences, parts of sentences, and words to identify topics of interest (Glaser, 1978). In particular, we took a content-based approach and sought descriptive codes for the respective units. In addition to the manual *open coding*, we used ML-based natural language processing techniques to detach from existing mental models and analyze the data more exploratory in the sense of computational grounded theory (Nelson, 2020). The goal of such computer-assisted procedures is to condense unstructured interview data into easily interpretable lists of coherent words, making it easier to identify patterns in the text. Moreover, this procedure counteracts biases and inconsistencies in human interpretation that can typically occur when analyzing large amounts of text data (Nelson, 2020). We loaded the transcripts into a Jupyter Notebook in a Pandas DataFrame. We preprocessed the interview data by removing stop words and expressions that occur frequently but have no informative value (e.g., position designations) and stemmed the words using the snowball stemmer. Next, we used the Python scikit-learn library to implement two common approaches to analyzing text data: topic modeling and clustering. These approaches are used to detect topics in text data by evaluating the co-occurrence of words (Nelson, 2020). We performed topic modeling by implementing a Latent Dirichlet Allocation model and a non-negative Matrix Factorization with 10 and 15 topics, respectively. Moreover, we clustered the data using a k-Means and a Mini Batch k-Means algorithm, both of which were run with 10 and 15 clusters to be formed. The resulting lists of relevant words were subsequently reviewed and interpreted by the authors. This exploratory approach helped us identify topics in the qualitative data that would otherwise have been overlooked. For example, one of the ideas received through the non-negative Matrix Factorization (with ten topics) was characterized by the words: *engineering, software, data, and science*. This combination of words has led us to consider the necessary IT development disciplines (i.e., software engineering, data science) as a code in further analysis (see Tension 3).

In the second step, we performed *selective coding*, in which the respective codes from open coding were grouped and traced back to common basic phenomena (Glaser, 1978). This step also included the codes achieved through the computer-assisted analysis. While creating the overarching categories, we considered *ambidexterity theory* as a meta-theoretical perspective. In grounded theory methodology, general theories can be identified as a meta-theoretical lens that – while not providing an exhaustive foundation for the study – can enrich understanding about the data (Glaser & Strauss, 1967) and help place findings in the context of existing theories (Urquhart et al., 2010). Following the principle of ambidexterity, we compared the codes to identify conflicting demands stemming from the integration of ML systems and conventional business processes and reveal associated tensions. In addition, we captured organizational structures that organizations follow to manage the respective tensions, as well as exploratory and

exploitative tactics beyond that. The resulting categories represent the core phenomena found in the data: the tensions, possible organizational structures, and tactics.

Finally, the *theoretical coding* approach helped us to identify relationships among the categories. Thereby, theoretical coding relates the basic categories generated by selective coding to each other (Glaser, 1978). In this sense, the concrete relationships between organizational structures, tensions, and tactics were captured, particularly with respect to achieving ambidexterity for combinations of tensions in a given organizational structure.

4.4 Results

In the following, we first address the organizational structures that, according to the interviews, are often used by organizations to develop and operate ML systems. We then show the tensions organizations face in developing and operating ML systems, how they differ by organizational structure, and how organizations move toward hybrid structures over time by using appropriate tactics.

Organizational Structures for ML Development and Operation

The analysis of the interviews confirms that there are two main types of organizational structures in practice. They are subsumed under the umbrella terms *centralized* and *decentralized* structure.

A **centralized structure** includes all formats in which a centralized unit is concerned with the development and operation of ML that operates largely independently of other organizational processes. The centralized unit thereby provides ML systems and corresponding services for the remaining units. This includes structures that are, e.g., designated by the terms *AI center of excellence* or *AI competence center* (E02, E03, E10, E13). The centralized structure allows organizations to achieve the required visibility of ML within the organization. It simplifies the specification of a central strategy for data, ML development, and its operation and paves the way to define corresponding processes uniformly (E01, E02, E03). In addition, the centralized structure allows organizations to pool ML expertise that has been limited to date: *“With AI, it makes sense to start very centrally because ML development and operation still requires real niche skills. That means you build up a pool of experts to reach a critical mass.”* (E02).

While organizations with a centralized structure establish a standalone unit for ML development and operation, an organization with a **decentralized structure** locates developers in the respective business units where they work in multidisciplinary teams with domain experts to combine technical and business knowledge (e.g., specialists with an IT background). A decentralized structure, which may be referred to as silo-based (E01, E03, E09, E11, E12), grows rather passively from the bottom up and offers the advantage of being particularly fast and flexible (E01), less costly, and close to the knowledge of the respective business unit (E09). According to interviewee E12, this structure for the development and operation of ML systems is most evident in those organizations that are already very tech-savvy: *“And the other approach is that it grows up from the bottom. That’s a bit rarer, but we see it in organizations that are already technically oriented.”* (E12).

Juggling Tensions and Tactics in ML Development and Operation

Next, we adopted an ambidextrous perspective to identify competing demands during ML development and operation that can be represented as tensions. We examine how these tensions manifest in the given organizational structures and explain the tactics that organizations thus

adopt within the different structures to alleviate the resulting pressure. Finally, we describe how organizations shift toward hybrid structural approaches due to the respective chosen tactics.

Figure 2 illustrates an overview of the concepts derived from the data. Due to the different demands imposed by the *exploitative-oriented* conventional business compared to the *explorative* requirements arising from ML systems, tensions emerge in organizations trying to develop and operate these systems. To enable ML development and operation despite these tensions, organizations adopt different organizational structures, i.e., *centralized* or *decentralized structures*. However, the chosen structure, in turn, necessitates the use of certain adapted tactics. Central settings are already much more aligned with the requirements stemming from experimental ML systems, so *exploitation tactics* (March, 1991) need to be employed in a complementary way. In the case of a decentralized structure, ML development and operation is already closely oriented toward business. Therefore, *exploratory tactics* (March, 1991) are used, which are intended to break away from daily routines. As a result, organizations move toward *hybrid structures* over time, allowing a compromise between the demands of exploratory ML processes and straightforward conventional business processes, respectively, providing an ambidextrous solution for more successful ML development and operation (E09, E11, E12).

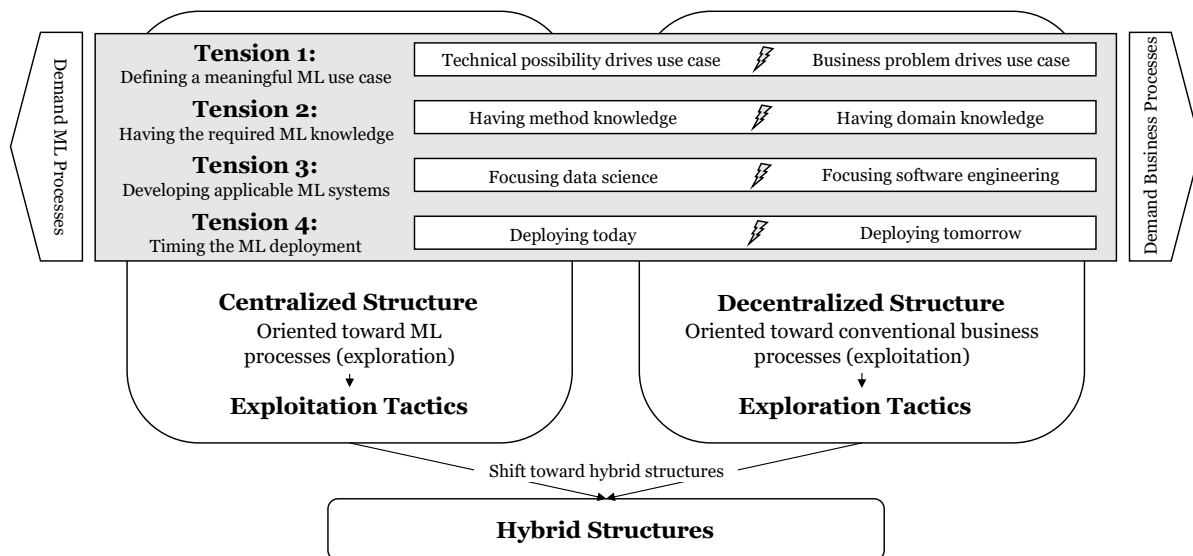


Figure 2. Ambidexterity in ML Development and Operation.

Tension 1: Defining a Meaningful ML Use Case

The first identified tension occurs in the early phase of the ML development process. To begin, an organization needs to identify a use case that it deems appropriate for applying ML systems (E15). The importance of use case selection is underscored by recent research (Jöhnk et al., 2020; Pumplun et al., 2019; Sturm et al., 2021c). However, the experts point out a problem that makes it difficult to identify meaningful use cases (E11, E15): *“It’s always a game: Just because you find the use case in business doesn’t mean that there is data, and vice versa”* (E11). Consequently, the technical requirements, e.g., in terms of data availability or algorithmic feasibility, may conflict with the demands posed by business, which can lead to a locking effect in ML development. Organizations must decide whether they prefer to drive the use case from the perspective of technical possibility (exploration) or take a problem-oriented approach (exploitation) and neglect technical realization for the time being (E13).

The experts also note that this tension can vary in the adopted organizational structure and can tend toward one pole or the other. More specifically, centralized units have a higher degree of thematic freedom in the development of ML use cases. However, they are more oriented toward the available data, algorithms, and technical infrastructure. In contrast, in decentralized structures, the focus is on the added value ML can bring to the business units: *“In a decentralized organizational structure, you are much closer to your problems, which have to be solved anyway. I think, as a central unit you have more freedom [...] to think out-of-the-box and to generally consider what could be done in the first place.”* (E09). In a decentralized structure, the developers of the ML systems have insights into the problems of the employees working for the business units or even face these problems themselves. Therefore, they are striving to help the employees within the business units to address their daily problems (E10). In contrast, centralized entities strongly encourage sharing about technical details (e.g., ML algorithms), which can lead to synergies and technology-driven ideas for new use cases (E10). As a result, the modes of thinking differ vastly between organizations pursuing a centralized structure compared to an organization with a decentralized approach: *“In the central system, the use cases are based on the mindset: ‘I have these tools, let’s make a cycle out of them and we can drive’. And in the decentralized structure people think more like, ‘I need a cycle, so let’s do something to make a cycle’”* (E10).

According to the different priorities observed in the different organizational structures, the experts suggest different tactics to relax the tension. In a centralized setting, proposed tactics should translate into better targeting of the exploitation-oriented business side to prevent the development of ML systems that end up being rejected by end users (E12). While the specialists in the central unit usually know what is technically feasible, they should, therefore, proactively gather the requirements of the business units, e.g., in the form of interviews: *“The easiest way to get purposeful use cases for the business is simply to go into the business units and interview them”* (E11). In contrast, everyday routines and problems determine ML use cases in a decentralized structure. As a consequence, particularly innovative use cases might be overlooked, which would theoretically be possible from a technical point of view. Therefore, organizations with a decentralized setting should encourage business units to explore and think more out-of-the-box, for example, through workshops or open time during working hours (E09).

Tension 2: Having the Required ML Knowledge

Another issue that can arise against the backdrop of scarce resources is the availability of different focal points of knowledge. Especially in the phase of data preparation and modeling, ML development requires not only methodological knowledge of algorithms and technical subtleties but also domain-specific expertise to enable a deep understanding of underlying data and business processes (E13, E16). Yet, because resources are limited, the provision of ML and domain knowledge cannot occur without limits. Organizations must decide what kind of skills they want to invest in (E09, E11).

Thereby, the availability and distribution of these types of knowledge varies in different organizational structures. Organizations that adopted a centralized structure bundle the scarce knowledge in ML and make it centrally available (E02). At the same time, this leads to ML knowledge being encapsulated from the domain knowledge of the business units (E15). In contrast, organizations that approach this problem in a decentralized manner have already combined ML and domain knowledge to a greater extent but will have limited ability to build specialized ML knowledge due to financial and time constraints. As a result, developers in a centralized setting may spend more time on technical training, sharing methodological knowledge, and implementing ML systems, while developers in a decentralized structure know

the domain better but lack the time and resources to pursue technical education: *“In the decentralized case, you’re more likely to have missing knowledge on the methodology side and in the centralized you’re more likely to have missing knowledge on the domain side”* (E09).

To overcome the respective challenges in the organizational structure pursued, the experts suggest various tactics. First, there must be a regular and intensive exchange between the central unit and the business units in a central setting. This can be done in various ways, from formats that facilitate communication (e.g., weekly status calls, workshops, the inclusion of visualization experts; E01, E04) to structural collaboration or staff exchanges: *“It is critical that the domain experts [...] and the technical methodological experts work closely together in a team. [...] And basically, it’s enough if you bring one domain expert into the entire development team or at least make him or her available for queries”* (E11). Second, it may be useful to establish other measures that facilitate the central unit’s ability to better understand the more exploitation-oriented business. For example, data dictionaries can be implemented to record definitions, representation rules, and relationships between data objects (E01). In contrast, ML developers in decentralized settings do not lack domain knowledge but often fail to acquire detailed methodological knowledge. Since decentralized structures often grow from the bottom up, driven by individuals with initial technical experience (E12), the experts assume that it is necessary to offer ML training opportunities, ensure freedom for self-directed upskilling, and promote open exchange on technical topics across business units. This requires decision-makers to create the necessary space for their employees’ personal development and relieve them of other daily routine tasks to enable exploration (E04, E12, E09, E10).

Tension 3: Develop Applicable ML Systems

Another challenge is the deployment and subsequent operation of the ML model in a way that users can apply (e.g., Amershi et al., 2019). According to E02, organizations currently struggle to transition the initial proof of concept into an applicable, integrated software system: *“I have a customer who builds proof of concepts very successfully. [...] But what they build are hacks. I can’t transfer them in any way into an industrialized form”* (E02). To effectively deploy ML systems, two development disciplines must come together, both of which are limited in their personnel capacity and time: data science, which aims to create and optimize the models at the core of ML systems, and software engineering, which is concerned with the integration of the models into the system landscape (E14). Expert E09 sees this as a trade-off between how much capacity and time should be allocated in data science and how much in software engineering during an ML project.

According to the experts, the concepts pursued in the different organizational structures for allocating resources vary significantly: In central units, data scientists are employed in particular, who deal with the statistical evaluation of the data and optimization of the algorithms (E09, E12): *“My feeling is [...] that many problems still lie in the field of software engineering. And that competence centers almost only hire data scientists who actually don’t have a focus on computer science, but rather on mathematics or statistics”* (E09). This is problematic since software engineering plays a critical role in the implementation of ML systems (E09). In contrast, the respective business units in a decentralized setting have a stronger focus on the future applicability of the ML systems due to their interdisciplinary composition: *“When it is a decentralized structure, then we may not need dedicated software engineers. We will only need the data scientists. Because software engineers are anyway working there”* (E10).

Developers working in the respective business units have closer contact with the later users, see their requirements, and work in close exchange. There are several ways to prevent a central

structure from focusing only on doing exploratory data science and neglecting the integrability of ML models into existing software systems: First, the conventional business unit and respective IT should be involved in the early stages of ML development. To this end, a bridge position can be created between the central unit and the business unit, which is dedicated to promoting and taking over continuous communication and exchange to facilitate exploitation (E12). In addition, software engineers should be actively recruited and located in the central unit to build a standard software solution, e.g., by creating programming interfaces, designing data pipelines, or enabling model versioning (E02, E13, E14). In a more decentralized setting, all-rounders will more likely take over the development and operation of ML systems (E09). In many cases, they already have prior engineering knowledge and more intensive contact with the potential end-users. Therefore, in a decentralized structure, tactics should focus on providing advanced training in the direction of data science and, in particular, mathematical-statistical methods (E04, E12, E09, E10). Some research has started to discuss this topic under the notion of citizen data scientists, who might compensate for the lack of dedicated data scientists in the future (Benbya et al., 2020a). Furthermore, it is possible to leverage ML services and platforms with automated components, such as Google's AI platform or Amazon's AWS AI services, which offer pre-built user interfaces, managed ML workflows, or ready-to-use validation tools that require less in-depth mathematical knowledge (E08, E09).

Tension 4: Timing the ML Deployment

The fourth tension comprises the matter of when to make an ML system available to users and start the operational phase. Many organizations aim to deploy their ML systems as early as possible. This allows new training data to be collected and the model to be improved over time. However, early deployments of ML systems can lead to a loss of trust among users, as these systems can exhibit initially low accuracy: *"A trade-off is always: When do you integrate something into the product, how long do you wait? [...] And the idea is always, the moment it's out in the world, we can learn even more about it, and make the model better. [...] But you have the problem that the benefit might be low in the beginning. And then, maybe, the business units don't want to use it anymore"* (E09).

In a centralized approach, this issue is particularly pronounced given the strict separation of development and operation vs. application and less exchange between the respective units. Therefore, the experts state that, especially in a central setting, they must meet a minimum quality (e.g., 80% accuracy) of the ML model until it is released to the business units (E09, E10). Conversely, in a decentralized setting, the experts see more of an opportunity to deploy today since even if the model is of lower quality, there is always a contact person on site. In addition, the experts assume that there is a higher level of trust as developers and users work more closely together (E10).

Therefore, especially in a centralized structure, it makes sense to set expectations for ML systems correctly from the beginning and actively communicate what can be expected from probabilistic ML systems to the business (E12). Therefore, users must be familiarized with ML systems and be aware that the system provided is still in its training process. For this purpose, it may be helpful, e.g., to create a test environment that can be accessed only by selected users who are aware that the ML system is still being improved (E01). This test environment could be operated until a minimum level of accuracy is reached by the ML system that justifies its use in a productive environment. According to interview participants, the tension between deploying the ML system today and tomorrow is less acute in a decentralized structure. Here, it makes sense to deploy ML systems earlier on to explore user behavior, collect data, and to leverage potentials quickly.

Indeed, sometimes a few percent improvements in process efficiency can lead to a positive return of ML, so it is not necessary to achieve high accuracy right at the start (E01).

Progressing Toward Hybrid Structures

Overall, it appears that the proposed tactics to alleviate tensions in the respective organizational structures induce a shift toward more *hybrid organizational approaches* (see Figure 2). While organizations in a centralized setting seek convergence of the otherwise loosely integrated central unit with the business (i.e., exploitation tactics), organizations taking a decentralized structure seek to gain more flexibility, methodological expertise, and focus on ML development and operation (i.e., exploration tactics). Consequently, organizations in a centralized environment primarily pursue exploitation tactics to better meet business demands, while decentralized organizations must become more exploratory to successfully develop and operate ML systems. Therefore, from an ambidextrous perspective, ML development and operation effectiveness can be achieved if a hybrid structure is pursued that satisfies both sides' demands, ML-specific and business-specific. In line with this, the experts emphasize that organizations are evolving toward hybrid structures in the long term, or at least have the goal of doing so (E04, E09, E10, E12, E16). Thereby, hybrid structures can take a variety of different forms that fall on different points of the spectrum between centralized and decentralized structures (and thus between exploration and exploitation respectively). For example, interview participants describe a range of approaches, from *temporary program structures, lab, or task force interventions* (e.g., six months of leave in a lab) to permanent implementation of *hub-and-spoke organizational structures* that include multiple central units located closer to the respective business units (E01, E03, E04, E11, E16). Regardless of how long-term the approach may be, organizations with hybrid structures seek to foster talent and ideas, achieve the necessary attention within the organization, provide infrastructure independent of the existing IT, and enable flexible pathways for ML development and operation without losing sight of the ML systems' business relevance (E01, E03, E04, E11, E16).

4.5 Discussion

The development and operation of ML systems differ significantly from the development and operation of prior software systems (Amershi et al., 2019). Realizing the development and operation of ML systems is non-trivial and requires new and specific competencies in organizations that already have to manage their scarce resources today. At present, most organizations are still in the early stages of ML development and are struggling to decide on the next steps to take. Despite previous research in IT innovation (e.g., Haffke et al., 2017a), the highly experimental development and operation of ML systems are still poorly understood. This study thus seeks to take a first step in identifying the relationship between organizational structures, tensions, and potential tactics that can contribute to successful ML projects.

At a theoretical level, our study contributes to understanding how organizations approach ML development and operation and why they do so in particular ways. First, this research discusses organizational structures that are followed by organizations to develop ML systems (i.e., centralized, decentralized) and conceptualizes tensions emerging in different organizational structures. In this regard, we were able to identify four different tensions, namely (1) *defining a meaningful ML use case*, (2) *having the required ML knowledge*, (3) *developing applicable ML systems*, and (4) *timing the ML deployment*. Our results overlap with tensions described in the literature that arise on the side of ML-utilizing organizations, such as the tension between letting

an ML system learn autonomously and having domain experts guide its learning process (van den Broek et al., 2021; Seidel et al., 2020). The key driver behind this tension in development and operation, however, seems to be budgetary constraints rather than optimal ML learning processes (see tension 2). Interestingly, we found no evidence that the fear of being replaced by machines was strong enough in development and operation to create tensions. Vaast and Pinsonneault (2021) note that data scientists who are deeply involved in creating ML systems engage in identity work, continuously redefining their occupational identity, and thus alleviate the tensions stemming from the autonomy of ML systems. While this might explain the absence of these tensions during development and operation so far, the increasing capability of ML systems to take over tasks from knowledge workers (e.g., Faraj et al., 2018) might intensify them in the future. Longitudinal studies might therefore be appropriate to identify how improvements in ML capabilities change the arising tensions during development and operation. Second, we identify exploitation or exploration tactics that organizations utilize to alleviate tensions, and we further contextualize the notion of ambidexterity in the development and operation of ML. Thereby, our identified tactics stem from both the domain of structural ambidexterity (e.g., workshops or open time during working hours as tactics to alleviate tension 1) and the domain of contextual ambidexterity (e.g., a bridge position between the central unit and the business unit to alleviate tension 3), providing insights into how organizations integrate structural and contextual ambidexterity and thus contributing to ambidexterity research on such integrated approaches as called for by Ossenbrink et al. (2019). Our study shows that in the organizational structures centralized and decentralized, the discussed tensions each intensify in the direction of one of their two poles, putting pressure on organizations to employ tactics to balance the respective tensions: Organizations with a centralized setting are already oriented toward the explorative requirements for ML development and thus engage in exploitation activities to satisfy business process demands. Conversely, organizations with decentralized approaches are well-aligned with the business side and employ exploration tactics to meet the requirements of ML development and operation. Organizations may, therefore, achieve ambidexterity by targeting hybrid approaches to alleviate problematic poles of the tensions while minimizing their overall intensity. Thereby, hybrid structures can take various forms that lean toward more central or decentral structures and thus to the respective poles of the tensions. In this way, organizations may be able to attend to the competing demands of both sides of the exploitation-exploration duality, which promises superior long-term performance (Raisch & Birkinshaw, 2008; Smith & Lewis, 2011; Smith et al., 2011). In this regard, future research could examine which hybrid structures for the development and operation of ML can be developed in detail and how exactly organizations shift to these hybrid structures – thereby possibly answering a call for ambidexterity research on how organizations navigate the spectrum between exploration and exploitation over time (Schad et al., 2017). Third, we contribute to IS research by applying computer-assisted analysis methods within an explorative interview study informed by grounded theory methodology. Our study shows how ML methods can be incorporated into IS studies to more effectively analyze qualitative data, provide additional avenues for exhaustive data exploration, and offer new insights that would otherwise be overlooked.

Besides these theoretical contributions, our study provides guidance to practitioners who intend to develop and operate ML systems. It highlights several organizational structures that are routinely used in the development and operation of ML. Organizations seeking to deploy ML systems can be guided by these organizational structures and adapt their own structures accordingly, depending on their previous setup. Furthermore, the study helps organizations identify potential tensions in their organizational structures early on that may arise from differing

requirements of integrating ML systems and business-driven processes. In this context, the study also highlights concrete actions that organizations of a given structure can take to alleviate tensions. In particular, we show that organizations must be careful to consider the requirements of both sides, those of the more explorative ML systems development and operation and those of the more exploitative-oriented business. Consequently, the findings that emerge from this study can inform practitioners developing a longer-term ML strategy. They are essential to allocate resources wisely, e.g., for the structural design of an organization, especially when considering hybrid structural approaches, or for the reduction of tensions, for instance, by providing more freedom for self-directed upskilling in decentralized structures. As a result, this study provides more structure to the discussion revolving around ML development and operation in practice.

As with any study, this study has limitations that necessitate further research. Our study follows an exploratory approach based on the analysis of qualitative interviews. While the interviews have provided us with a rich source of valuable information, a quantitative study might offer more insight into the generalizability of the derived concepts, or perhaps quantitative insights into the impact of tactics on the success of ML projects. Nonetheless, we are confident in the validity of our findings, as we carefully selected our experts using a theoretical sampling approach, which resulted in interviewing participants from organizations of different sizes and industries (Birks et al., 2013; Glaser & Strauss, 1967). Further studies could examine different organizational settings and investigate whether the tensions and tactics are prioritized differently under varying circumstances, such as at different points in ML development and operation.

4.6 Conclusion

Overall, our study takes a first step in showing how the arising tensions in the development and operation of ML systems can be theoretically conceptualized, whereas previous research has been a more technical discussion (Amershi et al., 2019; Studer et al., 2021). It contributes to an integrative perspective on ML development and operation guided by ambidexterity research. As results, we conceptualize four tensions that emerge during ML development and operation, which manifest differently depending on the adopted organizational structure. Furthermore, we shed light on the tactics that organizations use to alleviate tensions in particular structural settings and that ultimately drive organizations toward hybrid structural approaches and thus ambidexterity.

Acknowledgment

This research and development project is/was funded by the German Federal Ministry of Education and Research (BMBF) within the “Innovations for Tomorrow’s Production, Services, and Work” Program (funding number 02L19C150) and implemented by the Project Management Agency Karlsruhe (PTKA). The authors are responsible for the content of this publication.

5 Research Paper C: Delivering Value Through Explainable Artificial Intelligence

Title: A User-Centric Approach to Explainable AI in Corporate Performance Management

Authors: Vetter, Oliver A.; Efremov, Alexander

Published in: Proceedings of the 27th Pacific Asia Conference on Information Systems (PACIS 2023)

Awarded with the Best Complete Paper Runner-Up Award

Abstract

Machine learning (ML) applications have surged in popularity in the industry, however, the lack of transparency of ML-models often impedes the usability of ML in practice. Especially in the corporate performance management (CPM) domain, transparency is crucial to support corporate decision-making processes. To address this challenge, approaches of explainable artificial intelligence (XAI) provide solutions to reduce the opacity of ML-based systems. This design science study further builds on prior user experience (UX) and user interface (UI) focused XAI-research, to develop a user-centric approach to XAI for the CPM field. As key results, we identify design principles in three decomposition layers, including ten explainability UI-elements that we developed and evaluated through seven interviews. These results complement prior research by focusing it on the CPM domain and provide practitioners with concrete guidelines to foster ML adoption in the CPM field.

Keywords: Explainable artificial intelligence, corporate performance management, UX

5.1 Introduction

Machine learning (ML) use in organizations has grown rapidly in recent years (Jordan & Mitchell, 2015). In a study by McKinsey, 56% out of 1.843 participants from different industry fields reported using ML in at least one use case (Chui et al., 2021). One area in which organizations can stand to benefit from utilizing ML systems is corporate performance management (CPM), a company practice that deals with strategic and tactical activities, such as planning, budgeting, and forecasting. ML-based techniques have much potential for supporting humans in these tasks, especially in forecasting, as ML can learn from previous data and take outside effects into account (Makridakis et al., 2023). The opaque and abstract character of ML models is a significant obstacle, however, particularly for decision-making (Adadi & Berrada, 2018). This issue has given rise to the study area of explainable artificial intelligence (XAI), which aims to elaborate to users how ML models function or how a particular prediction was generated (Adadi & Berrada, 2018). Research has developed and evaluated XAI techniques to gain insights into ML models and algorithms in order to generate these explanations (Adadi & Berrada, 2018; Guidotti et al., 2019).

Despite the potential of ML in forecasting use cases, the technology is rarely used in the CPM domain. For instance, in a study by McKinsey, ML had the lowest adoption rates in the CPM-related business functions of strategy and corporate finance at 7% and 6%, respectively (Chui et al., 2021). Despite the research advances in the field of XAI, 34% of respondents in emerging economies and 44% in developed economies perceived the topic of explainability as a relevant risk when using ML (Chui et al., 2021). A possible reason for this gap could lie in the design of ML-based systems. As shown by DARPA's XAI program, a high-quality user experience (UX) and user interface (UI) for XAI systems are crucial to foster user understanding (Gunning et al., 2021). Moreover, there is no one-size-fits-all strategy for XAI, as various users in different circumstances demand diverse solutions (Gunning et al., 2021). UX-/UI-centered research has therefore evaluated different user groups, motivations for explanations, and formulated design principles (see Laato et al., 2022; Liao et al., 2020). Yet, many studies focus on providing explanations to data scientists, which is why recent literature calls for matching XAI techniques with more lay users and their needs (Liao et al., 2020; Brennen, 2020). This lack of focus on the end users may explain the gap between the low current adoption of ML systems and their considerable potential in the CPM sector, as low usability might discourage non-data-scientists from adopting and using ML. In this study, we thus aim to provide CPM- and user-specific answers to the following research questions:

Research question 1 (RQ 1): What are the goals of the user groups of ML-based CPM systems, and what kind of explanations do they require to achieve them?

Research question 2 (RQ 2): How can an ML-based CPM system effectively and efficiently provide these explanations to the users?

To answer these research questions, we follow the design science research process proposed by Peffers et al. (2006). Thereby, our study builds on preceding research concerning XAI techniques and their UX-/UI-centered design by developing and evaluating a user-centric approach to provide explanations in an existing ML-based CPM information system (CPM system) which is utilized in a forecasting use case and being developed by a European provider of enterprise service management solutions. To this end, we identify the goals and requirements of the user groups of the CPM system. We subsequently derive design principles (DPs) to meet these requirements in the UX, which we apply, refine, and evaluate through six interviews with management and CPM experts. With this study, we thus make important contributions to UX-/UI-

centered research for XAI systems by distinguishing and validating DP for distinct user groups of ML-based CPM systems. Furthermore, to our knowledge, this is the first study that elaborates DP for UX-/UI-design tailored specifically for users in the CPM context and their unique requirements. Thus, we provide scholars with fertile ground for future research on user-centric XAI approaches in the CPM or adjacent domains while offering guidelines to practitioners that they can utilize to bridge the gap between CPM experts unfamiliar with ML and the potential promised by ML-powered CPM systems.

5.2 Theoretical Background

This section first describes the business practice of CPM and then expands on the utilization of ML solutions for planning and forecasting, taking XAI approaches and UX/UI-centric research into consideration.

5.2.1 *Corporate Performance Management*

CPM has been described as a system combining management processes with corresponding business intelligence (BI) information systems (Miranda, 2004). BI information systems enable companies to collect and analyze the data enabling CPM practices (Miranda, 2004; Frolick & Ariyachandra, 2006). The objective of CPM is generally to support corporate decision-making in ensuring the company performs well in its success metrics, such as revenue or profit (Frolick & Ariyachandra, 2006). Because of its quantitative nature, CPM thus usually focuses on metrics that can be expressed in financial figures, such as revenue and profit, as overarching goals (Frolick & Ariyachandra, 2006). On a deeper level and depending on the company's business model, CPM practitioners can also look into figures such as specific or aggregated sales numbers for a specific product or the costs a particular department generates (Frolick & Ariyachandra, 2006). There are different frameworks and approaches to summarizing the processes that make up the CPM of a company: generally, they all include processes for the planning, analysis, and monitoring of the predefined performance metrics (Richards et al., 2019; Frolick & Ariyachandra, 2006).

In the context of CPM, planning can be described as the process of gathering relevant information for CPM decision-making, such as budget allocations (Rogers et al., 1999). This information is used to base strategies around them (Rogers et al., 1999). For this purpose, it is important to use internal and external information to predict how certain figures will perform in the future (Richards et al., 2019). This step of forecasting is essential to the planning process, and the correct execution can be a critical step to outperform the company's competition (Frolick & Ariyachandra, 2006). Extensive research has been done on different types of forecasting (Bontempi et al., 2013). This ranges from stock price prediction to the CPM-relevant use case of sales prediction (Pavlyshenko, 2019). In this study, we focus exclusively on use cases entailing the forecasting of business figures in the context of CPM. In practice, forecasting is often done in the form of a time series problem (Bontempi et al., 2013). This means a collection of historical data, all in the form of a time series, is used to predict future values based on these variables (Bontempi et al., 2013). There are different models and algorithms to transform the input variables into the desired output variables. Simple or advanced statistical methods, like forecasts, are often robust methods and offer their own advantages (Adya & Collopy, 1998; Makridakis et al., 2023; Spiliotis et al., 2019): They do not heavily rely on the amount and the quality of their input data and are also not very computation heavy because of the simplicity of their algorithms, making them the dominant methods for forecasting in the past. They do, however, also have some weaknesses, as they only

prescribe the input data and do not recognize causation, for instance (Barker, 2020). The rise of data availability and quality, together with an increase in readily and cheaply available computation power, laid the groundwork for ML-based approaches to these forecasting use cases (Makridakis et al., 2023).

5.2.2 *Explainable Artificial Intelligence for Forecasting*

ML can be described as the algorithmic generation of a model from provided data by extracting patterns within the data (Russell & Norvig, 2021). Most modern artificial intelligence (AI) systems are implemented using ML technologies (Brynjolfsson & Mitchell, 2017). Therefore, we use the term ML to refer to ML-based instances of AI in this study. Time series forecasting can be done with supervised learning, a subcategory of ML (Bontempi et al., 2013). This is due to the circumstance that historical data is used, which is regularly labeled (Bontempi et al., 2013). A good example is historical sales numbers, as the label would be the sales volume mapped to the date the sale was concluded (Ma & Fildes, 2021). There exists a lot of research on ML-based approaches to forecasting, mostly focusing on designing the right type of algorithm for the defined use cases. Prior research includes ML-based solutions for energy forecasting, sales forecasting, or other financial figures (Ghoddusi et al., 2019; Pavlyshenko, 2019; Wasserbacher & Spindler, 2022). This research shows that ML-based approaches can perform quite well in planning use cases. In recent years deep learning algorithms are starting to catch up or surpass traditional algorithms in terms of performance (Hewamalage et al., 2021). Deep learning refers to especially complicated ML techniques, often associated with models with multilayered circuit structures, which are referred to as artificial neural networks (Russell & Norvig, 2021). They also exacerbate one of the biggest disadvantages of ML: The actual or perceived lack of transparency and the lack of explainability (Gunning et al., 2021). This disadvantage has brought out another important aspect of ML, XAI, which will be a focus of our study (Adadi & Berrada, 2018).

ML models are often described and perceived as a black box, and XAI tries to address this (Brennen, 2020). Therefore, XAI can be summarized as approaches seeking to explain aspects of ML-based systems to their users and stakeholders (Langer et al., 2021). For terminological clarity, we first describe the concepts of interpretability and explainability as understood within this paper. Interpretability can be defined as the grade to which the model and its predictions are interpretable to the user (Russell & Norvig, 2021). This means that based only on the inspection of the model, a human can derive the reasoning behind a certain output and could also predict the output for a different input. Linear regressions or decision trees are interpretable because the human can simply go along the tree with the respective input or calculate according to the regression parameters (Russell & Norvig, 2021). Following this definition, interpretability can only be achieved by choosing an interpretable algorithm. Examples of non-interpretable algorithms are deep learning algorithms due to their complex structure and the vast number of parameters and layers they use (Castelvecchi, 2016). They and other uninterpretable ML algorithms are commonly referred to as black box models or algorithms (Castelvecchi, 2016). In comparison to interpretability, explainability is not something inherent to the algorithm. It can be provided by posthoc processes, e.g., by another algorithm being trained on top of the model to be explained. This new algorithm is in itself interpretable but performs similarly to the one to be explained. Because the explanation occurs after the model is already trained, methods designed to provide explainability can also be described as posthoc techniques (Barredo Arrieta et al., 2020; Russell & Norvig, 2021).

Most posthoc techniques are model-agnostic, i.e., they can be applied posthoc on every type of ML model. These posthoc techniques themselves can be divided into two categories: Global and local techniques (Murdoch et al., 2019): Global techniques operate on a dataset level, providing explainability into global relations and patterns the model has learned. Local techniques operate on a prediction level, providing explainability into individual predictions. Barredo Arrieta et al. (2020) also provide further, not mutually exclusive, distinctions of posthoc techniques, such as explanations by simplification, feature relevance or simply visual explanations: Explanations by simplification generally try to extract rules based on how the model works and explain those rules to the users. Feature relevance refers to a wide array of explanations regarding the input variables of the model, i.e., the features. They use different approaches like game theory to learn how to explain the features, and afterward provide information on, e.g., the influence of the features on the model, the importance of the features to the model performance, or the interaction of the different features. Visual explanations in turn can be used to present the information gathered from the feature relevance techniques. As presented in this subsection, model-agnostic, posthoc techniques allow for abstraction from the utilized ML algorithms. Consequently, we focus the presented UX/UI approach on these to ensure the transferability of our results for most ML-based solutions.

5.2.3 *UX-focused XAI Research*

When designing a UX-focused approach, different aspects need to be kept in mind. First, it is important to provide explanations into adjacent fields in the ML-based systems, such as the model performance or the used data (Liao et al., 2020). This also motivates a more general approach in terms of looking at the whole UX of the ML-based system, instead of just the parts that are ML-related in the narrow sense. Liao et al. (2020) provide key factors that influence user requirements and design recommendations. Among others, these include the motivation of the users to obtain explainability, the type of users in terms of, e.g., domain knowledge or prior experience with data science topics as well as their role in the ML-based system, and lastly the decision context in which the explanations are provided. Because of the focus on the forecasting use case and posthoc model-agnostic techniques, other factors identified, such as the data and algorithm type, are omitted for this study. In a data science-heavy domain, motivations for explainability can be divided into debugging the model, identifying biases, and building trust (Brennen, 2020). Liao et al. (2020) argue that the sheer volume of different contexts and user motivations makes it difficult to predefine general explainability needs. They, therefore, describe a question-driven approach to generating UX-Guidelines for ML-based systems. Focusing on strictly posthoc techniques, they aggregated explainability methods, which resulted in the addition of counterfactual explanations and example-based explanations to the list of explainability methods and techniques. They then built a question bank with questions that users of ML-based systems might have in terms of explainability (Liao et al., 2020). After considering a broader look of users on the ML-based system, they formulated ten question categories, which can be summarized as questions regarding input and output data of the ML-based system, performance of the ML model, “How”-questions concerning how the ML-based system and the ML model generate their predictions, and questions concerning why specific values were or were not predicted as well as what could happen by differing parameters of the ML-based system.

As discussed by Liao et al. (2020) it is important to consider the roles of the users. Meske et al. (2022) define five stakeholder groups in XAI systems, which can be divided into three groups based on their interaction with the ML-based system: The first group includes the three

stakeholder types of regulators, managers, and developers. It focuses on general regulatory certification, managing and controlling, and the responsibility for the development of the ML-based system. The second group, named users, relates to the actual end users of the ML-based systems. The third group, individuals affected by ML-based decisions, refers to stakeholders who may have no direct interaction with the ML-based system but still experience the consequences of its deployment. Another aspect to consider while designing the UX of an ML-based system utilizing XAI is the quality metrics such an approach should adhere to. Meske et al. (2022) further discuss different personalized quality criteria, i.e., metrics that explanations should fulfill. They argue that explanations should generally be interpretable by the user, in a way that users should be able to comprehend them and that they seem plausible. Moreover, an important aspect is the effort required to get the explanations. Zhou et al. (2021) also discuss the clarity, broadness, simplicity, completeness, and soundness of explanations. Oh et al. (2018) evaluated their UX research via general usability metrics, such as the ease of use as well as the ease of learning the use, and more ML-specific metrics, such as the comprehensibility and the controllability of the ML-based system. Lastly, research also already provides several principles and guidelines for the design of the UX or UI in the context of XAI. Laato et al. (2022) recommend always considering visualizations. These types of recommendations were also provided in general UI research, to reduce the cognitive load of users and thereby increase usability satisfaction (Hu et al., 1999). Similarly, the use of coloring in the UI, when well combined with other UI elements, can be helpful (Hu et al., 1999). This is further emphasized when presenting business information, such as in CPM use cases (Bačić & Fadlalla, 2016). Storytelling and the right use of symbols can also be beneficial (Bačić & Fadlalla, 2016). To summarize, there are metrics for the general perception of the UX or UI, e.g., concerning the visualizations and general usability, as well as metrics specific to the explainability goal and the ML domain that evaluate if the provided explanations are understandable and actionable. Our work builds on these suggestions, specifying and validating them for users in the CPM field.

5.3 Methodology

For this study, we followed the design science research process proposed by Peffers et al. (2006), which is appropriate for research on applicable solutions to an existing problem. It includes the six phases problem identification and motivation (1), objectives of a solution (2), design and development (3), demonstration (4), evaluation (5), and communication (6). This process can be iterative and jump backs to any of the phases are possible. Design science in the IS domain can aim to deliver DPs that include prescriptive statements on how to perform activities to solve the problem (Gregor et al., 2020). They can have three different foci: First, they can aim to describe what users should be able to do with the artifact (principles about user activity). Second, they can aim to describe what features should be built into the artifact, such as requirements on a technical or functional level. And third, they combine the definitions of the first two fields, by describing, what users should be able to do with the artifacts as well as the features the artifact should possess (principles about user activity and an artifact) (Gregor et al., 2020). Further, DPs should address and define the actors involved in them, as well as the decomposition of the principle in a hierarchical manner due to the high complexity of IS (Gregor et al., 2020). The object of examination chosen for this study is an existing ML-based CPM system developed by a medium-sized European supplier of enterprise service management software. The system allows for the creation of ML models to forecast business figures based on historical data and data from external providers.

The first phase, problem identification, and motivation consisted of literature research and one qualitative interview. The goals of the literature research were to gain knowledge on ML-based use cases for the CPM domain with a focus on forecasting, the possibilities of XAI, and the existing user-centered research concerning the adoption of XAI in ML-based systems. The results of the literature research are presented in the previous section of this work. To extend the results of the literature in terms of the CPM context, an initial interview with a CPM expert (IIC) was conducted. Following the guidelines for qualitative interviews, the interview allowed for flexibility and was confidential (Myers & Newman, 2007). The semi-structured interview aimed to refine the collected knowledge on the different user groups of ML-based CPM systems, as well as on the requirements they have in terms of explainability. The interview was transcribed with a content-driven approach, excluding filler words and correcting obvious grammatical errors. Because assumptions regarding user groups and requirements were made before the interview, a directed content analysis approach was used (Hsieh & Shannon, 2005). The aforementioned assumptions were used as predefined codes in the following coding process as described by Saldaña (2009).

The second phase, objectives of a solution, aims to define objectives, which the artifact in the form of the DPs will try to accomplish. To achieve this, the results of the literature research and the results of IIC were analyzed to first define the relevant user groups and their main goal in the ML-based CPM system. Based on that, user group-specific requirements for explainability were derived. These included the knowledge that the users should gain from the explanations and also exemplary actions that users should be able to perform with the help of the provided explanations, with the latter focusing on decision-making actions.

In the third phase, design and development, UX/UI-centric DPs for incorporating XAI into ML-based CPM systems were created. To later demonstrate and evaluate the principles in the next phases, prototypical artifacts in the form of UI mock-ups were designed following the DPs. This phase can thus be split into two parts: the development of the DPs and the design of mock-ups. First, derived from the main goals and objectives, one DP was created for each user group. This principle addressed the general need to enable the user group to achieve this goal. Focusing on three user groups, this produced three DPs (DP 1 – DP 3). Based on the research done in phase one, explainability requirements were derived and assigned to the appropriate DP. This resulted in six explainability requirements (XRs; XR 1 – XR 6). Then, based on explainability techniques, ten explainability elements (XEs) were defined (XE 1 – XE 10), which were mapped according to their possibility of satisfying the explainability requirements. The definition of the XE contains the used explainability technique and some further meta-information. This includes the way in which the XE should provide its explanations and what kind of information should be provided alongside the technique. In a separate step, three general UX/UI DPs (UDPs) were formulated (UDP 1 – UDP 3) based on prior research and practical guidelines provided by the company. After the principles and associated definitions were created, representative UI mock-ups were designed. For their creation, we first considered which XR the respective XE is aiming to fulfill and, therefore, which information the CPM system and thus the mock-up must convey, before elaborating on the presentation and visualization of the information. To allow for a fast and flexible design, they were created in Microsoft PowerPoint. For each XE, 2-3 mock-up alternatives were created using applicable UDPs to examine how the XE is perceived by the user groups (e.g., see Figure 4). The mock-ups used representative dummy data and information to illustrate how actual explanations would be provided via the XE. During this phase, we collaborated closely with both the ML development lead as well as the product owner of the examined CPM solution, involving them in multiple feedback loops. Especially the product owner had valuable insights on the requirements

and desires of the CPM clients of the company as well, which were incorporated into the DPs and the mock-ups.

The fourth phase, demonstration, and the fifth phase, evaluation, were conducted in the same interview setting, and are thus covered together. First, an interview approach to evaluate the DPs was developed. Based on the research results, generic and specific metrics for an evaluation of the XE were defined. These focused on more general quality (GQ) metrics as well as on XAI-specific quality (XQ) metrics (see Section 5.2). The first part, therefore, concentrated on general UI quality requirements with the latter focusing on quality requirements in regard to the provided explanations. The evaluated quality metrics are the following: Easy interpretation, meaning that the process of interpreting the representation is without much effort (GQ 1); Intuitive interpretation, meaning that the process of interpreting the representation is possible without a lot of knowledge required (GQ 2); Easy-to-learn interpretation, meaning that the knowledge to interpret the representation can be easily gained (GQ 3); Satisfactory interpretation, meaning that the process of interpreting the representation leaves the interpreter in a satisfied state regarding his goal with the interpretation (GQ 4); Transparent presentation of the explanations, meaning that the representation provides all the expected information (GQ 5); Complete explanations, meaning that all the relevant explanations for the regarded use case were provided (XQ 1); Plausible explanations, meaning that the occurrence of the explanations makes sense and can be understood (XQ 2); Trust-building explanations, meaning that the explanations help the user understand the ML-approach in such a way that the user gains trust in it (XQ 3). Additionally, as each XE should help the users achieve their main goal by fulfilling the assigned XRs, specific quality metrics are defined for each XE/XR combinations: Knowledge measures (KMs) describe whether the XE presents the right knowledge, i.e., the right explanations to achieve the XR. Actionability measures (AMs) describe, whether based on the provided knowledge, decisions in connection with the XR can be made.

Table 7. Overview of Interviewees for Mock-up Evaluation.

Interviewee ID	Role in organization	User group
IIC, IC 1	Management in a planning department	Model creator
IC 2	IT-Consulting for CPM solutions	Model creator
IC 3	Sales and presales for CPM solutions	Model creator
IM 1	Middle management	Model user, data consumer
IM 2	Upper management	Model user, data consumer
IM 3	Middle management	Model user, data consumer

The interviews were conducted with three CPM experts (IC 1 – IC 3) and three interviewees with management roles (IM 1 – IM 3). Table 7 shows their respective user role and their role in the organization. The CPM experts were asked to evaluate nine XEs and the management group eight XEs, according to the mapping of the XEs to their respective DPs (see Figure 3). Because of the similar XRs for the user group of model users and data consumers (see Section 5.4), the second interview round was used to evaluate both DP and the associated XE to XR couplings. We conducted all interviews according to the following structure: 1) The general context of the interview was established. The interviewer explained the general approach for the interview and the goals of the study and the interviewee provided further background on their knowledge

regarding CPM-systems and ML if necessary. 2) The interviewees were presented the user groups of ML-based CPM-systems, as well as the goals of the user groups, and which explainability requirements were derived from them. 3) The interviewees were guided through a fictional case study for ML-based forecasting of ice cream sales. This included a fictional management role for the interviewee, fictional decisions for which the ML-generated forecasts were needed, and how the ML-based CPM system would be used to generate a ML model as well as predictions from the model. 4) The interviewees were guided through the evaluation methodology. For this, all the generic quality metrics were presented in the form of hypothetical statements and all the other steps taken in one evaluation iteration were explained. 5) The interviewees were iteratively asked to evaluate each of the XE represented by the mock-ups. In a first step, two or three slightly different mock-ups for the respective XE were presented to the interviewee in random order. After an alternative was shown, the interviewee was asked to signal when the next can be presented. If needed, the interviewer provided explanations on what was shown in the representation. The alternatives served to evaluate some of UI-focused DPs, as the mock-up alternatives generally displayed the same information but in a different way or on different aggregation levels. The interviewee noted the mock-up alternative which they liked best. Furthermore, the alternatives were used to convey that there is not one unique solution to use the XE. This way, they could evaluate the XE independent from a case where they did not like a particular illustration even though the provided information was fitting. Next, the interviewee rated the respective XE according to the generic and specific quality metrics. For this, we used a balanced Likert scale with five possible expressions (1) Strongly disagree; 2) Rather disagree; 3) Partially agree; 4) Rather agree and 5) Strongly agree) and instructed our interviewees to rate the XE from 1 to 5, with 3 being the midpoint (Likert, 1932). Depending on the results, the interviewer asked follow-up clarifying questions. All other qualitative comments by the interviewee were also recorded and transcribed. The evaluation was conducted by analyzing the quantitative assessment of the mock-ups via the Likert items and the qualitative comments and feedback provided by the interviewees. The refined results of this study are presented in the following section.

The communication of the obtained results through this study concludes the design science research process to develop a user-centric approach to XAI in CPM systems.

5.4 Results

As literature has shown, XAI-based approaches to provide explainability should keep the users, their motivation, and the context in mind. This study, therefore, builds on the identified three user groups creators, end-users, and affected individuals of the ML-based systems or ML models (Meske et al., 2022). In order to account for the context of a forecasting process in a CPM system, we refine them as follows: By focusing the creator group onto ML models instead of the whole ML-based system, our study defines the user group of model creators. They are in charge of creating, validating, and administrating ML models for different use cases. The model can be a standard statistical model or an ML model. Exemplary company roles are administrators or data engineers/scientists in the planning departments. Secondly, this work splits the end-users group into model users and data consumers. Model users utilize the created and deployed model to generate forecast data. They validate the generated data, may make some adjustments, and report them in the CPM system. Exemplary company roles are planners or business analysts. By using the model they create data, which in turn is used by the data consumers. Data consumers use the reported forecast data to evaluate the company's plans and performance and make decisions

based on them. They may also choose to relay information to non-user individuals affected by the ML-generated decisions. Exemplary company roles are managers or other decision-makers. As non-user individuals affected by the decisions of the ML model in practice typically do not gain access to the CPM systems or its explanations, we could not derive meaningful DPs for this group and thus omitted it from this study.

5.4.1 Design Principles and Derived Explainability Elements

The DPs identified in this study are structured in three decomposition layers. Incorporating the main goals of the user groups, we first create three overarching DPs, one for every user group for which the ML system is designed. These are shown in Figure 3 and describe which user group is addressed and what the explanations should be able to accomplish to help the users achieve their goals through the ML-based CPM system. Next, as discussed in the Section 5.2, research has identified different requirements that users can have based on the questions they ask in the context of ML-based systems (e.g., Liao et al., 2020). Among them are questions relating to the input and output data of the ML-based system, the performance of ML models, how the ML models work, or how certain predictions are generated. Taking them, the CPM use case, and the IIC results into account, this study further identifies six XRs that ML-based systems must meet in order to satisfy the questions posed by the three user groups and to thus help them reach their goals. This resulted in the following XRs, which we mapped to the overarching DPs according to which user group's demands they reflect, as shown in Figure 3:

XR 1: Provide explanations to help understand and evaluate the input data. Users should know what kind of data was used, how big the data sample is, and how good the quality of the data is. While this is important for any kind of data analytics, it is crucial for ML models due to their lack of transparency.

XR 2: Provide explanations to help understand and evaluate different ML models. Users should be able to see and understand the used ML algorithms, as well as the input and output data.

XR 3: Provide explanations to help assess the influence of internal and external drivers. Drivers denote features that contain information on driving factors crucial for the decision-making process and can stem from internal or external sources. Statistically, this distinction is irrelevant, but from a business perspective the drivers should be separated, as companies usually are only able to adjust internal drivers (such as running marketing campaigns as opposed to the external temperature).

XR 4: Provide explanations to help compare different models. This can include the comparison between different ML algorithms or other models, such as simple average functions.

XR 5: Provide explanations to help assess the applicability of the ML models to different use cases. Use cases can be the same type of forecasting done with other data (such as a different region for a sales forecast) or other figures that could benefit from a similar approach.

XR 6: Provide explanations to help understand and evaluate the predictions generated by the ML model. For this purpose, it should be explained which influencing factors were considered and how the prediction was generated.

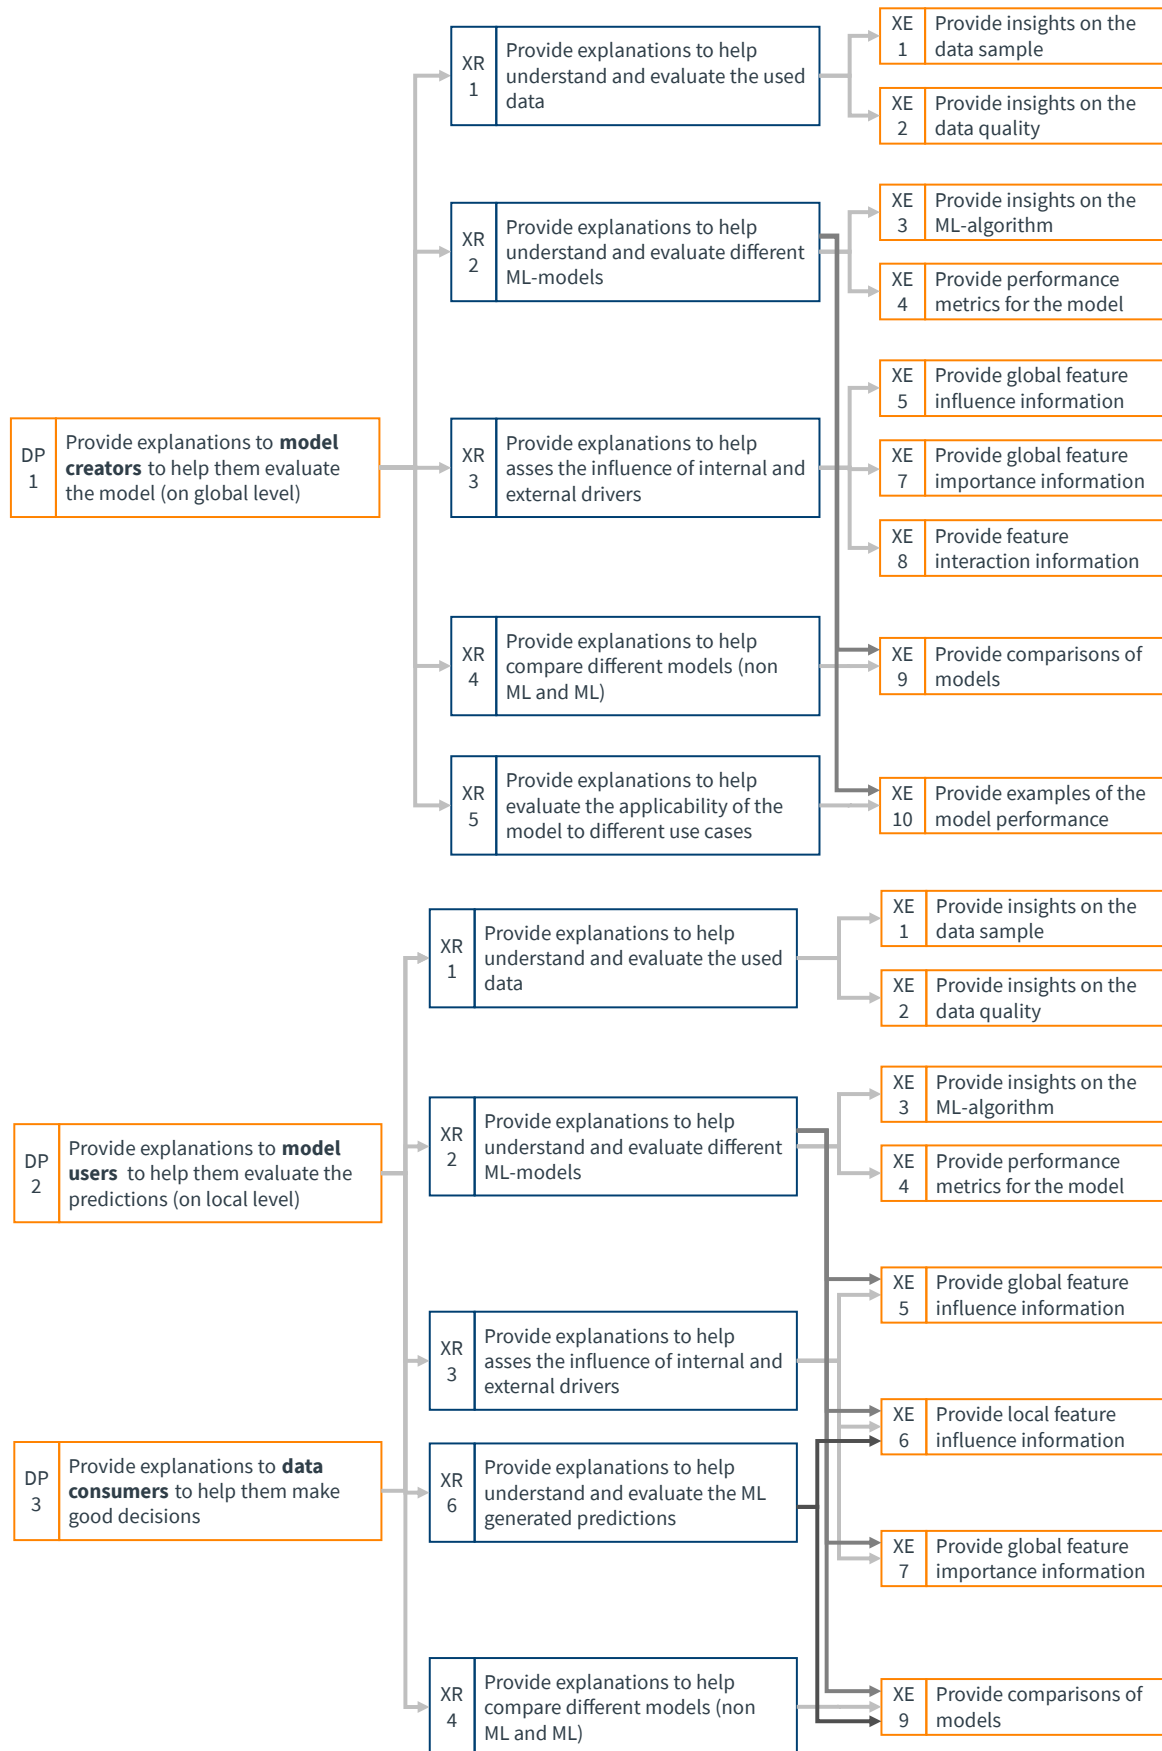


Figure 3. Mapping of XE to XR and DP.

Based on the research on explainability techniques and refinement through our interviews, we furthermore formulate XEs, which can be described as representations and visualizations of generated explanations or other information in the UI. These are based on XAI techniques, but can also provide non-XAI- or even ML-related information. The provided mapping of the XEs to the XRs completes the third decomposition layer of our DPs. The XE are described in the following together with the measures utilized to evaluate how well the XE provides explanations to satisfy their respective XRs:

XE 1: Provide insights on the size of the data sample, i.e., the data amount. This includes information on the number of used time series of the actual figure (e.g., previous sales numbers) and of the internal and external drivers. Further information can be provided on the timeframe of the time series and various metrics concerning the amount of data points (e.g., average amount of data points per timer series). To fulfill XR 1, the XE should help in evaluating the data amount (KM 1.1) and enable decision-making concerning the data amount, such as deciding whether the amount is enough or more data needs to be acquired (AM 1.1). Only for DP 2 and DP 3, to fulfill XR 1, the XE should further help in evaluating the time frame (KM 1.2).

XE 2: Provide insights on the quality of the data. In this case, the quality can be evaluated by the amount of missing data points or anomalies per time series. To fulfill XR 1, the XE should help in evaluating the quality of the used data (KM 2.1) and enable decision-making concerning the data quality, such as deciding to improve the data quality further, exclude certain data series or data points, and many more possibilities (AM 2.1).

XE 3: Provide insights on the general context of how the ML model will be applied. High-level information concerning the input data, output data, and transformation process in the form of an ML algorithm is provided. General information on the ML algorithm can be provided (e.g., the advantages and disadvantages of the used algorithm and a short summary of its functionality). In order to satisfy XR 2, the provided information should help in evaluating the correctness of the context (KM 3.2) and in making decisions concerning the correctness of the context (e.g., changing some of the context or keeping it as is) (AM 3.2).

XE 4: Provide performance metrics for the model. These metrics can vary depending on the use case, but should generally give an indication of how well the model performs, i.e., how well the model predictions match the actual data of testing data set. To satisfy XR 2, the metrics should help in evaluating the model quality (KM 4.2) and in deciding whether the model is usable or needs further refining/should be discarded completely (AM 4.2).

XE 5: Provide information on the influence of all internal and external drivers (features) used in the model on the predictions on a global level. This means that a selected feature's influence on all the results of the model is shown, depending on the expression of this feature. To satisfy XR 3, the presented driver influence should help in evaluating the influence of the drivers in the model (KM 5.3.1) and the driver-dependent biases (KM 5.3.2). Based on the explanations, decisions regarding the correctness of the feature influence can be made (e.g., use the driver, modify the feature selection, or discard the driver completely) (AM 5.3.1). Also, decisions can be made regarding the improvement of the features, such as paying external providers for better data on external drivers (AM 5.3.2). Only for DP 2 and DP 3, to fulfill XR 3, the XE should further help in evaluating the causal relationships of the drivers (KM 5.3.3) and enable decision-making regarding the usability of the model (AM 5.2).



Figure 4. Shown Mock-ups for XE 6.

XE 6: Provide information on the influence of all internal and external drivers (features) used in the model on a selected prediction (local) generated by the model. This can be accomplished via different means, but generally, it should help evaluate why a certain value was predicted based on how the features are expressed. Figure 4 depicts the alternatives shown to interviewees, which compare the predicted values with the actual values in a chart. In the lower half, it shows bar charts on how different drivers influenced the selected data point, as indicated by the shaded column in the above chart. In order to satisfy XR 6, the explanations should help evaluate the local influence of drivers on the selected value (KM 6.6) and enable decision-making regarding the use of the forecast data (AM 6.6). Further, to satisfy XR 2, decisions should be facilitated regarding the usability of the ML model (AM 6.2). To satisfy XR 4 decisions should be able to be made in regard to the correctness of the importance of the drivers (such as using or not using the drivers) (AM 6.3.1) and the improvement of the features, as described for XE 5 (AM 6.3.2).

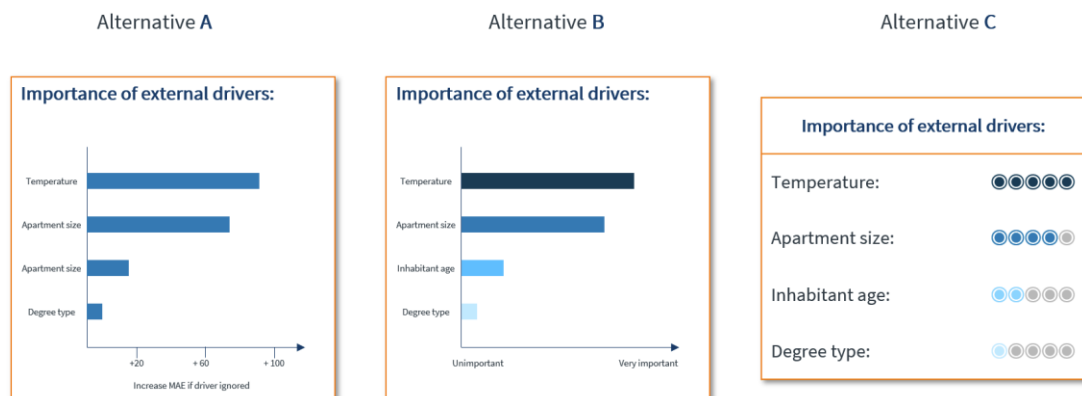


Figure 5. Shown Mock-ups for XE 7.

XE 7: Provide information on the importance of the selected internal and external drivers (features) for the model's performance. High importance means that the model would perform worse if this feature would not have been considered in the model. Figure 5 depicts the alternatives for XE 7, which show the importance of different drivers in a bar chart, with the bar width indicating the loss of a performance metric if the driver were not to be included in the model. To satisfy XR 3, the explanations should help evaluate the importance of drivers to the model (KM 7.3). Decision-making concerning the correctness of the importance of the drivers (such as using or not using the drivers) (AM 7.3.1) and the improvement of the features, as described for XE 5 (AM 7.3.2)

should be possible. Only for DP 1 and DP 2, to satisfy XR 2, decision-making should be enabled in regard to the usability of the ML model (AM 7.2).

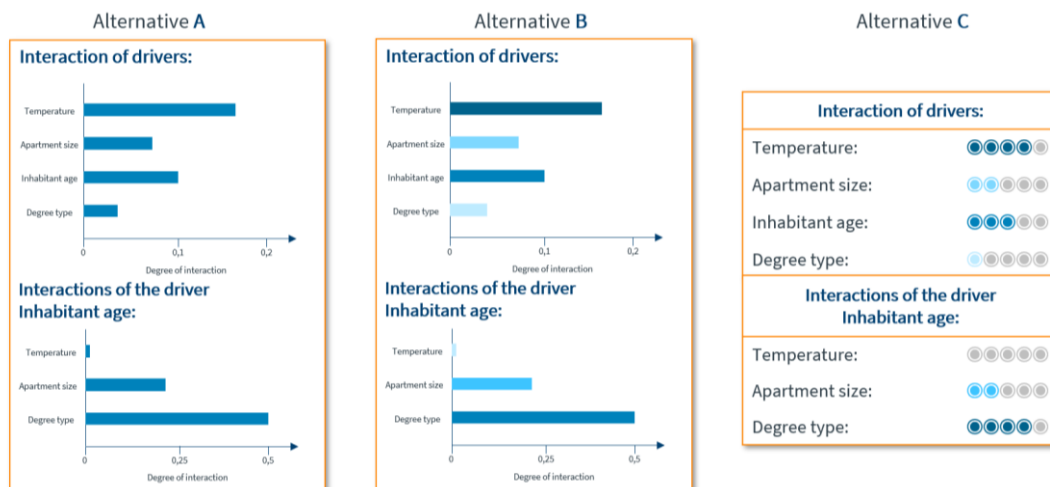


Figure 6. Shown Mock-ups for XE 8.

XE 8: Provide information for the interaction of the selected internal and external drivers (features). A high overall interaction suggests that the features correlate in such a way that they can describe patterns that would not be recognized if the drivers were used standalone. High interactions can therefore indicate the right combination of the features so that the model makes up “more than the sum of its parts”. Figure 6 displays the alternatives for XE 8, with the overall interaction of different drivers shown in the upper half as well as the interaction of a selected driver with the others in the lower half. The interaction is visualized via different variations of bar charts, with the bar width indicating the degree of interaction. To fulfill XR 3, the explanations should help in evaluating the relationship between different drivers (KM 8.3), facilitate decisions regarding the correctness of the drivers (AM 8.3.1) as described for XE 7, and the improvement of drivers (AM 8.3.2) as described for XE 5 and XE 7.

XE 9: Provide comparisons of different models. ML models could be compared with other models or other means of achieving a forecast, such as a simple moving average, which averages the actual data in a moving time window. Also, the actual data can be included to better assess the accuracy of the models. For a better comparison, the focus is on visualizing the forecasts of different models. To satisfy XR 2, the explanations should help evaluate the model quality (KM 9.2) and facilitate decisions regarding the usability of the model (AM 9.2) as described for XE4. To satisfy XR 4, the provided explanations should help compare different models in terms of quality (KM 9.4) and choose the appropriate model (AM 9.4). Only for DP 1 and DP 2, to satisfy XR 6, decision-making should be enabled in regard to the usability of forecast data (AM 9.6).

XE 10: Provide examples of the model's predictions based on different performance classifications. Classifications can be done based on the performance metrics and could show the best, medium, and worst time series in terms of performance. A time series in this case means a sub-dataset of the overall use case where the model is applied, e.g., different regions of a sales forecast. To satisfy XR 2, the explanations should help evaluate the model quality (KM 10.2) and make decisions regarding the usability of the model (AM 10.2) as described for XE 4 and XE 9. To achieve XR 5, the explanations should help evaluate the applicability of the model in the intended area (KM 10.5) and enable decision-making regarding the selection of areas where the model will be applied (AM 10.5).

The following general UX/UI DPs were derived to analyze different guidelines based on the user needs:

UDP 1: Use visualizations where necessary and possible. Especially if data can be illustrated in charts of various forms, it should be.

UDP 2: Utilize color coding to better differentiate numbers or graphics. A good example is the indication of good or bad performance metrics based on an intuitive “street light” classification (Red for bad, yellow for medium, green for good). Other options include intensifying the display color of bars in bar charts.

UDP 3: Abstract and aggregate information if necessary. Examples are the aggregation of big amounts of data into a standardized value or simply displaying continuous values in a normalized five-point illustration.

These were regarded as guidelines subordinate to the importance of the information provided by the XE, so they were only applied when they seemed fitting. The main focus of these principles is to provide good usability. Based on the evaluation of prior research (see Section 5.2), data can be better understood when visualized. Further colors have the potential to make decision makings easier. Lastly, deriving inspiration from lots of web services or e-commerce tools used by lay users, this study proposes the use of aggregated five-point scales, similar to those used in rankings from, e.g., products on shopping sites. One point on the scale indicates a poor expression of the considered value, and five points indicate the opposite.

5.4.2 Evaluation

Each of the XE was discussed and evaluated with the experts using mock-ups of slightly different implementations of the XE as the basis for the discussion (e.g., with and without color-coding of the displayed metrics). For instance, the results of the evaluation of XE 7 (which aims to provide insights concerning the importance of the drivers for the model) are shown in Table 8, whereby the results from the CPM-expert group are in blue colors and the results from the Management group are red. IC 1, IM 2, and IM 3 chose the option with bar charts and no color coding, IC 2 with bar charts and color-coding, and IC 3 as well as IM 1 with five-point scales and color-coding. When asked to evaluate XE 7, IC 2 noted that the measurement of mean absolute error (MAE) increase when withholding the driver in the model as an indicator for the driver’s importance was interesting. All of the general metrics were evaluated with at least “rather agree”, with the exception of IC 1, who only partially agreed with GQ 4. The explainability-specific metrics all had a majority of at least “rather agree” and in the case of XQ 2 and XQ 3 a majority of “strongly agree”. IC 1 and IM 1 only partially agreed with the completeness of the explanations. All interviewees strongly agreed that the explanations help in evaluating the drivers. A majority strongly agreed that decisions based on these explanations could be made in terms of selecting the correct and important drivers. In terms of deciding on driver improvement actions a majority rather agreed. For the decisions concerning the model usability most of the Management-interviewees only partially agreed. IC 2 suggested a different color coding and was missing the data source for plausibility. IM 3 wanted a tooltip that explains what the MAE is. IM 1 mentioned the feature interaction effects that were shown in XE 8 but not to the Management-group.

General feedback to the XE 7 mock-ups included that the color coding was helpful, but red, yellow, and green coloring could be a problem in terms of accessibility, as color-blind people would not be able to interpret them (IC 3). IC 2, as mentioned in most of the XEs individual evaluations, was missing the data sources for more plausibility. IC 2 also stated that because of their statistics and

mathematics background, their evaluations could be a little “nitpicky”. IC 2 suggested providing more explanations for the terms and metrics displayed in the XEs, as people with no statistics or data science background may not be able to understand them. IM 1 generally found all the information helpful but suggested layering it for example in the form of tooltips, so to allow the users to get the detailed information only if they need it. They stated that they mainly want to see an indication if something is good or bad at first glance. IM 2 also liked the XEs and emphasized the potential of combining them. IM 3 was positively surprised by the possibilities of what XAI technologies can accomplish. They also stated that it was difficult to evaluate the data amount and quality XEs because they are hard to understand for lay users. When they saw the following XEs they could also retroactively better understand the data-centric XEs. In order to answer the research questions, the general and XE-specific feedback and the overall evaluation results are discussed in the following.

Table 8. Evaluation Results from the Interviews for XE 7.

Interviewees and chosen alternatives:		IC 1: ● A IC 2: ● B IC 3: ● C			IM 1: ● C IM 2: ● A IM 3: ● A		
		Strongly disagree	Rather disagree	Partially agree	Rather agree	Strongly agree	
GQ 1	Easy interpretation				●●●●	●●	
GQ 2	Intuitive interpretation				●●	●●●●●	
GQ 3	Easy to learn interpretation				●	●●●●●	
GQ 4	Satisfactory interpretation			●		●●●●●	
GQ 5	Transparent presentation				●●●●	●●	
XQ 1	Complete explanations			●●	●●●●●	●●●●	
XQ 2	Plausible explanations				●●	●●●●	
XQ 3	Trust-building explanations				●●	●●●●	
KM 7.3	Evaluate driver importance					●●●●●●	
AM 7.3.1	Decisions driver importance				●●	●●●●	
AM 7.3.2	Decisions driver improvement		●	●	●●	●●	
AM 7.2.	Decisions model usability (IM)			●●		●	

The general and explainability quality metrics were asked for every XE together with the XE-specific knowledge and actionability metrics and were used to validate and refine both every XE on its own (as described in detail for XE 7 above) as well as the overarching UDP. In the following, we first qualitatively discuss the overall evaluation results and then elaborate on further insights that were not directly integrated into the XE, XR, and UDP detailed above. Regarding the evaluation of the explainability elements, it is important to note that because of the small sample size (n=3 in each interview group), the results should only be interpreted on an indicative level. Nonetheless, the XEs were generally well perceived in terms of helping to fulfill the explainability requirement that they were mapped to. In the first interview group, representing the model creators, in at least 85% of the cases the knowledge and actionability metrics, which aimed to

evaluate, whether the XE would help satisfy the XR, were evaluated as at least “rather agree”. In over 50% they were even evaluated with “strongly agree”. Even when XEs were not as well evaluated concerning the generic metrics, the interviewees still felt that they provide the right knowledge and help make decisions. The generic metrics were also rated with at least “rather agree” in over 80% of the cases and “strongly agree” in around 50% of the cases. The second interview group, representing the model users and data consumers, evaluated the XE with similar results. The knowledge metrics were evaluated with at least “rather agree” in over 90% of the cases (over 60% strongly agreed), and the actionability metrics in over 80% of the cases (over 40% strongly agreed). The general metrics also showed good evaluation results of over 85% “rather agree” (over 30% strongly agreed).

In the case of low-rated XEs, interviewees often felt that some information to understand the ML-related metrics or terms was missing, complicating the use by lay users (e.g., IM 1 and 2 regarding XE 1 or IM 1, 2, and 3 in regards to XE 4). In some cases, the visualization of the charts was not optimal (e.g., IC 2 and IM 1 in regards to XE 5 or IC 1 regarding XE 8). One interviewee (IC 3) also emphasized the importance of providing data sources for all calculated metrics to make the explanations plausible. This was also proposed in research by Laato et al. (2022). Because this was not mentioned by the other interviewees, it should be implemented with caution, so as not to overload lay users with information. The possibility to “drill down” as proposed by Laato et al. (2022) to get the requested information on demand should be evaluated. In terms of the chosen mock-up alternatives the conclusion is not as clear: For XE 5, regarding the global driver influence, all interviewees chose the same option, and for XE 9, regarding the model comparisons, 5 out of 6. The rest only had majorities in their respective interviewee groups in 7 out of 9 cases (IC) or 7 out of 8 cases (IM). This further supports the assumption that the design of the XE should be user-group specific. Based on the qualitative comments of the interviewees during the selection and afterward, it should be evaluated how the best features of the different options could be combined. Again, the often-suggested option of tooltips or “drill downs” should be evaluated, to accommodate different levels of details in one UI. Alternatively, things like color-coding which were better received by the Management interviewees could be made configurable in the ML-based system, depending on the user utilizing them. The UX/UI-DPs should therefore be evaluated further, in a context where they could be applied depending on the user’s personal preference. In conclusion, the conceptualized DPs provide the right XEs to the identified XRs, as indicated by the generally favorable evaluation results of the KMs and AMs. While the evaluation results for the general and explainability-specific UX metrics were almost equally as well-rated, there is still some improvement potential for the UI design. More interactable explanations, such as tooltips for ML-related metrics and terms, could offer much potential, as well as the combination of the different XEs at the appropriate steps in the user journey.

5.5 Discussion

In order to contribute relevant knowledge to IS-research, design science studies should fall into one of three categories (Gregor & Hevner, 2013): *Exaptation* means that they extend existing knowledge to new problems. *Improvement* describes studies that develop new solutions to existing problems. *Invention* studies develop novel solutions for new problems. The user-centric approach developed in this work combines, refines, and focuses prior research to develop a solution to provide explainability in ML-based CPM systems. Therefore, it provides a new solution to a known problem, thus classifying it as an *improvement* work.

This study makes two major **theoretical contributions**. First, this work refines prior stakeholder and user groups identified by, e.g., Langer et al. (2021) and Meske et al. (2022) to a more ML model-focused level. In particular, it specializes the user groups for the CPM context, connects them with their respective XR, and thus invites future research in the CPM field to evaluate XAI approaches using our framework for CPM users. Within our framework, we follow Brennen (2020) and take the lay user groups model users and data consumers into account. Additionally, we encourage further studies to refine user groups of XAI systems for other contexts to lay the groundwork for research on the user-centric design of respective XAI approaches. Second, we provide DPs representing UX/UI building blocks that describe what kind of information should be displayed to the user, which can be used by scholars to inform future user-centric XAI research. Our DPs relate directly to our identified user groups following the advice of Liao et al. (2020) to consider user requirements to derive DPs. Through this study, we thus validate existing research on the UX/UI design of XAI approaches (e.g., Zhou et al., 2021; Oh et al., 2018; Laato et al., 2022) for the CPM field and refine CPM-specific design by centering our DPs around the CPM users and their requirements in particular. The positive evaluation by our lay user groups indicates that our XEs are suitable for providing explainability to these users in particular, following the suggestion of Brennen (2020). Moreover, our derived XR can be matched to the XAI questions in the UX-centered research of Liao et al. (2020) with the exception of XR 6, which may therefore be specific to the CPM domain. However, our DPs contribute to XAI research not only in the CPM context. A plethora of similar forecasting use cases exist in other domains such as sales forecasting for supply chain management (e.g., Bi et al., 2022) or wind speed forecasting for power grid balancing (e.g., Yang & Chen, 2019) with model users similar to planners in the CPM context, for instance. As our DPs are tailored to provide explainability specifically to CPM users, we argue that they can apply to any domain with use cases sufficiently similar to the forecasting of business figures and user groups whose characteristics and goals align with those of the CPM user groups (e.g., data consumers that utilize the model predictions to support their decision-making processes, for instance for planning or performance monitoring purposes). In this regard, we invite future research to adopt our evaluation framework to validate and extend our DPs for contexts other than CPM.

Additionally, we make two primary **practical contributions** with our study. First, organizations can utilize our developed evaluation framework to assess their ML-based CPM systems currently in use for their ability to provide explainability to their users. We invite practitioners to rethink whether their CPM systems are truly understood by users, especially lay users, and offer considerations on the requirements and evaluation criteria to be fulfilled to enhance user understanding. Second, our DPs can be used as concrete guidelines to foster explainability in practice when developing ML-based CPM systems by incorporating an XAI approach. Getting more explainability into ML-based solutions may help organizations build trust in the technology in CPM, as they provide knowledge that without them would be lost in the “black box”. By utilizing our user-centric approach, practitioners may therefore address the perceived risk of low explainability (see Chui et al., 2021) and thus foster the use and adoption of ML in the CPM domain.

This study is not without limitations, however. First, although the technical possibilities of posthoc techniques and other aspects of ML or XAI were kept in mind, they were not evaluated in detail. Completely interpretable models were omitted from the start, and their potential for specific and small use cases should be determined, as the development of the conceptualized XEs could potentially come with a technical overhead. Secondly, while this study made a first step toward user-centric XAI design in CPM, its scope is too small for empirical validation of the identified XEs. Further quantitative studies could thus build on our evaluation framework to

refine and validate the XEs with a larger sample of CPM users from organizations of various sizes and industries. Lastly, the evaluation of the results was done via UI mock-ups for standalone XEs. As mock-ups are static, they cannot truly show the UX with interactive elements. Therefore, interactive explanations were largely omitted and could hold the potential for additional explanatory power to be unlocked by future research.

5.6 Conclusion

In this study, we followed the design science research process as presented by Peffers et al. (2006) to create UX/UI designs for the ML-based CPM system developed by a medium-sized supplier of enterprise service management software, which we evaluated and refined through six interviews with CPM and management experts from the firm. As results, we were able to derive DPs, which aim to provide explainability according to the user's goals and requirements in order to facilitate knowledge exchange between users and ML-based CPM systems. To create our DPs, we, therefore, specified the user groups of ML-based CPM systems as model creators, model users, and data consumers and identified their respective goals. The DPs themselves consist of a three-layer decomposition structure. The highest levels (DP 1 – DP 3) describe the goals of the user groups on which the explanations should focus. The second layer derives more specific explainability requirements (XR 1 – XR 6), such as the need to describe data sample size and quality that can be connected to different DPs. Lastly, the explainability elements (XE 1 – XE 10) that deploy techniques from the field of XAI such as feature importance or simply present metadata, such as model performance metrics, to fulfill different XRs constitute the third level. Further, more general UX/UI-DPs are provided (UDP 1 – UDP 3). They suggest the use of visualization, color-coding, and aggregations to point scores. The DPs are validated on an indicative level from interviews, for which we elaborate an evaluation framework. Our framework includes generic metrics concerning the UX/UI design, as well as XE-specific metrics, evaluating whether the right knowledge was provided and the right actions can be taken based on the explanations. The evaluation results show a high agreement of users with the defined quality metrics. Another finding from the interviews suggests that the same XE could be deployed with a user-group-specific design to further enhance usability. Therefore, the DPs developed in our study offer first concrete guidelines for designing XAI approaches in CPM to practitioners while providing scholars with both a CPM-specific evaluation framework and user-specific DPs for future XAI approaches to be refined and expanded on through research in the CPM domain.

Acknowledgment

This research and development project is/was funded by the German Federal Ministry of Education and Research (BMBF) within the “Innovations for Tomorrow's Production, Services, and Work” Program (funding number 02L19C150) and implemented by the Project Management Agency Karlsruhe (PTKA). The authors are responsible for the content of this publication.

6 Research Paper D: Composition of Machine-Learning-Driven Business Models Enabling Value Capture

Title: What Constitutes a Machine-Learning-Driven Business Model? A Taxonomy of B2B Start-Ups With Machine Learning at Their Core

Authors: Vetter, Oliver A.; Hoffmann, Felix; Pumplun, Luisa; Buxmann, Peter

Published in: Proceedings of the 30th European Conference on Information Systems (ECIS 2022)

Abstract

Artificial intelligence, specifically machine learning (ML), technologies are powerfully driving business model innovation in organizations against the backdrop of increasing digitalization. The resulting novel business models are profoundly shaped by ML, a technology that brings about unique opportunities and challenges. However, to date, little research examines what exactly constitutes these business models that use ML at their core and how they can be distinguished. Therefore, this study aims to contribute to an increased understanding of the anatomy of ML-driven business models in the business-to-business segment. To this end, we develop a taxonomy that allows researchers and practitioners to differentiate these ML-driven business models according to their characteristics along ten dimensions. Additionally, we derive archetypes of ML-driven business models through a cluster analysis based on the characteristics of 102 start-ups from the database Crunchbase. Our results are cross-industry, providing fertile soil for expansion through future investigations.

Keywords: Business Models, Machine Learning, Artificial Intelligence, Taxonomy

6.1 Introduction

The rapidly advancing digitalization is leading to more and more data being collected by organizations. In 2020, 64.2 zettabytes of data were generated or replicated worldwide – an amount ten times larger than in 2012, with no end to the growth in sight (IDC, 2021). The high availability of data has also fuelled another technological trend, the growing use of machine learning (ML) to support or automate organizational processes (Jordan & Mitchell, 2015). ML is a technology that can be used to implement instances of artificial intelligence (AI) by learning patterns based on data that can be applied to make predictions (Brynjolfsson & Mitchell, 2017; Mitchell, 1997; Russell & Norvig, 2021). This learning process (i.e., training) is largely independent of human influence and is thus highly experimental in character (Amershi et al., 2019; Choudhury et al., 2021). ML fundamentally offers the potential to significantly change organizational processes and enable business models (BMs) in the business-to-business (B2B) segment that were previously inconceivable. For example, Salesforce utilizes ML in their *Einstein* solution to provide sales and marketing departments with insights and predictions to better understand their customers, drawn from past customer interactions (Salesforce, 2022). At the same time, the characteristics of ML hinder the creation of genuine business value for the organization from this technology (e.g., Burström et al., 2021). Harnessing the power of ML for an organization is therefore difficult to achieve and differs greatly from building BMs based on conventional technologies. Business model innovation (BMI) has always been a demanding and multi-faceted process, but ML exacerbates this challenge by adding another experimental component (Choudhury et al., 2021; Schneider & Spieth, 2013). Current research that might guide the development of ML-driven B2B-focused BMs, however, is in its infancy and is primarily focusing on specific use cases in dedicated domains, such as manufacturing (e.g., Burström et al., 2021).

In order to alleviate some of the complexity of ML-driven BMs and to establish a concise structure to guide researchers and practitioners in BMI, we aim to create an overarching taxonomy of B2B BMs enabled by ML technologies. Since established organizations often pursue multiple BMs whose boundaries become blurred, we focus particularly on start-ups whose core BM is still clearly discernible (Hartmann et al., 2016). More specifically, we focus on B2B start-ups due to the B2B segment's high potential to benefit from ML technologies (MIT Technology Review Insights, 2018). Therefore, in the interest of a meaningful taxonomy on B2B BMs, we exclude BMs in business-to-consumer (B2C) markets due to a variety of differences to B2B BMs (e.g., in customer approach (Iankova et al., 2019), value creation for customers (Grewal et al., 2021), or influence of innovative services (Dotzel & Shankar, 2019)). To reveal how B2B start-ups operate their organization, we ask:

Research question 1 (RQ1): *What are the characteristics of start-up B2B BMs that use ML at their core along different dimensions, and how can they be combined into an overarching taxonomy?*

Research question 2 (RQ2): *What are the ideal-typical archetypes of ML-driven B2B BMs based on recurring characteristics of start-ups that use ML at their core?*

According to Nickerson et al. (2013), taxonomies are artifacts that organize a set of objects according to their characteristics to help researchers and practitioners better comprehend complex domains. In relation to BMs, taxonomies serve to create a high-level abstraction of the BMs' essences. By creating the taxonomy and deriving corresponding archetypes through a cluster analysis, we contribute to structuring the research field of ML-driven B2B BMs. In particular, we highlight relevant dimensions and characteristics by which ML-driven BMs can be

distinguished and thus provide other scholars a starting point to better define the object of organizational research and frame future studies. In addition, we provide practitioners with a clear overview and insights into archetypical BMs that they can use to develop new B2B BMs of their own in a more systematic way. In doing so, we support organizations in general, and start-ups in particular, to better recognize trends in the market, innovate their own business, and differentiate themselves from competitors.

The remainder of the paper is structured as follows: We start by describing research in the area of ML-driven BMs as well as existing taxonomies from the related field of data-driven BMs. We then discuss the methodology we adopted to develop the taxonomy and determine the archetypes of ML-driven B2B BMs. Finally, we present the resulting taxonomy and archetypes, discuss their value for theory and practice, and point to avenues for future research that can address limitations of this work.

6.2 Theoretical Background

In the following sections, we will present a brief overview of BM theory in the context of ML technologies and then report the current state of research on taxonomies related to ML-driven BMs.

6.2.1 Machine-Learning-Driven Business Models

The concept of BMs has received much attention by scholars in various literature streams such as e-commerce, strategy, or innovation management (Zott et al., 2011) and is to date considered a useful perspective for novel insights and further theory building in management literature (e.g., Lanzolla & Markides, 2021; Prescott & Filatotchev, 2021). In essence, a BM is a concept that illustrates the business logic of an organization and depicts how the organization creates and delivers value to customers as well as the associated architecture of revenue, costs, and profits (Teece, 2010). Various definitions of the term have been introduced and discussed in literature (Zott et al., 2011; Al-Debei & Avison, 2010; Casadesus-Masanell & Ricart, 2010). For the purpose of this research, we adopt the definition by Osterwalder and Pigneur (2010), which states that “a business model describes the rationale of how an organization creates, delivers, and captures value” (p. 14). BMs are often conceptualized through the components that constitute them, e.g., the value proposition or the revenue stream (e.g., Al-Debei & Avison, 2010; Teece, 2010; Zott & Amit, 2010; Remane et al., 2016). Many of the BM components described in literature can be categorized in four types of components present in most BM conceptualizations (Burkhart et al., 2011): *Offering factors*, describing how the organization creates value for stakeholders; *Market factors*, detailing for whom value is created; *Internal capability factors*, describing activities and competences of the organization, and *economic factors*, including all economic-related aspects of the organization.

In information systems (IS) research, the BM concept is seen as the missing link connecting business strategy, processes, and information technology (IT) (Veit et al., 2014). In turn, the recent trends of increasing availability of relevant data and technological advances in data analysis carry the potential to profoundly change existing BMs of organizations in the future (Veit et al., 2014) by forcing them to adapt their BMs to survive against globalized competition (Hanelt et al., 2015). While most organizations will certainly benefit from data and data analytics, some BMs go one step further and utilize data as their *key* resource, eventually becoming data-driven BMs (Schüritz & Satzger, 2016). For the term *data-driven BM*, we adhere to the definition of Hartmann et al.

(2016) as “a business model relying on data as a key resource” (p. 1385). AI, as “the science and engineering of making intelligent machines” (McCarthy, 2007, p. 2), offers the opportunity to leverage additional potential in the context of digitalization (Dingli et al., 2021). The most popular technology to realize AI systems is ML, which uses learning algorithms to derive patterns from observed data and saves them in ML models, which in turn can be used on new data to solve problems (Russell & Norvig, 2021). As ML is the foundation of most modern AI systems (Jordan & Mitchell, 2015; Brynjolfsson & Mitchell, 2017), we use the term ML to refer to ML-based instances of AI for terminological clarity throughout this paper. According to Hahn et al. (2020), ML-driven BMs are a subgroup of the previously introduced data-driven BMs due to their reliance on data, yet can be differentiated from the latter as they rely on ML as self-improving technology to draw applicable patterns from the data. The authors thus conceptualize a BM as ML-driven if it utilizes ML technologies in at least one of its BM components. Modern ML technologies urge organizations to reshape or develop entirely new BMs (Lee et al., 2019; Wamba-Taguimdje et al., 2020), as this technology significantly differs from other digital technologies and poses new challenges for organizations alongside diverse opportunities: First, ML technologies can complement, constrain, or substitute for humans at work (Murray et al., 2021). Second, being capable of human feats such as conversation, they blur the boundary between human and machine capabilities (Schuetz & Venkatesh, 2020). Third, the data-based learning approach not only renders ML technologies more complex and thus inscrutable but can also result in unexpected outcomes (Benbya et al., 2020a). Given these striking differences from other digital technologies (Benbya et al., 2021) and considering that BMs will undergo ML-induced transformation (Burström et al., 2021), we argue that there is a need to study these new BMs driven by ML technologies.

6.2.2 *Taxonomies of Data-Driven Business Models*

Taxonomies are a widely used tool to analyze and represent complex systems and their interrelationships in a structured way. Through the holistic disclosure of the components of the system and its properties, different manifestations can be classified and compared with each other (Nickerson et al., 2013). As literature on ML-driven BM taxonomies is scarce, we draw on the broader field of data-driven BMs to identify transferable aspects. Table 9 summarizes the general and industry-specific taxonomies that exist to date. The illustration is based on the categorization by Dehnert et al. (2021), which we expanded to include new publications on data-driven BM taxonomies and the only taxonomy found on the subject of ML-driven BMs (marked in cursive). Among the added taxonomy papers are Woroch and Strobel (2021), who develop a data-driven BM in the context of the Internet of Things and Weking et al. (2020), who establish a taxonomy of industry 4.0 BMs enabled by the Internet of Things and smart factories among others. The taxonomy of Baecker et al. (2021) examines data-driven value creation within organizations and focuses on the required underlying data, the gained business value, and the approach to create it. Lastly, Anton et al. (2021) present the first available taxonomy of ML-driven BMs with a focus on start-ups in the energy sector. By examining how ML technologies shape BMs in this sector, they contribute to a better understanding of organizations that are implementing ML-driven BMs as the energy sector continues to transform.

Table 9. Classification of Existing Taxonomies Adapted From Dehnert et al. (2021).

	General	Industry-specific
Data-driven business models	<p><i>Baecker et al. (2021);</i> <i>Dehnert et al. (2021);</i> <i>Woroch and Strobel (2021);</i> Passlick et al. (2021); Bock and Wiener (2017); Naous et al. (2017); Hartmann et al. (2016); Engelbrecht et al. (2016); Schroeder (2016); Schüritz and Satzger (2016)</p>	<p><i>Manufacturing: Weking et al. (2020);</i> Logistics: Möller et al. (2020); Manufacturing: Müller and Buliga (2019); Urban: McLoughlin et al. (2019); E- Commerce: Dorfer (2016)</p>
ML-driven business models		<p><i>Electric Power Industry: Anton et al. (2021)</i></p>

The preceding discussion shows that existing taxonomies cover primarily data-driven BMs. However, as mentioned in the previous section, these can not be immediately applied to ML-driven BMs, as those utilize self-improving ML technologies (Hahn et al., 2020) that affect organizations differently (see Benbya et al., 2021). Nevertheless, only Anton et al. (2021) have studied ML-driven BMs to date, albeit specifically for the electric power industry. To the best of our knowledge, there is no cross-industry and thus universally applicable taxonomy for ML-driven BMs yet – a gap we aim to bridge.

6.3 Methodology

To develop the taxonomy of ML-driven BMs, we utilize the development approach proposed by Nickerson et al. (2013). The method is well-accepted in IS research, having been used by several researchers for taxonomies in related fields (e.g., Anton et al., 2021; Dehnert et al., 2021; Möller et al., 2019; Remane et al., 2016). Taxonomies comprise a set of dimensions that, in turn, contain characteristics that can describe the objects under study (Nickerson et al., 2013). As a first step in the taxonomy development process, Nickerson et al. (2013) suggest specifying a meta-characteristic as the “most comprehensive characteristic” (p. 343) that should reflect the purpose of the taxonomy, from which all other characteristics can be derived logically. Next, Nickerson et al. (2013) propose an iterative process to add, change, or subtract dimensions and characteristics during each iteration. The iterations can either be carried out as empirical-to-conceptual or as conceptual-to-empirical approaches. In the former, researchers analyze a subset of objects – such as real-world start-ups – to obtain their characteristics and group them into dimensions. Researchers following the conceptual-to-empirical approach conceptualize the dimensions and characteristics based on the researchers’ knowledge and existing literature and then examine real-world objects to revise the taxonomy. The iterative development process ends when all predefined ending conditions are met after an iteration. After four iterations of taxonomy development, we further conduct a cluster analysis to derive archetypal BMs from the studied start-ups based on their identified characteristics.

6.3.1 Taxonomy Development

As previously stated, the meta-characteristic defines the purpose of the taxonomy, and for this research, we determined it as *distinguishing elements of B2B-focused, ML-driven BMs*. This wording best specifies our goal to identify the different essential components to reveal the core business logic behind ML-driven BMs while at the same time discerning distinctions between different instances found in reality. We decided to adhere to the eight objective (e.g., no new dimensions and characteristics added in the last iteration, at least one object is classified under every characteristic, all objects or a representative sample of objects has been examined) and five subjective (e.g., conciseness, robustness, extendibility) ending conditions proposed by Nickerson et al. (2013).

We chose the conceptual-to-empirical approach for our first iteration. Because little research on ML-driven BMs was available, we turned to literature on general BMs as a starting point. A large volume of published works describes possible configurations of BMs (e.g., Al-Debei & Avison, 2010; Hedman & Kalling, 2003; Osterwalder & Pigneur, 2010). As the Business Model Canvas by Osterwalder and Pigneur (2010) contains the majority of BM components discussed in literature (Passlick et al., 2021) and is additionally well-regarded in practice, we chose it as a starting point for the taxonomy. We drew from literature on data-driven and ML-driven BMs (see Table 9) to select the dimensions from the Business Model Canvas promising the highest discriminatory power for ML-driven BMs as starting dimensions, namely *value proposition* (which we split into *value promise* and *key offering* in the third iteration; see Möller et al., 2019), *customer segment*, *channel*, *key resources*, *key activities*, *revenue stream*. We supplemented these dimensions with aspects from data-driven BMs that are transferable to ML-driven BMs: We specified the dimension *key resources* into the more ML-relevant dimensions *data source* and *data type* (Engelbrecht et al., 2016; Möller et al., 2020; Azkan et al., 2020; Hartmann et al., 2016). Similarly, we specified the dimension *channel* into the most relevant aspect for distinguishing ML-driven BMs: The *deployment channel* (Passlick et al., 2021).

Since the taxonomy should not only be academically motivated but also consider emerging ML-driven BMs in practice, we conducted the following three iterations in compliance with the empirical-to-conceptual approach. We employed the Crunchbase database for our data collection and searched it for suitable start-ups (Crunchbase, 2021). We focused exclusively on start-ups because the available population is larger compared to established organizations, and start-ups presumably have purer BMs that are not hampered by old legacy systems (Hartmann et al., 2016). In particular, start-ups possess only a single or small number of BMs, facilitating their analysis (Sabatier et al., 2010). Our search terms to browse Crunchbase were *machine learning* and *artificial intelligence*, which we used to search the start-ups' tags as well as their short descriptions. We included the latter search term, since start-ups often operate under the more general buzzword *artificial intelligence* when referring to ML-based AI. We wanted to find start-ups that struck a balance between being young enough to pursue a singular BM still while being old enough that we could omit organizations that went bankrupt quickly after launch. Therefore, we focused on start-ups founded in 2018 and 2019. Our search yielded a total of 2,057 start-ups. Due to the large dataset size, we followed the recommendations by Nickerson et al. (2013) and randomly selected subsamples for each iteration (see Möller et al., 2019; Möller et al., 2020). We removed and replaced all start-ups that went bankrupt, did not realize an ML-driven, B2B-focused BM, or did not have sufficient information in German or English on their homepages (see Remane et al., 2016; Täuscher & Laudien, 2018). The subsample sizes are 22 for the second and 40 for the third and fourth iteration, resulting in a data set of 102 start-ups, an excerpt of which can be found

in Appendix 1. Consistent with Hunke et al. (2019), we started with a smaller number of start-ups in the second iteration and subsequently included more entities, as we wanted to roughly identify dimensions and characteristics of ML-driven BMs first, while relying on more information to refine and elaborate the taxonomy in more detail in the later iterations. In each iteration, we analyzed the ML-driven BMs of the subsample for their characteristics. In particular, we checked whether these characteristics were consistent with the previously found taxonomy or whether additions, revisions, or deletions of characteristics or dimensions would improve the usefulness of the taxonomy. We gathered the required information about the BMs in our dataset from publicly available sources such as the start-up's website, articles, blog entries, or other online presences. Because "gross elements of business models are often quite transparent" (Teece, 2010, p. 179), a start-up's BM can be inferred using such reliable public sources (see Hartmann et al., 2016; Möller et al., 2019). We employed multi-researcher triangulation to ensure a high degree of objectivity (e.g., Hsieh & Shannon, 2005; Flick, 2004a). As few start-ups disclosed data on their revenue streams or utilized ML form, we derived the information for the corresponding dimensions with the help of pertinent literature and validated them empirically with the start-ups on which data was available. The taxonomy was finalized after the fourth iteration of the development process, as all ending conditions were met.

6.3.2 Cluster Analysis

We performed a cluster analysis on our dataset to derive information on which archetypes of ML-driven BMs commonly appear in practice. Cluster analysis seeks to form groups of objects based on their similarities (Bailey, 1994), thus in our case, assembles the start-ups into archetypal BMs based on their similarity along the dimensions of the taxonomy. Regarding the design of our cluster analysis, we followed preceding research (Remane et al., 2016; Möller et al., 2019; Anton et al., 2021) and carried out the two-step procedure of Punj and Stewart (1983). The first step consists of the agglomerative hierarchical clustering algorithm of Ward's minimum variance method (Ward, 1963). The procedure starts with every object being a separate cluster and then iteratively merging the two closest clusters based on the calculated distance between them (Eszergár-Kiss & Caesar, 2017). We used Euclidean distances as distance metric, as it is suitable for binary variables and Ward's method is well defined for Euclidian distances (see Rencher, 2002). The results of Ward's method show that either a 6 or 7 cluster solution is optimal for our dataset. In the second step, we used the k-means partitional clustering method. The k-means algorithm finds a partition that minimizes the sum of squared distances between the empirical mean of each cluster and the objects in the respective cluster for an a priori defined number of clusters (Jain, 2010). We chose the 7 cluster solution for our final archetypes because it outperformed the 6 cluster solution both on the elbow curve and the Davies-Bouldin index (Davies & Bouldin, 1979). We implemented the data preparation and the k-means clustering in RapidMiner Studio, and the Ward's method in Python using the library SciPy.

6.4 Results

In this section, we present our final taxonomy and the derived BM archetypes.

6.4.1 Final Taxonomy

The final taxonomy is shown in Table 10 and consists of ten dimensions, which in turn contain three to six different characteristics. Each of the 102 start-ups in our dataset is described by at

least one of the characteristics in each dimension. Following Nickerson et al. (2013), only dimensions in which ML-driven BMs differ are included in the taxonomy, as characteristics that are identical among all ML-driven BMs are of little use to a taxonomy due to their lack of discriminatory power (Anderberg, 1973). To be able to represent the large variety of BMs from our industry-overarching dataset, our dimensions allow BMs to exhibit more than one characteristic (see Hunke et al., 2019; Möller et al., 2020). The following paragraphs describe the identified dimensions and characteristics in depth.

Table 10. Final Taxonomy, Visualized as a Morphological Box.

Dimensions	Characteristics					
Value promise	Cost and time reduction		Quality increase		Insight increase	Innovation increase
Key offering	Aggregation & filtering	Information enrichment	Detection	Optimization	Forecasting	Generation
Client influence on ML system	No influence	Selection of settings	Feedback loop	Development of model	Ownership of model	
Customer segment	Primary sector		Secondary sector	Tertiary sector	Quaternary sector	
Key activity	Consulting	Data science	Data sourcing & engineering	Software engineering	Hardware development	
ML form	Supervised learning		Unsupervised learning		Reinforcement learning	
Deployment channel	Edge	On-premise software	Hosted software		Plug-in	
Data source	Client data		Provider data		Publicly available data	
Data type	Structured		Semi-structured		Unstructured	
Revenue model	Pay-with-data	Subscription	Pay-per-X	Gain sharing	One-time fee	

Note. To improve readability, we have removed the characteristic “unspecified.”

The dimension **value promise** describes what type of value the BM creates for its clients and can take four different characteristics. Hereafter, the term *client* denotes an organization utilizing the services of an ML-driven BM. BMs characterized by the first characteristic *cost and time reduction* either replace human labor for menial tasks or assist humans in their work, thus allowing them to complete workflows more quickly or cost-efficiently. The *quality increase* characteristic denotes BMs that provide effectivity improvements for their client’s products, services, or processes, usually by supplementing them with some form of intelligent behavior. These BMs aim to modify their clients’ services to deliver better results or to offer additional, previously impossible or infeasible features, e.g., enriching video material with ML-generated metadata. Clients of BMs with the characteristic *insight increase* are supplied with ML-generated knowledge derived from data and designed to improve the client’s decision-making process either through faster or more informed decisions. The type of use case thereby determines the content of the provided information: Performance metrics calculated through ML methods might support clients in management decisions, while ML-created risk estimations might aid real estate investors in finding trustworthy debtors in their day-to-day operations. Lastly, *innovation increase* describes BMs with ML systems aiding the client in exploring previously uncharted territory. Providers aim to improve their clients’ search for innovation or novel inventions. An example from our dataset

is ML supporting pharmaceutical companies in their drug discovery process by suggesting possible solutions or identifying gaps in pre-existing knowledge.

Another critical dimension for ML-driven BMs is the **key offering**, or in other words, the type of service they provide to their clients to create the previously described added value. *Aggregation & filtering* organizations provide their clients with an ML solution that analyzes large amounts of data, omits irrelevant data, and condenses the essential information into meaningful output values for the client organization. An example would be a system that screens job applications and highlights each candidate's most relevant experiences for the position. Solutions with the *information enrichment* characteristic also analyze data but aim to expand the given data(set) with supplementary data. They either extend unstructured data with structured data (e.g., analyzing clinical images and displaying additional diagnostic information) or integrate information from complementary sources into the system (e.g., an ML system that crawls the social media sites of job applicants and derives their trustworthiness). Furthermore, the characteristic *detection* describes systems that continuously monitor data streams and alert the client when certain patterns or suspicious activities are detected. Prominent examples include credit fraud detection systems that raise alarms when credit cards are used irregularly or ML-based visual inspection systems that call attention to defective products. BMs with *optimization* offerings apply their ML systems to solve specific, well-defined problems that have a clear desired output but are challenging to solve with conventional methods. They may involve traditional optimization problems like scheduling or vehicle routing problems. However, the ML systems might also attempt to find the ideal candidate for a given job position or optimize the bidding process for E-commerce advertisements by only bidding when individual clients are likely to buy (minimizing costs while maximizing the likelihood of sale). As the name suggests, *forecasting* organizations offer their clients glimpses into the future. They attempt to predict future states of given dependent variables by incorporating large amounts of data. Examples include predictive maintenance solutions which aim to predict when equipment will need maintenance or renewable electricity generation forecasting, which calculates future power production by analyzing weather data. Organizations represented by *generation* offer ML systems for tasks with high degrees of freedom that use input data to independently create complex, context-specific output that resembles the solution a human might have conceived. Chatbots are a prime example of a generation ML solution, answering user queries with solutions relevant to the user's interests. Another exemplary generation system performs automated legal document generation, creating contracts based on tabular input data.

The **client influence on ML system** dimension differentiates BMs based on the extent to which they individually adapt their ML systems to satisfy their clients' needs, which is a double-edged sword. BMs can allow their clients different degrees of contribution, but the more influence clients have, the more difficult it generally becomes to scale the BM. This is because higher degrees of influence usually necessitate that ML models are (re-)trained individually for each client, which increases the effort of each sale. Possible characteristics start with clients having *no influence* at all. These clients either possess no ML knowledge of their own, do not want to draw on their ML resources, or implement a use case that can be fulfilled with a one-size-fits-all solution. With *selection of settings*, the clients still have minimal influence on the ML system but can alter certain predefined settings to cause changes in the system (e.g., a formality setting for a chatbot to adapt the system to different use cases). *Feedback loop* means that clients can evaluate the output of the ML system and feed their evaluation back into the system, which in turn learns from the additional data and corrects itself over time. To have more control over the finished system, clients might opt for a *development of model* organization. These organizations include their clients in the

development process either through a joint team, regular interactions, or platforms simplifying ML development. Lastly, organizations with the *ownership of model* characteristic hand the finished ML model over to their clients, who gain full access and can analyze the model, improve, or re-train it for other use cases.

The dimension **customer segment** records the economic sector in which the target clients of the BM are allocated. According to the definition of Kenessey (1987), which is widely used today, a distinction can be made between four sectors: The *primary sector* supplies raw materials for a product and includes, among others, the harvesting of wood in forestry, fishing, or the generation of hydropower. The *secondary sector* is responsible for processing raw materials from the *primary sector*. It thus includes the manufacturing industry, craft production, and the energy industry, among others. The *tertiary sector* is comprised of all services provided by private organizations or government institutions such as transportation services, utilities, and wholesale or retail trade. The concept of the *quaternary sector* has gained in importance in the context of the transformation to an information society and subsumes all industries that deal with creating, processing, and selling information (data or knowledge). These include IT services and communications technology. Which sector a BM's clients stem from greatly influences their available data and IT infrastructure.

The dimension **key activity** describes “the most important things a company must do to make its business model work” (Osterwalder & Pigneur, 2010, p. 36). Naturally, the vast majority of ML-driven BMs require *data science* as key activity. However, some organizations get by with minimal data science activities, for example, when the applied ML models are already very mature, like it is the case for computer vision solutions. *Consulting* indicates that conveying ML-related knowledge in close contact with the clients is essential for the BM. *Data sourcing & engineering* characterizes BMs that spend much time on gathering, curating, and supplying data in order to provide their services, as it might be the case for organizations offering insights on financial markets for which they need carefully curated data. The key activity *software engineering* is assigned to organizations whose ML solutions are embedded into highly complex software that must naturally be developed and maintained. Similarly, *hardware development* describes organizations that rely on and must develop complex physical devices to execute the output of their ML system, with computer-vision-powered robots being one example.

Organizations can primarily apply three **ML forms** in their BM (Russell & Norvig, 2021), each with distinct capabilities and uses, as well as unique requirements regarding expertise and development. *Supervised learning* systems are given sets of input-output pairs and then learn a function that predicts the appropriate output, or label, when given new inputs. In *unsupervised learning*, the machine learns to find patterns in the input data without being given any explicit feedback. Lastly, in *reinforcement learning*, the system performs certain actions and is then given either rewards or punishments as feedback, which it uses to learn which actions lead to more rewards and alter its actions accordingly.

How ML-driven BMs deliver added value through their products or services generally differs between four **deployment channels**. The channel *edge* means, that the BM's ML system is run on physical devices that are often supplied as a product package. A Chatbot that is implemented on a special tablet for direct in-store customer interaction would be an example for deployment via edge. *On-premise software*, on the other hand, denotes ML systems that are run on the client's network hardware, often on servers or in their cloud. Conversely, *hosted software* is executed on the BM's hardware (thus increasing cloud provisioning costs), with their clients gaining access through a website or application programming interfaces (APIs). Lastly, *plug-in* ML solutions

integrate seamlessly into pre-existing software or platforms, with an example being a human agent augmentation plug-in for a contact center platform.

The data required for running a BM's specific ML solution can stem from three different **data sources**. *Client data* means the client either has pre-existing datasets with the information required or records new data to be used in the system. The ML models are thus either fully trained on client data or come as pre-trained models and are re-trained with client data. In contrast, BMs denoted by *provider data* sell their ML system along with their own supply of data for training and running the model. Furthermore, models can also utilize *publicly available data*, which any interested party can acquire from data platforms such as Kaggle (2021), from data vendors, or from other public sources.

ML systems can require data in many different **data types**, which can be subsumed under three major categories (Sint et al., 2009; Abiteboul et al., 1999). *Structured data* denotes any type of data with an underlying structure, such as tabular data in a database. *Unstructured data*, on the contrary, does not exhibit an identifiable structure and includes images, video, audio, and free text. *Semi-structured data* has no separate, explicit description of its structure, yet it does demonstrate some structure within the data (e.g., e-mails consisting of subject, sender, and text). The type of data carries many implications for a BM: Each type necessitates different kinds of expertise within the BM and requires different preprocessing efforts for value creation.

The **revenue model** dimension depicts how the BM generates revenue in order to cover costs and thrive as an organization. Since many start-ups withheld information on their revenue model until after a demo or sales talk, the characteristics are based on Schüritz et al. (2017), Osterwalder and Pigneur (2010), and Hartmann et al. (2016) and validated with cases from our dataset. In the case of *pay-with-data*, there is no cash flow from the client to the provider; instead, the client gives the provider access to their data in return for the service. Said data can then either be sold by the BM or be used to re-train existing or train future ML models. The characteristic *subscription* is assigned to organizations whose clients must pay a monthly fee to gain access to their services. The subscription rate may vary depending on the service level selected, with the possibility of offering a basic version of the service free of charge in a Freemium model. *Pay-per-X* denotes revenue models where the clients pay a dynamic fee based on performance measurements. These measures can range from the amount of input or output data requested to the occupied computational resources and can also include the number of utilized billable hours. Dynamic fees are also incorporated in *gain sharing* models; however, in this model, they are directly dependent on monetary success measures of the ML system. An exemplary fee, in this case, is a commission that depends on the value of a mediated contract (e.g., between employers and employees). Lastly, in *one-time fee* models, clients pay one time for the ML system and associated services (e.g., including maintenance services for the first few years).

6.4.2 Business Model Archetypes

Our cluster analysis grouped the 102 start-ups of our dataset into seven clusters that each contained 12 to 19 start-ups. Each cluster has a centroid for every characteristic, representing the distribution of BM characteristics in the respective archetype, depicted in Table 11. As characteristics are deliberately not mutually exclusive, the percentages in each dimension do not add up to 100%. Instead, they show how many start-ups in the archetype exhibit each characteristic. We omitted the characteristic *unspecified* from the analysis as it does not contribute to an archetype's distinctness. Due to the high amount of unspecified ML forms (83%) and

revenue models (75%), we consequently decided to omit these dimensions as well (see Möller et al., 2019). Table 11 also shows the consistency between the Ward's method and the k-means clustering for each archetype (see Anton et al., 2021; Möller et al., 2019). We analyzed the cluster centroids and validated them with the BMs contained in each archetype. The resulting descriptions for each archetype are presented in the following paragraphs and illustrated with archetypical examples from our dataset.

Intelligence for services: Start-ups in this archetype provide opportunities to integrate ML-enabled functionalities and intelligent behavior into their clients' services. Clients thus benefit through their own services achieving superior results. Due to the focus on ML-driven enhancement of services, clients of this archetype typically stem from the tertiary or quaternary sector. The most common key offerings, information enrichment and optimization, are each used by 32% of start-ups in this archetype to improve the services of clients – however, which key offering a BM chooses largely depends on its client's services, with cases of all key offerings existing in the sample. The majority of ML solutions in this archetype use unstructured data, with 53% of start-ups utilizing their client's data, 26% supplying their own data alongside their ML system, and 21% using publicly available data. An example for start-ups of this archetype is Bidnamic (2022), whose ML system supports retailers selling via search engine by calculating optimal prices for each individual product and search term.

Automated sensing: This archetype contains BMs with ML systems that can interpret unstructured data quicker or cheaper than humans can. Often, their solutions are computer vision systems that might be deployed on a physical visual sensor, recording their surroundings and analyzing the gathered data. The conducted analysis depends on the key offering, with most systems (71%) focusing on information enrichment – extracting additional information from the data, and passing it on for further use. If BMs deploy their ML systems on complex edge devices, they must perform the necessary hardware development as key activity (21%). Not all start-ups in this cluster are computer-vision-based, but they all process a type of unstructured data. EAIGLE (2021), as archetypical BM, provides a computer vision solution for automated visitor sign-in, including visitor health screening.

Robotic process automation: Robotic process automation uses software to imitate repetitive tasks that would otherwise be carried out by humans (Santos et al., 2020). Start-ups in this archetype automate routine workflows of their clients with ML systems to achieve cost and time reductions. The provided key offering depends primarily on the respective business process being automated, with detection (38%), optimization (31%), and aggregation & filtering (31%) being most common in our sample. In 54% of our cases, the data required for automating the workflows are structured data. Additionally, 62% of start-ups in this archetype perform extensive data sourcing & engineering tasks to reduce human labor as much as possible. In our dataset, start-ups in this archetype mainly serve the secondary and tertiary sectors. One example of such start-ups is Circuit Mind (2021), whose ML system automatically selects components and generates possible circuit schematics for electronics.

Table 11. Distribution of the Start-Ups' Characteristics Within Each Archetype.

Dimensions	Consistency (algorithms):		73.68%	64.29%	53.85%	66.67%	68.75%	91.67%	68.75%
	Characteristics	Sample distribution	Intelligence for services	Automated sensing	Robotic process automation	ML development partner	Constructive assistant	Internal business diagnostics & prediction	Environmental diagnostics & prediction
		n	102	19	14	13	12	16	12
Value promise	Cost and time reduction	47%	5%	100%	100%	33%	81%	17%	6%
	Quality increase	20%	89%	0%	15%	0%	0%	0%	6%
	Insight increase	39%	5%	0%	15%	67%	6%	100%	100%
	Innovation increase	5%	5%	0%	8%	8%	13%	0%	0%
	Unspecified	2%	0%	0%	0%	17%	0%	0%	0%
Key offering	Aggregation & filtering	38%	16%	0%	31%	42%	19%	83%	88%
	Information enrichment	23%	32%	71%	23%	8%	0%	17%	6%
	Detection	15%	11%	21%	38%	8%	0%	0%	25%
	Optimization	18%	32%	14%	31%	33%	6%	8%	0%
	Forecasting	15%	11%	0%	15%	17%	0%	50%	19%
	Generation	24%	21%	0%	23%	8%	100%	0%	0%
	Unspecified	3%	0%	0%	0%	25%	0%	0%	0%
Client influence on ML system	No influence	63%	79%	64%	77%	25%	38%	83%	69%
	Selection of settings	17%	11%	21%	8%	17%	25%	8%	25%
	Feedback loop	3%	0%	14%	0%	0%	6%	0%	0%
	Development of model	10%	5%	0%	0%	42%	19%	0%	6%
	Ownership of model	5%	0%	0%	15%	25%	0%	0%	0%
	Unspecified	5%	5%	0%	0%	8%	13%	8%	0%
Customer segment	Primary sector	6%	0%	29%	0%	0%	0%	8%	6%
	Secondary sector	14%	11%	0%	54%	0%	6%	8%	19%
	Tertiary sector	38%	53%	14%	46%	0%	50%	67%	31%
	Quaternary sector	13%	26%	21%	8%	0%	0%	8%	19%
	Unspecified	33%	21%	36%	0%	100%	44%	8%	31%
Key activity	Consulting	7%	5%	0%	0%	42%	6%	0%	0%
	Data science	84%	95%	79%	85%	75%	88%	58%	100%
	Data sourcing & engineering	26%	16%	7%	62%	0%	13%	0%	81%
	Software engineering	40%	21%	50%	38%	25%	69%	83%	6%
	Hardware development	3%	0%	21%	0%	0%	0%	0%	0%
Deployment channel	Edge	10%	5%	36%	15%	0%	6%	0%	6%
	On-premise software	22%	37%	29%	23%	25%	6%	25%	6%
	Hosted software	47%	26%	29%	46%	33%	69%	67%	63%
	Plug-in	15%	11%	7%	15%	0%	44%	0%	19%
	Unspecified	16%	21%	7%	15%	42%	6%	8%	13%
Data source	Client data	74%	53%	86%	77%	92%	75%	92%	56%
	Provider data	17%	26%	14%	15%	8%	25%	8%	13%
	Publicly available data	25%	21%	0%	23%	17%	6%	25%	75%
	Unspecified	4%	11%	0%	0%	8%	6%	0%	0%
Data type	Structured	26%	11%	0%	54%	33%	0%	92%	19%
	Semi-structured	4%	0%	0%	0%	8%	6%	8%	6%
	Unstructured	49%	79%	100%	8%	0%	75%	0%	50%
	Unspecified	23%	16%	0%	38%	58%	19%	0%	31%

0%



100%

Distribution of characteristics

ML development partner: ML development partners work towards providing their clients with user-friendly access to the technology of ML. They can offer access to a variety of ML services and aid in developing ML systems that are specifically tailored to each client and their data. To achieve this high degree of individuality, BMs either provide a platform that does not require extensive ML-specific expertise for clients to develop their own ML models, or BMs conduct a collaborative development process. In the latter case, the start-up is often engaged in extensive consulting activities (42%) and can grant its clients different degrees of influence on the ML system as desired. Due to the focus on providing the technology without a specific business context, the target group of these start-ups is not limited to a specific sector. An archetypical example is AltaML (2021), with their ML experts working together with clients to realize ML solutions for various use cases.

Constructive assistant: All organizations in this BM archetype offer generation ML solutions (100%), most of which are designed to achieve cost and time reductions for client organizations (81%). These ML systems are usually given instructions through unstructured data (75%) and then aim to create output in an unstructured data format that emulates the way a human would have completed the task. Chatbots are a prominent example of such solutions. Many use cases for these ML systems exist in cross-sectional functions or functions with contact to end consumers. Consequently, they are primarily offered to clients in the tertiary sector (50%) or are not sector-specific. Lastly, users of constructive assistant solutions should not require experience with ML systems, so extensive integration efforts into comprehensive software systems through software engineering are often necessary (69%). These ML systems are delivered as plug-ins for software of other providers in 44% of cases. Scissero (2021) is an exemplary start-up, with ML software supporting legal departments by analyzing or suggesting drafts of legal documents.

Internal business diagnostics & prediction: All start-ups of this archetype aim to support their clients' decision-making processes by supplying relevant information about the respective client's internal business activities. A software package that combines an IS with ML capabilities analyzes the client's internal data and extracts essential facts for decision-makers. 50% of start-ups in the archetype additionally forecast measures to reduce uncertainty in decisions. The archetype focuses mainly on clients in the tertiary sector, and its ML systems primarily utilize structured (92%) client data available on internal processes (92%). Start-ups of this archetype integrate their ML technologies into ISs, so software engineering is often a key activity (82%). As an archetypical example, CognitOps (2021) utilizes client data to assist warehouse managers with operational decisions, e.g., in scheduling.

Environmental diagnostics & prediction: Start-ups in this category also assist their clients' decision-making processes through insight increases – usually through an aggregation & filtering key offering (88%) as well. However, unlike the archetype *internal business diagnostics & prediction*, they provide information on elements external to the client organization, such as financial markets or public opinion. This seemingly small distinction has further implications for BMs, distinguishing this archetype: Due to the high relevance of publicly available data (75%), sourcing data and setting up pipelines to provide curated data for the ML systems is essential for 82% of the BMs. In 50% of cases, the incoming data is also unstructured, requiring interpretation. Start-ups in this archetype slightly favor targeting the tertiary sector (31%), or all of them with an ML solution for a cross-sectional function. Nevertheless, there are some with a focus on the secondary or quaternary sector, making this a widely diffused archetype. The example for this archetype, Sanctify Financial Technologies (2021), aids in asset management decisions by analyzing non-financial news articles on potential investments.

6.5 Conclusion, Limitations & Outlook

To examine how ML technologies influence BMs of organizations, we use the methodical approach of Nickerson et al. (2013) to develop an industry-overarching taxonomy of ML-driven B2B BMs by analyzing a sample of 102 start-ups. The taxonomy describes possible characteristics of start-ups along ten dimensions and addresses **RQ 1**. Moreover, we aggregate the start-ups into seven BM archetypes that represent designs commonly found in practice and thus address **RQ 2**.

Our study provides several **theoretical contributions**. Literature has been calling for further taxonomic research related to data-driven BMs (e.g., Veit et al., 2014; Müller & Buliga, 2019; Omerovic et al., 2020). With ML giving rise to opportunities and challenges unique to the technology (Benbya et al., 2021) when transforming BMs (Björkdahl, 2020; Burström et al., 2021), we extend the discussion to ML-driven BMs. In this field, the presented taxonomy with the described dimensions and characteristics fosters a deeper understanding of the anatomy of ML-driven BMs. It allows researchers to specify ML-driven BMs in a unified manner and distinguish them from each other. By standardizing the vocabulary in this topic area, we facilitate the scientific exchange between researchers and future work in the context of ML-driven BMs. Using a common language, new ideas can be objectified, and considerations can be shared among scientists to build a more profound theoretical understanding of ML-driven BMs. This allows further systematization of research in this field. Furthermore, the presented artifacts provide a basis for future taxonomic research. Researchers can validate or extend them for narrower scopes like specific application domains (see Anton et al., 2021), specifying and extending the dimensions and characteristics. Another contribution of our paper is the method-based identification of BM clusters and the subsequent derivation of archetypes of ML-driven BMs. The seven derived archetypes provide deeper insight into the structural composition of commonly implemented ML-driven BMs. Additionally, our research also offers **practical contributions**, of which we focus on two: First, the taxonomy in combination with the archetypes of BMs offers a comprehensive overview of the market. Organizations and other stakeholders can benefit from this increased understanding by improving investment decisions or assessing their own as well as their competitors' BMs. Second, practitioners can further benefit from our artifacts by using them as supporting tools for the conception of novel ML-driven BMs through the structured recombination of BM components (Bouwman et al., 2020), therefore employing our research results to facilitate BMI.

As with every research, our study is subject to certain **limitations**. Despite the previously mentioned advantages, our focus on start-ups as a data basis excludes established organizations, leading to several consequences: In particular, the identified archetypes can only be transferred to established organizations to a limited extent, as these may already have higher resources and established structures in place and can offer different services accordingly. Further, the developed taxonomy may need to be adapted and expanded to reflect the specifics of established organizations, e.g., in terms of the key activities undertaken. Even though this would be a valuable line of further research, we are confident that the taxonomy stemming from the start-ups is broadly applicable to more mature organizations as well, as we have intentionally abstracted from specifics of start-ups within the taxonomy and kept dimensions and characteristics rather generic. Moreover, despite the methodical approach followed, many steps in the taxonomy development process required the researchers' own judgment (Nickerson et al., 2013). Therefore, another group of researchers might encounter different dimensions and characteristics from this study. Similarly, archetypes may vary based on the chosen number of clusters, which in turn depends on the employed algorithm (see Mojena, 1977). Our dataset includes a large variety of ML-driven

BMs as per our purpose. However, aggregating them into a manageable number of distinct yet general archetypes is challenging, which can be seen in the consistency between algorithms of some clusters (see Table 11). Nevertheless, as taxonomies and archetypes are never perfect, they should instead be assessed based on whether they are useful (Nickerson et al., 2013; Remane et al., 2016); a quality revealed when researchers and practitioners start using them.

Acknowledgment

This research and development project is/was funded by the German Federal Ministry of Education and Research (BMBF) within the “Innovations for Tomorrow’s Production, Services, and Work” Program (funding number 02L19C150) and implemented by the Project Management Agency Karlsruhe (PTKA). The authors are responsible for the content of this publication.

7 Research Paper E: Capturing Value by Realizing Machine-Learning-driven Business Models

Title: As Much Art as Science – Examining the Realization of Business Models Driven by Machine Learning Through a Dynamic Capabilities Perspective

Authors: Vetter, Oliver A.; Mehler, Maren F.; Buxmann, Peter

Published in: Proceedings of the 31st European Conference on Information Systems (ECIS 2023)

Abstract

Machine learning (ML) technologies open up enormous potential to be unlocked through entrepreneurial activities in organizations, causing countless novel business models with ML at their core to emerge in the market. As ML technologies differ significantly from other digital technologies both in their characteristics and their effect on organizations, little is currently known about the complexities of the realization process for business models driven by ML and why only some organizations execute it successfully. By building on a qualitative study grounded on cross-industry insights from 20 expert interviews, this paper contributes to a greater understanding of the realization process by identifying ML-specific complications, before aiming to determine the underlying reasons for successful business model realization. We adopt a dynamic capabilities perspective and conceptualize eleven microfoundations that explicate how organizations build, implement, and transform business models driven by ML.

Keywords: Machine Learning, Business Model Realization, Dynamic Capabilities

7.1 Introduction

Machine learning (ML) unlocks possibilities to support or entirely automate processes within organizations (Jordan & Mitchell, 2015) and further provides powerful opportunities for entrepreneurship by enabling entirely new services and business models (Chalmers et al., 2021; Davenport et al., 2020). ML denotes a technology that can be utilized to create instances of artificial intelligence (AI) by allowing algorithms to learn patterns hidden within data and then make predictions for new data (Russell & Norvig, 2021; Brynjolfsson & Mitchell, 2017; Mitchell, 1997). The novel business models with ML at their core are distinct from other types of business models enabled by information technologies (IT), as recent literature shows that ML not only opens up new possibilities for value proposition, but potentially impacts the overall business logic (as we will further discuss), e.g., by continuously learning with new data (e.g., Weber et al., 2022; Vetter et al., 2022). However, these studies have mostly examined the ideation of ML-driven business models, disregarding how organizations develop and realize them. Only little research extends beyond the ideation phase and covers the realization process of the broader category of business models enabled by data, focusing either on models of the development process (e.g., Hunke et al., 2017) or on the resources that are required for it (e.g., Lange et al., 2021). However, in fast-moving business environments, organizations need more than the ownership of expertise, processes, and resources to maintain sustainable competitive advantage (Teece, 2007). For instance, they must field a strategy that allows them to both defend their position in the market against competitors and contingencies (Casadesus-Masanell & Ricart, 2011). Following Teece (2018), we argue that an organization's dynamic capabilities allow it to connect strategy and business models to navigate complex and fast-changing environments through the creation, realization, and refinement of business models (Teece, 2018; Winter, 2003); a process that is "as much art as science" (Schoemaker et al., 2018, p. 27). Complicating development of suitable business models further, ML exhibits some characteristics that differ significantly from other digital technologies and come interleaved with distinct managerial challenges within organizations (Benbya et al., 2021; Berente et al., 2021), e.g., due to their ability to supersede humans at work (Murray et al., 2021). Moreover, the development of ML is highly uncertain and experimental in character (Choudhury et al., 2021; Amershi et al., 2019). This is problematic as organizations failing to counteract ML-specific challenges while navigating dynamic business environments may fail to acquire and maintain competitive advantages. We thus deem dynamic capabilities a suitable lens to examine the realization process of business models driven by ML and pose the following research question:

How do organizations build dynamic capabilities to empower their ML-driven business model realization process?

To answer this question, we conduct an explorative study and investigate the microfoundations that substantiate the dynamic capabilities enabling the realization of business models shaped by ML. In doing so, we contribute to extant dynamic capability literature in the context of digitalization by verifying and expanding known microfoundations with ML-specific aspects and identifying new microfoundations emerging due to unique ML characteristics. Furthermore, by specifying these ML-specific dynamic capability microfoundations, we open up the field for future organizational research on entrepreneurial activities to capture value from ML technologies. Additionally, we inform practitioners on the effect of ML technologies on the organization and its business environment during business model realization and offer guidelines for practitioners on how to create dynamic capabilities that aid the ML-driven business model realization process.

Thereby, we support organizations in analyzing their present business model realization efforts and in building capabilities for future endeavors.

7.2 Theoretical Background

Next, we first present research on the development of business models with ML technologies at their core before covering the theoretical foundation of dynamic capabilities in the business model context.

7.2.1 Realizing ML-Driven Business Models

Researchers have put great emphasis on examining the business model concept, which depicts the business logic of an organization, including how it creates and delivers value to its clients, as well as the corresponding architecture of revenue, costs, and profits (Teece, 2010). Numerous definitions of the term exist in the literature (Al-Debei & Avison, 2010; Wirtz et al., 2016; Zott et al., 2011; Birkinshaw & Ansari, 2015), of which we adhere to the definition by Osterwalder and Pigneur (2010) for this study, which reads as follows: “A business model describes the rationale of how an organization creates, delivers, and captures value” (p. 14). The available conceptualizations of business models often consist of their constituting components, such as the value proposition or the revenue stream (e.g., Al-Debei & Avison, 2010; Zott & Amit, 2010; Remane et al., 2016). In information systems (IS) research, the business model is regarded as the missing link, acting as an intermediary between strategy, processes, and IT (Veit et al., 2014). Furthermore, the concept is seen as a useful lens for examining competitive advantage in management literature and is thus valuable in theory building to generate novel insights (e.g., Lanzolla & Markides, 2021; Prescott & Filatotchev, 2021).

AI, which denotes “the science and engineering of making intelligent machines” (McCarthy, 2007, p. 2), enables organizations to leverage additional potential in the context of digitalization (Davenport et al., 2020; Dingli et al., 2021; Makridakis, 2017). With the large availability of data and advancements in data analytics, AI technologies have thus regained importance in recent years (Ågerfalk, 2020; Berente et al., 2021). As most modern AI systems have ML technologies at their core (Brynjolfsson & Mitchell, 2017; Jordan & Mitchell, 2015), we use the term *ML* to refer to ML-based instances of AI in this study for terminological clarity. While the unique opportunities unlocked by ML drive the development of new types of business models in organizations (Weber et al., 2022; Vetter et al., 2022; Wamba-Taguimdje et al., 2020), ML technologies also exhibit characteristics that differ considerably from other digital technologies, presenting ML-utilizing organizations with novel challenges (Benbya et al., 2021; Berente et al., 2021). First, the self-learning algorithms utilized in ML to learn and improve automatically from data (Amershi et al., 2019; Russell & Norvig, 2021) are incapable of reacting to environment states which they have not been trained with (Dennett, 2006) and require human guidance for framing the respective tasks and interpreting the results (Seidel et al., 2020; Salovaara et al., 2019). Second, ML systems can act autonomously (Berente et al., 2021; Baird & Maruping, 2021) and can even take over tasks previously firmly in human grasp (Schuetz & Venkatesh, 2020; Benbya et al., 2021). Lastly, modern ML systems have become increasingly complex, thus making their behavior inscrutable and difficult to understand for humans (Faraj et al., 2018; Asatiani et al., 2021). This is especially problematic as their introduction in organizations can create unexpected and unintended outcomes (Benbya et al., 2020b; Benbya et al., 2021) that can lead to a variety of ethical issues, e.g., discriminating ML systems which include inappropriate factors in their decision-making

(Martin, 2019a). With the behavior of many ML systems being inscrutable for humans, preventing such unforeseen ethical, legal and practical consequences proves difficult (Asatiani et al., 2021). Investigating how this variety of unique challenges affects the business model development process is thus part of this study.

In light of the ongoing digitalization, IS researchers often utilize the business model concept to examine how advancements in IT, such as progress in big data and data analysis technologies, transform how established organizations create value and which novel types of ventures they enable (e.g., Hartmann et al., 2016; Steininger, 2019). For this study, we follow previous research and denote such business models as data-driven when they utilize data as a key resource (Hartmann et al., 2016). A subset of data-driven business models are ML-driven business models, which have infused at least one of their business model components with ML technologies (Vetter et al., 2022; Hahn et al., 2020). IS literature has thus far focused on the ideation of data-driven business models, largely disregarding how organizations develop and realize such business models and which strategies they employ (Lange et al., 2021; Wiener et al., 2020). Hunke et al. (2017) propose an innovation process depicting the major phases organizations must go through when developing data-driven business models and the tasks they must perform in each phase. Rashed and Drews (2021) differentiate the realization process into four pathways that established enterprises can take, depending on their data understanding and incentive for change, and elaborate corresponding implementation strategies. Similarly, Shollo et al. (2022) specify three ML value creation mechanisms through which organizations can reach their own organizational goals and present resources and conditions that must be met to shift between mechanisms. Finally, Lange et al. (2021) adopt the resource-based view of the firm (see Barney, 1991) and present the resources necessary for different phases of the business model realization process, which they subsume under four capabilities, along with challenges and enablers to fully utilize the organization's resources.

7.2.2 *Dynamic Capabilities for Digitalization*

However, assembling the resources and capabilities needed to achieve sustainable competitive advantage is only one part of business model development (Teece, 2018). While these operational, ordinary capabilities aid organizations in operating a business model efficiently, e.g., in following a specified manufacturing program, an organization's overlying dynamic capabilities determine the success in creating, implementing, and transforming business models (Teece, 2018; Winter, 2003; Ricciardi et al., 2016). In their seminal paper, Teece et al. (1997) define dynamic capabilities as the "ability to integrate, build, and reconfigure internal and external competencies to address rapidly-changing environments" (p. 516). While early dynamic capabilities research disputed whether they are indeed firm-specific as proposed by Teece et al. (1997) or common among organizations and whether they necessarily confer superior performance (e.g., Eisenhardt & Martin, 2000; Peteraf et al., 2013), recent literature proposes that dynamic capabilities may consist of both elements that are common across organizations and aspects that are idiosyncratic (Barreto, 2010; Yeow et al., 2018). We follow Yeow et al. (2018) in conceptualizing dynamic capabilities as both broad organizational capabilities and specific actions and in adopting, at the broad level, the three dynamic capability clusters, or higher-order capabilities, proposed by Teece (2007): *sensing*, *seizing*, and *transforming* (Teece, 2018). Each of the higher-order dynamic capabilities can be disaggregated into various second-order dynamic capabilities, or microfoundations (Teece, 2018), which represent the actions and processes enacted by individuals (including managers) within the organization, that build and maintain dynamic

capabilities (Vial, 2019; Helfat & Peteraf, 2015; Yeow et al., 2018). The strength of an organizations dynamic capabilities then determines how successfully it innovates and adapts to rapidly changing markets and technological progress (Di Stefano et al., 2014; Warner & Wäger, 2019; Eisenhardt & Martin, 2000). Due to the highly disruptive and fast-changing nature of the digitalization in general, dynamic capabilities are seen as a suitable theoretical foundation to examine the mechanisms that enable organizations to engage in digitalization (Warner & Wäger, 2019; Ellström et al., 2022; Vial, 2019). Currently, literature is thus calling for additional research on how organizations build dynamic capabilities to propel their digital transformation forward and which microfoundations exactly constitute these dynamic capabilities in practice (Vial, 2019). Furthermore, dynamic capabilities enable business models in the sense that they allow organizations to rapidly design, test, and revise novel and modified business models, while simultaneously being enhanced by the organizational flexibility allowed by the business model of the organization (Schoemaker et al., 2018; Teece, 2018), which is especially relevant for organizations operating in rapidly-changing, complex, and uncertain environments (Schoemaker et al., 2018). Given the disruptive potential of ML-driven business models in particular (Davenport et al., 2020; Chalmers et al., 2021; Townsend & Hunt, 2019), in combination with the unique managerial challenges posed by ML (see Section 7.2.1), we argue that dynamic capabilities are therefore a compelling perspective for investigating the determinants of successful ML-driven business model realization. As the allocation of dynamic capabilities into the three clusters proposed by Teece (2007) is widely accepted in the literature (Ellström et al., 2022; Yeow et al., 2018; Warner & Wäger, 2019), we adopt it for this study and elaborate on the three groups in the context of realizing ML-driven business models in the following.

Sensing is the “identification, development, codevelopment, and assessment of technological opportunities in relationship to customer needs” (Teece, 2014, p. 332). Sensing and shaping opportunities, as well as threats, is done through activities involving scanning, creation, learning, and interpretive activities, which need to be embedded into suitable organizational routines (Teece, 2007). Organizations must be aware of the ecosystem surrounding them, including customer needs and technological possibilities, as well as the structural evolution of markets and likely competitor responses (Teece, 2007) to be able to manage market uncertainty and detect opportunities (Teece, 2009). This dynamic capability is especially relevant for organizations with digital technologies deeply embedded in their strategy (Yeow et al., 2018), as it enables recognizing and understanding unexpected trends in fast-changing environments to be able to adapt accordingly (Warner & Wäger, 2019).

Seizing denotes the “mobilization of resources to address needs and opportunities, and to capture value from doing so” (Teece, 2014, p. 332). Once an opportunity is sensed by an organization, it must be addressed through new products, processes, services, or a combination of these, to ensure the value capture through appropriate investments (Teece, 2007). As the outcomes of such investment decisions are often highly uncertain, organizations must develop strong decision-making and evaluation abilities that foster innovation (Teece, 2007). When incumbent organizations are introduced to new technologies, they often experience a gap between the configuration of ordinary capabilities present in their organization and the optimal configuration required to fully utilize the technology (Karimi & Walter, 2015), necessitating seizing capabilities to incorporate the technology into the organization to allow capturing value from corresponding opportunities (Ellström et al., 2022).

Transforming is the “continued renewal” (Teece, 2014, p. 332) of the organization’s business model along with its resource base. Organizations must retain their ability to reconfigure their

assets and organizational structures as their size grows and market and technologies change to maintain evolutionary fitness and achieve sustained profitable growth (Teece, 2007). To sustain competitive advantage in dynamic environments, established routines need to be revamped constantly, which involves top management leadership skills, business model redesigns, and potentially even organizational restructuring for radical innovations (Teece, 2007; Helfat et al., 2007). The transforming capability thus aids in aligning existing resources to new strategies as well as in building or accessing required new resources (Yeow et al., 2018). Consequently, with digital technologies and especially ML being relatively new and many organizations lacking the associated expertise and routines, transforming capabilities are crucial for organizations that pursue a strategy infused by these technologies and, therefore, must acquire the corresponding resources (Yeow et al., 2018; Rindova et al., 2016).

7.3 Methodology

To answer our research question, we carried out a qualitative expert interview study (Bogner et al., 2009; Gläser & Laudel, 2004). In contrast, some comparable studies on dynamic capabilities in established enterprises select case studies for their qualitative approach (e.g., Yeow et al., 2018; Mousavi et al., 2019), which allow for in-depth investigations of phenomena in few organizations that are exemplary for the respective topic (Sarker et al., 2012). However, as novel business models are volatile, their success difficult to determine in the first years, and their established processes and resources to draw data from typically scarce, we decided to rely on multiple perspectives on our topic stemming from various organizations. Therefore, we chose to conduct an expert interview study, a suitable approach for incorporating a wide range of expertise on ML-driven value creation (e.g., Lange et al., 2021; Shollo et al., 2022) and an appropriate method for illuminating areas of research that have not yet been thoroughly explored (Corbin & Strauss, 2015; Myers & Newman, 2007). For the interviews, we concentrated on experts from the fields of digital business and data science with insights into the technical aspects as well as the business aspects of realizing ML-driven business models. The experts were chosen from organizations of various sizes (including established enterprises and start-ups) and in different industries (including experts from consulting firms or accelerators with an overview of a wide range of client organizations) and recruited through LinkedIn and other networks. All selected organizations initiated ventures to realize their own ML-driven business models or advise their clients in respective endeavors.

In total, we conducted 20 expert interviews in the third and early fourth quarter of 2022. Table 12 shows a list of all interviewed experts, along with information on the corresponding organizations. Where applicable, we specified experience into time spent developing ML-driven business models and total time spent developing other types of digital business models in brackets, as ML-driven business models only recently gained popularity (Weber et al., 2022). *Scope* denotes whether the ML-driven business model was an entirely new, independent business model (*New Business*), developed out of an existing business model (*Extension*), or transformed from a previous business model (*Transformation*). The consulting firms in our sample mostly supported Transformation or Extension projects, yet could advise all scopes (*Flexible*). Interview E03 was conducted in person, and interview E17 over the phone, while all other expert interviews were held over the online conferencing tools Zoom or Microsoft Teams. All interviews lasted between 37 and 54 minutes with the exception of interview E03, which lasted 87 minutes due to the large amount of helpful anecdotes shared by the expert. The average duration of all interviews is 48 minutes. All expert interviews were either conducted in German or English.

Table 12. List of Interviewed Experts.

ID	Position	Experience	Industry	Size	Scope
E01	Product Lead ML	8 yr.	IT Services	Large	Extension
E02	Co-Founder	3 yr.	Software	Small	New Business
E03	Managing Director	5 (12) yr.	IT Services	Small	New Business
E04	Co-Founder & CEO	6 (10) yr.	IT Consulting	Small	Flexible
E05	Global Head of ML	6 (17) yr.	IT Services	Very large	Extension
E06	ML Consultant	2 (5) yr.	IT Consulting	Small	Flexible
E07	Co-Founder & Head of ML	5 yr.	IT Services	Small	New Business
E08	Co-Founder	5 (10) yr.	Software	Small	New Business
E09	ML Consultant	5 (7) yr.	ML Accelerator	Small	New Business
E10	Director of ML Strategy	6 (11) yr.	ML Initiative	Medium	Extension
E11	Lead ML Product Manager	6 (21) yr.	Software	Very large	Transformation
E12	Co-Founder & CEO	2 yr.	IT Consulting	Small	Flexible
E13	Manager Digital CX	6 yr.	Technology	Very large	Transformation
E14	Senior Data Scientist & Senior Consultant	4 yr. & 5 yr.	Consulting	Very large	Flexible
E15	Data Scientist	6 yr.	Automotive	Very large	Transformation
E16	Senior ML Consultant	3 yr.	Consulting	Very large	Flexible
E17	Senior ML Strategy Manager	4 (12) yr.	Consulting	Very large	Flexible
E18	Founder & Managing Director	4 yr.	IT Services	Small	New Business
E19	Co-Founder & CFO	3 (9) yr.	IT Services	Small	New Business
E20	Project Lead	6 yr.	ML Accelerator	Small	New Business

Note. Very large: empl. > 1000, large: empl. > 250, medium: empl. > 50, small: empl. ≤ 50.

In preparation for the interviews, we created a semi-structured interview guideline (Myers & Newman, 2007). The interview guideline first contained introductory questions on the experiences of the interviewee and on the business models under study. It then progressed into thematic questions informed through extant research (Teece, 2007; Teece, 2018; Yeow et al., 2018; Warner & Wäger, 2019; Leemann & Kanbach, 2022; Lange et al., 2021), revolving around each of the three dynamic capability clusters for ML-driven business model realization. We recorded and fully transcribed all interviews. Next, we carried out a content analysis of the interview data based on the transcripts, choosing an inductive approach to avoid imposing preconceived ideas on dynamic capability theory on the data and instead let dynamic capability microfoundations emerge from the data (Hsieh & Shannon, 2005; Mayring, 2007; Myers, 1997). More specifically, we followed Gioia et al. (2013) in our coding process. To perform the coding, we used the MAXQDA software. We went through the transcripts line by line and assigned codes to sentences or parts of sentences that described the respective units, preferably in informant terms. Next, we iteratively grouped these emerging first-order concepts into meaningful second-order concepts, representing dynamic capability microfoundations. Finally, we assigned the latter to aggregate dimensions corresponding to the three dynamic capabilities clusters. During the analysis, we employed multi-researcher triangulation to achieve rigor and a high degree of objectivity (e.g., Hsieh & Shannon, 2005; Carter et al., 2014). In total, we identified three microfoundations of sensing, four of seizing, and four of transforming dynamic capabilities (see Figure 7). Furthermore, various factors uniquely complicating the development of ML-driven business models emerged from the data, which we grouped into two aggregate dimensions based on where they manifest. We thus identified four complexities taking effect within the business

models themselves and three relating to dynamism in the organizations' business environments, which we elaborate on in Section 7.4.1.

7.4 Results

In the following, we first address particularities unique to ML-driven business models and their business environments that complicate business model realization and intensify the need for dynamic capabilities. Next, we elaborate on how the examined ML-driven organizations built and utilized their dynamic capabilities during business model development, using the tripartite clustering by Teece (2007).

7.4.1 Particularities of Realizing ML-Driven Business Models

The analysis of the interview data confirms that there are various aspects complicating the realization process of ML-driven business models, evoked by the ML technology at their core, which we subsume under the categories *environmental dynamism* and *business model complexities*.

Environmental dynamism delineates that both the technology of ML as well as the market surrounding corresponding business models are in constant flux. First, on the technological side, the speed at which new ML approaches are being developed is enormous (E11), even outpacing other digital technologies, according to the interviewees (E03, E18). Therefore, organizations and even start-ups that want to position themselves in the field of ML must build competencies for research and development from the start to avoid getting left behind (E03). The rapid releases of novel ML services additionally demand much more flexibility from organizations during the development process of business models: *"So these classic [business model] planning processes, that's very difficult – that's usually obsolete very quickly. The world changes too quickly"* (E18). Second, the experts note that for many ML solutions, the respective market is still nascent (E08, E18), following the hype evoked by successful ML proof-of-concept projects (E20). As the market is currently developing, the IT infrastructure around ML systems is becoming increasingly modular and organizations are moving to specialize in application areas or ML components (E08, E11, E15, E18). *"There will be companies that try and cover the full spectrum of the machine learning flow [for one specific application area] – or you'll get companies that do best in breed on one specific component"* (E08). Thereby, the ever-growing product portfolio of the large cloud service providers enables the purchase of specific components of ML development, such as computing power or pre-trained ML models, lowering the entry barriers for organizations with limited resources of their own at the cost of operating expenditure (E16). Third, as ML models absorb information about the organization's environment in the form of data, changes in the business environment not only become problematic due to the possibly negative effects of the changes, but additionally due to their impact on the ML model within the organization's ML system (E14, E16). When changes occur, the arising technical problems of data drift (new data is unlike training data of ML system) and concept drift (interrelations within new data is unlike interrelations in training data) must be addressed, which requires humans to interpret the data and evaluate whether the business idea remains relevant and which technical changes are necessary to ensure the ML system learns the desired patterns (E08, E14, E16).

Business model complexities denotes factors complicating business model realization from within the business model due to the unique effects of ML technologies. First, with ML being a general-purpose technology (see Brynjolfsson & McAfee, 2017), corresponding solutions can be

utilized across various different sectors and business functions (E03, E05, E19). ML-developing organizations must thus be able to recognize application opportunities for their ML systems across industries and organizational boundaries (E03). Second, ML start-ups are often very tech-driven because of the high degree of required technical expertise in-house, frequently leading to some neglect of the business side (E09, E19). However, business expertise is elementary in developing business models, for instance, to quickly gather market feedback (E09). Additionally, ML solutions and their non-deterministic output require intensive engagement with the respective client to make them understand “*how ML works and what they can do with it and where it really helps them*” (E01), necessitating the building of expertise in sales and marketing (E01, E02). Third, while traditional digital solutions were separable into data and software, allowing the molding of standard processes into software (E10), that is no longer the case for ML approaches, in which data and software are inextricably connected (E08, E10, E12). This begets some technical challenges during development: Not only is data difficult and costly to acquire (E10, E12, E14, E17), but whether the process of abstracting information from data, storing it in an ML model, and then applying that in the real environment works and how well it works is difficult to ascertain in advance (E01, E07, E08, E10). Deploying an ML system to a new application area, therefore, not only requires retraining (E02, E10) but also extensive evaluation efforts (E06, E08, E10) that continue in the maintenance phase afterwards (E06). As one expert put it: “*Mike Tyson said: ‘Everyone got a strategy until they get hit in the mouth.’ Your model is going to get hit in the mouth. What do you do with it when it does happen? How do you feed that into a system or [to] a person that can interpret that, update the model, get the new model, deploy it down so the next time it sees an exception, it can handle it?*” (E08). The uncertain performance and the high retraining efforts when transferring ML systems between clients consequently make ML-driven business models very difficult to scale (E02, E07, E10). Fourth, ML systems have the capacity to highly individualize services to single end-consumers, e.g., in social media networks (E20), or to autonomously make high-impact decisions, e.g., in self-driving cars (E15). Coupled with their opacity (E15, E17), ML systems can cause more severe ethical ramifications than traditional software (E11, E14, E20). ML-driven organizations, therefore, need strong governance to foresee ethical issues early on when developing ML systems and to constantly review potential problems even after development to ensure operational safety (E14, E15) and data privacy (E11, E13, E15).

7.4.2 *Dynamic Capabilities for Realizing ML-Driven Business Models*

In what follows, we present the identified microfoundations explicating how organizations successfully realize ML-driven business models despite the unique challenges described in Section 7.4.1. Figure 7 gives an overview of the microfoundations that emerged from the data and the first order-concepts that informed them. Thereby, organizations first sense worthwhile opportunities via *sensing*, then design and commit to an appropriate business model via *seizing*, and lastly reconfigure the organization via *transforming* (Teece, 2018; see Warner & Wäger, 2019). While the dynamic capabilities are thus used sequentially when realizing a new business model, organizations must sense and seize opportunities continuously and transform parts of the organization periodically for long-lasting success, with stronger dynamic capabilities leading to faster and closer alignment to customer needs (Teece, 2018).

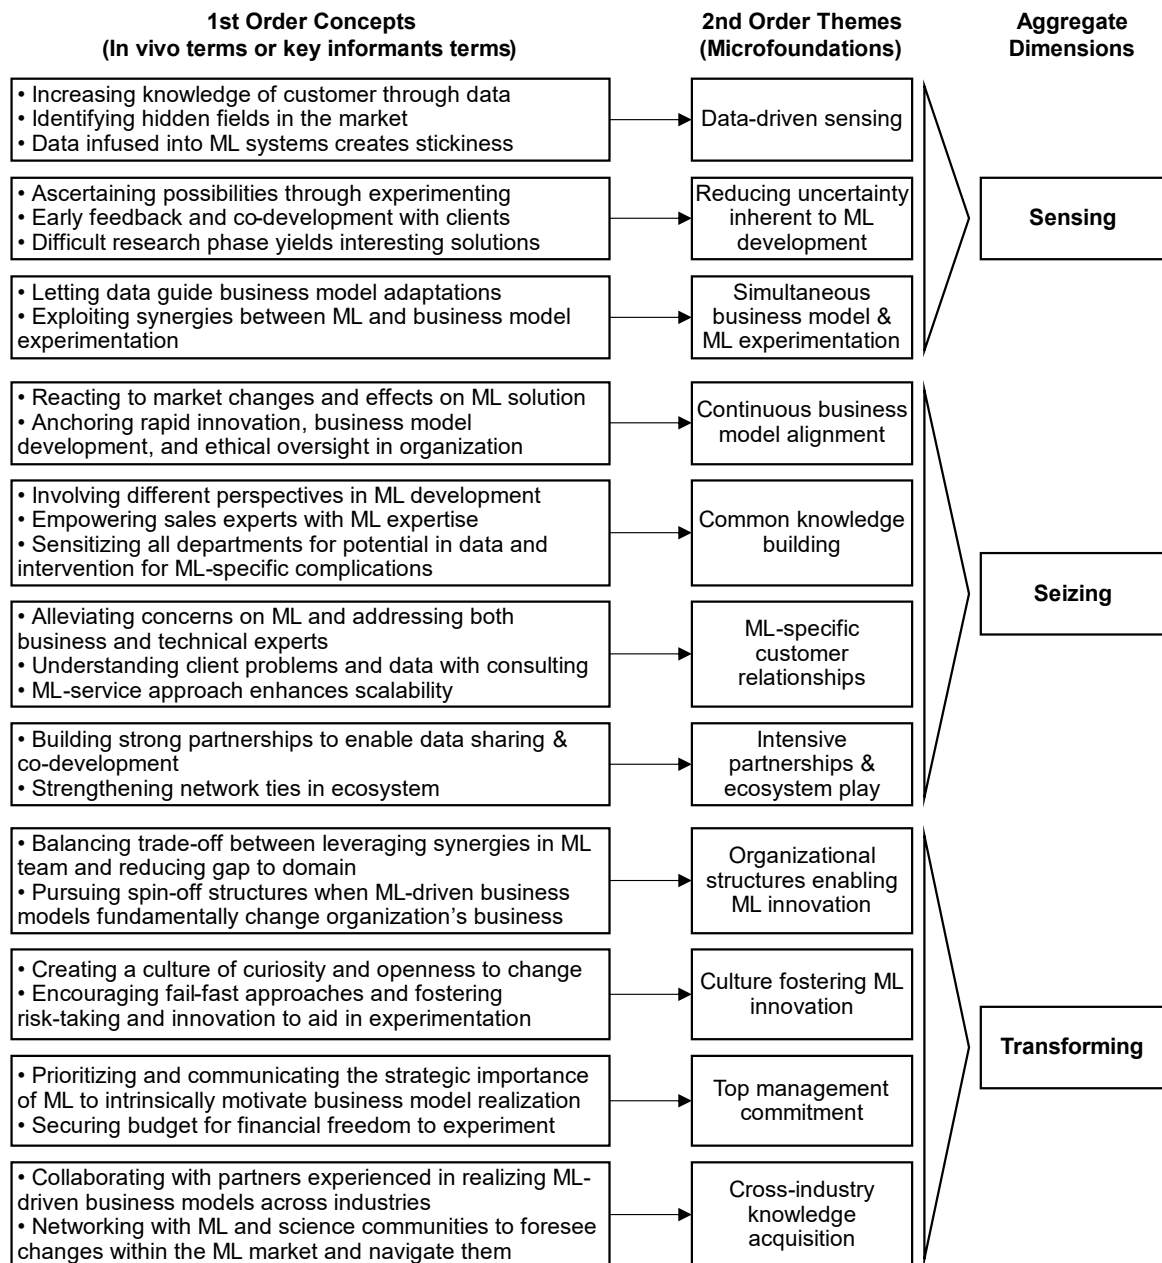


Figure 7. Microfoundations of Dynamic Capabilities for Realizing ML-Driven Business Models Emergent From the Interview Data.

Sensing

Due to ML systems requiring large amounts of well-curated and current data, ML-driven business models can utilize their database to engage in *data-driven sensing*. By analyzing the available data manually or through automated processes, organizations can thereby gain insights into the characteristics and especially the needs of their customers (E01, E02, E05, E13, E16, E17). Through the organizations' improved knowledge of their customers, they increase their capability for detecting opportunities to serve additional needs on the market. "[Name of former employer] always say that they get so much from data; they understand their target groups much better because they analyze massive amounts of data and, in fact, find many hidden fields" (E02). While the capacity for data-driven sensing is beneficial to various types of business models, "when you're based around machine learning, [listening to data is] that much more important because that is driving how you win in the market. Also, your data is a form of stickiness with customers. I wouldn't

stop using Spotify because I fed it huge amounts of data, and the machine learning algorithm knows what to recommend” (E05). Another microfoundation of organizations’ sensing capability for ML-driven business model realization is *reducing uncertainty inherent to ML development*. Compared to traditional agile software development, where “*you set an amount of story points or a certain view of what you’re going to do in the next sprint, and you know [...] what you will get at the end [of that sprint]. You don’t have that from the machine learning point of view because there’s data in the equation”* (E05). Therefore, ML development requires prolonged experimentation phases to ascertain the possibilities of the envisaged ML solution (E01, E02, E05, E14, E18). Agile principles known from other digital products (E04, E16) such as small, incremental development steps (E03, E13, E16) while embracing a fail-fast mindset, quickly testing many ideas and allowing ideas based on mismatched assumptions to fail early (E07, E13), aid in reducing ML-induced uncertainty during development as well. Moreover, as setting up a functioning ML model is difficult early on (E09, E13, E17, E20), ML-driven business models often co-innovate, working closely with pilot customers to test assumptions early and to gain access to the customer’s data (E03, E06, E09). While this laborious experimentation phase is considerably longer still for problems that are not yet well-solved, these cases also yield the most interesting ML solutions (E04, E18). The last sensing microfoundation, *simultaneous ML and business model experimentation*, describes organizations remaining adaptable with their business model in parallel to their technical experimentation (E05, E09, E14, E18). First, organizations are well-advised to let the information derived from data also inform their business goals and to adapt their business model accordingly (E05, E19). “*The data needs to guide you, and data is telling you about your business model and its effect in the market”* (E05). Second, the question of which to develop first, business model or ML system, is “*a bit of a hen and the egg question: [...] Do you first spend a lot of time in making a business model because that takes resources as well or do you make a proof of concept?”* (E14). Waiting for the experimentation phase is often not an option either, as the organization must prove the rentability of the idea to justify the high costs of ML development (E02, E09, E11, E13). Organizations should therefore jump back and forth between the two, continuously evaluating the fit between their business model idea and the results of the ML experimentation process (E02, E12, E14, E20). Thereby, these two learning processes amplify each other, giving rise to positive synergies between them (E09, E14, E19): “*Along the process, often you can find that you’re able to do more things. You’re able to actually predict some other things that you didn’t think about at the start”* (E14). Due to the characteristic of ML being applicable across sectors and functions, organizations can thus “*recogniz[e] that the tool we have created can do much more than what we have done with it so far”* (E03). Organizations may find their solutions to be applicable in new, previously unaddressed sectors (E03, E19) or discover that their ML system intended for use by end consumers unlocks the potential to lower costs when used internally (E18). The reverse of the latter is even more common: “*We develop something for a specific section in the value chain – for example, for production, error management processes is a big topic for us – and then we suddenly realize our application has a lot of potential to be used in after sales”* (E18).

Seizing

Organizations seeking to capitalize on opportunities with their ML-driven business model realization process need the capacity to react to market or environment changes through *continuous business model alignment*. While all types of business models should be adapted regularly to reflect market conditions (E09, E11), being able to respond quickly to developments is especially critical for business models in rapidly changing markets such as the ML market (E03, E11). As dynamism in the market also enters ML products through data, causing data or concept

drift, the need for continuous alignment is further exacerbated: *“It starts with an invoice template, which changes in some country because some additional information is necessary, and I have to learn to extract that too if I want to extract something. [And it ends] with the fact that suddenly we have to keep a distance of 2 meters, we all have to wear masks, and the life behavior has completely changed, and therefore certain business processes have completely changed”* (E11). To enable fast business model adaptations, organizations may establish departments such as innovation labs to rapidly test hypotheses and business ideas, while simultaneously anchoring the process of industrializing and scaling promising ideas within parts of the organization as well, e.g., in an ML factory unit (E17, E20). The potential ethical issues arising from some ML solutions demand additional awareness and that organizations are quick to respond to new regulations and to adapt their business processes accordingly in a responsible manner (E04, E05, E06), which can be fostered through involving people with differing perspectives (E06). Another microfoundation fueling the seizing capability of organizations realizing ML-driven business models is their ability for *common knowledge building* among departments due to three reasons: First, for the development process of ML systems to run smoothly, it needs the involvement of experts with different skillsets and perspectives, such as data engineers, ML engineers, sales experts, and domain experts (E17, E20). To allow experts from the business side to give input during development, organizations must ensure that all employees across organizational structures have some understanding of ML technologies and how they function (E01, E02, E08, E11, E14, E19). Second, that shared technical understanding is also necessary beyond ML development. As the output of ML models is non-deterministic and often uncertain in advance (E01, see Section 7.4.1), ML-driven business models need highly tech-savvy sales experts to be able to explain the unique characteristics as well as their ramifications to clients (E02, E08, E11, E20). Such skilled sales also *“have to be able to transport [...] trust in the technology, because that has not been there until today”* (E11). Third, a shared technical understanding sharpens the senses of all members of an organization for ML-specific issues (E10, E14, E18). Decisions made around a deployed ML system can unintentionally alter its results, e.g., in the visual inspection example of an expert, where *“someone says ‘Hey, we’re going to install new lights’ or some other stuff, which at first has no direct dependency, but of course then maybe the image is influenced and then all of a sudden my model no longer works”* (E10). As ML systems can not interpret their surroundings (see Section 7.4.1), the humans around them must thus possess an awareness of such issues and the technical understanding for reasonable interpretations and for requesting intervention if necessary (E10, E14). Moreover, with more members of an organization paying attention to the potential within their data, detecting promising ML use cases and exploiting synergies in data silos across departments becomes significantly easier (E14). Next, the unique properties of ML solutions further demand that corresponding business models engage in *ML-specific customer relationships*. Due to the technology’s complexity and the potential impacts of ML adoption, like ML taking over creative tasks from humans, solutions must be marketed with some finesse, laying a focus on communicating the added value of solutions and addressing both budget owners as well as experts with technical understanding (E01, E03, E08). To better understand how to approach and work with their customers, many ML start-ups in the B2B sector start with a consulting-oriented business model before transforming to a product- or service-based model (E03, E07, E08, E10, E12, E19). This not only helps them learn from the clients’ business problems but also from their data to enable improvements to the ML system (E07, E10). The consulting model further alleviates marketing challenges caused by the uncertainty of ML project outcomes: *“The problem with machine learning, the failure rate for machining applications is still 60-70%, somewhere there. That failure rate needs to come down to like 20%, and then enough companies will start adopting this on*

a scale that makes sense. And then once that happens, you will have people switch to understanding how to license the software and not need consultation behind it as well" (E08). Switching to product-based business models is aspired by many ML start-ups due to the higher scalability and disruptive potential (E08, E17, E20), whereby keeping the ML model hidden within the system a secret and only providing and selling its generated solutions is promising to be the "*holy grail*" (E17) of approaches, allowing for the best monetization (E08, E17). The fourth seizing microfoundation identified in the interviews is *intensive partnerships & ecosystem play*. For development partners of ML-driven business models to be willing to give out their sensitive data for the ML development process, mutual trust and strong contractual agreements must be in place (E01, E16). Customers only enter such partnerships "*if they know you as a provider or find you reputable and [...] if they have sufficient need to do something in that area*" (E01). Partnerships thus also help in knowing the customer and validating assumptions about the customer's need (E01, E16). Furthermore, the high dynamism of the ML market coupled with the increasing modularity of cloud and ML services (see Section 7.4.1) amplifies the benefits of strong collaboration and joint learning through ecosystem play, especially for small and medium-sized enterprises (E17, E18). Organizations should thus rather strive to "*build up many solutions [in a network] and see the strengths of them than acquire everything [themselves], [having] to build up a huge body of knowledge over many years. So I think I have to develop the ability to enter into this collaboration and say that 'he won't take away my water, he can support me' and then I think you have to develop these joint business models*" (E18).

Transforming

As ML-developing organizations experience influences of the technology across all departments (E11, E18), reconfiguring their structure towards *organizational structures enabling ML innovation* unlocks the full potential of ML. When developing novel ML business models within the organization, centralized teams are no longer recommended due to the resulting gap between domain experts and the ML team (E07, E16, E18). More suitable configurations include the hub-and-spoke model, which pools ML expertise in a hub that closely cooperates with experts throughout the organization that are close to the business side and endowed with technical understanding (E05, E07, E16, E18). For some established organizations, however, realizing an ML-driven business model that cannibalizes their own business can prove unsustainable within their company boundaries (E04, E05, E17). For example, "*lawyers sell hours. No lawyer in the world wants to become more efficient; that's the last thing he wants*" (E04). To circumvent this issue, such organizations may pursue corporate spin-offs to capture new parts of the market (E04, E05, E17). "*We are working with an insurer at the moment who's been around for hundreds of years as a company, and they're creating a new business which will be the claims experience itself which will be entirely driven by machine learning. [...] So, what they've done with that model is they've created a separate company, and they will eventually move out of the insurer, and they funded that company themselves. So, the insurer will still be a majority stakeholder, but they recognized that they couldn't create such a different business model within an old insurance business model*" (E05). This approach further helps in attracting talent as well as novel partnerships, maybe even with former competitors (E17), but also represents a large investment that takes a long time, "*about three to five years*" (E05), to pay off. Nevertheless, the experts advise highly disruptive business models: "*If [the ML-driven business model] becomes fundamental like that, spin it off*" (E04). Besides appropriate structures, ML-driven business model realization further benefits from *organizational culture fostering ML innovation*. While concerns often emerge when adopting novel technologies, especially ML-driven business models are often met with resistances (E06, E14, E16) due to their unique implications for humans at work (see Section 7.2.1). Realizing ML-driven

business models, however, calls for an organizational culture fostering curiosity and “*willingness to embrace change, to accept what makes ML solutions very different than other ones, which is the amount of uncertainty that [they] might imply. [...] There needs to be aware[ness] also of the risks that [uncertainty] might imply*” (E06). Members of the organization should be willing to take risks, accepting of failure, and championing trial-and-error approaches (E02, E05, E11, E13, E15, E17, E18). Building such a culture further sensitizes members of all departments for added value in data (E17) but may call for extensive change management efforts in established enterprises (E06). Another microfoundation of the transforming capability is the *top management commitment*. As the realization of ML-driven business models must be carried by all parts of an organization, committing to and communicating its strategic importance from the top management is crucial (E06, E13, E14, E16, E17). This is especially important in traditional organizations like “*insurance companies, for example: They don’t have any pressure to change at the moment. Business is good. They have nice margins on their traditional insurance policies. [...] Of course, there are a few fintech players who are just starting to act in a more innovative, faster, and more customer-centric way. But the pressure from outside has not yet been felt. That means there has to be an internal, intrinsic motivation to develop new business models, and there has to be oomph behind it from the board*” (E17). Furthermore, the top management must secure the budget for ML development (E01, E02, E06, E13), which is both costly (E06, E09, E11) as well as risky due to the outcome being uncertain (E09, E18). Thereby, laying a large focus on performance indicators and costs “*destroys the innovation*” (E13). Having top managers that understand the need for extensive and untied experimentation during ML development and provide the necessary budget thus benefits the business model realization (E06, E11, E13, E14, E18). Lastly, the microfoundation *cross-industry knowledge acquisition* describes an organization’s ability to gather insights through its network of partners. Due to ML being a general-purpose technology, these insights are not necessarily limited to the organization’s market, which makes the exchange with partners equipped with vast experience on different ML projects, such as consulting firms, very attractive (E05, E16, E17). Such exchange is fostered through strong internal champions who understand ML technologies and are driven to motivate for novel ML business ideas (E04, E06, E07, E19). Additionally, as the ML market is relatively new and changes within it can be far-reaching (see Section 7.4.1), e.g., due to new regulations, networking with the ML community and the scientific field is essential for maintaining competitive advantages (E02, E09).

7.5 Discussion

By adopting a dynamic capabilities perspective, we uncover the microfoundations that explain how organizations build the sensing, seizing, and transforming capabilities that empower them in realizing ML-driven business models. More specifically, we conceptualize three sensing, four seizing, and four transforming microfoundations (see Figure 7), answering our research question.

Our study makes multiple **theoretical contributions**. First, motivated by scholars stating that organizations must build dynamic capabilities for digital transformation efforts (Warner & Wäger, 2019; Wamba et al., 2017; Shollo et al., 2022; Mikalef et al., 2020) and to confront turbulent environments (Wilden & Gudergan, 2015), we contextualize and verify the suitability of the dynamic capabilities perspective for the realization of ML-driven business models. Second, with dynamic capabilities playing a large role in the success of organizations realizing business models (see Teece, 2018; Winter, 2003; Ricciardi et al., 2016), we illuminate the processes and practices underlying the dynamic capabilities of these organizations. More specifically, we theorize on microfoundations on which ML-specific dynamic capabilities are grounded, answering a call for

research on dynamic capabilities for digitalization (Vial, 2019). Thereby, we extend extant research on realizing data-driven business models revolving around models of the development process (e.g., Shollo et al., 2022; Rashed & Drews, 2021; Hunke et al., 2017) and the necessary resources and ordinary capabilities (e.g., Lange et al., 2021) by delineating how organizations build the ML-specific dynamic capabilities that enable them to orchestrate their resources in the realization process for sustainable success (see Teece, 2018). Third, we corroborate and expand the results of previous research stating that experimenting with business model ideas helps organizations reduce uncertainty (Andries et al., 2013). While organizations creating value from other digital technologies can tinker with novel business model ideas at very low costs (Huang et al., 2017; Lange et al., 2021), the development of ML technologies requires high up-front investments. Our results thus suggest that experimentation for the two development processes, ML development and business model development, must be interleaved to reduce the associated uncertainties simultaneously. Fourth, our results contribute to research building on the dynamic capabilities perspective in several ways. Our identified microfoundation of *data-driven sensing* gives support to recent literature reporting that organizations must increasingly digitize their sensing capabilities to quickly understand changes in highly dynamic markets (Nambisan et al., 2017; Warner & Wäger, 2019). The identified seizing microfoundation *continuous business model alignment* further underscores research emphasizing the need for rapid strategic realignment in such markets characterized by high velocity and uncertainty (Teece & Leih, 2016; Peteraf et al., 2013). Additionally, while the benefit of collaboration in inter-organizational networks to innovation has been shown in the literature (e.g., Schilling & Phelps, 2007), our results suggest an exacerbated need for ecosystem play in the ML market due to its increasing modularity and partnerships involving the exchange of data requiring high mutual trust between partners. Regarding transforming microfoundations, ML-driven business models need the involvement of all members of the organization as per our results and thus necessitate an organizational culture that emphasizes openness to change, cross-departmental collaboration, and sensitizes towards ML- and data-specific topics. These findings coincide with literature on digitalization efforts, which similarly require change management efforts to overcome resistances to change (Singh & Hess, 2017; Ellström et al., 2022; Warner & Wäger, 2019). Lastly, while many types of digital business ideas benefit from the agility offered by developing the new business model next to traditional structures in a bimodal organizational structure (Haffke et al., 2017b; Lange et al., 2021), we found that the increased capacity of ML-driven business models to cannibalize an organization's own business may more often demand realization outside of previous organizational boundaries. Furthermore, our study makes two major **practical contributions**. First, we guide organizations by conceptualizing the novel challenges that arise when realizing business ideas driven by ML technologies. Practitioners can thus circumvent these issues and their implications when designing their business model realization process. Second, practitioners can utilize the presented microfoundations as guidelines on how to build dynamic capabilities that empower ML-driven business model realization processes. We thereby enable decision-makers to identify and make the necessary changes within their organization to align it for optimal capacity to realize ML-driven business models.

Our study is subject to **limitations** that invite future research. As our focus lay on capturing a practitioner's perspective, we utilized the broad scope of our qualitative study to identify microfoundations grounded on expertise on a variety of business model realization cases. Yet, this leaves the quantitative measurement of the microfoundations' effectivity in building dynamic capabilities or the associated success of business model realization endeavors unresolved to date; a limitation shared by comparable research on dynamic capabilities (see Warner & Wäger, 2019;

Mousavi et al., 2019). Future research could therefore operationalize our microfoundations to both verify them and measure their effects in real-time longitudinal studies. Particularly their impact on the speed and degree of an organization's alignment to customer needs (see Teece, 2018) during different phases of the realization process could thereby be of interest. Moreover, future research might investigate whether involving consultants in the realization process of business models alters how dynamic capabilities are built. To further expand on this work, future research could examine the interaction of ordinary capabilities (see Lange et al., 2021) and dynamic capabilities during business model realization.

7.6 Conclusion

In summary, our study represents a first step in examining how organizations navigate the ML-driven business model realization process despite ML-specific complications. Based on data from 20 expert interviews with practitioners, we thus shed light on the dynamic capabilities that empower organizations in realizing ML-driven business models. As results, we conceptualize various microfoundations (see Figure 7) on which these dynamic capabilities are grounded and that are crucial for counteracting the identified ML-induced complications in the business environment and the business model itself. Yet, further research is needed to unpack the full potential of ML-driven business models.

Acknowledgment

This research and development project is/was funded by the German Federal Ministry of Education and Research (BMBF) within the "Innovations for Tomorrow's Production, Services, and Work" Program (funding number 02L19C150) and implemented by the Project Management Agency Karlsruhe (PTKA). The authors are responsible for the content of this publication.

8 Overarching Discussion and Conclusion

Modern ML systems hold immense economic potential for organizations to exploit. The vast number of associated entrepreneurial opportunities, in combination with the disruptive nature of ML innovations, both enable and pressure organizations to adapt their existing business models or develop entirely new ones driven by ML (e.g., Chalmers et al., 2021; Davenport et al., 2020). In this regard, while ML systems offer novel and unique opportunities in countless application scenarios along the value chain (e.g., Benbya et al., 2021; Burström et al., 2021), they also pose new challenges that differ from those of other digital technologies (e.g., Benbya et al., 2020b; Berente et al., 2021). When organizations thus integrate ML into their business models, they must adapt the way they create and deliver value for their customers while capturing value themselves, to harness opportunities unlocked by ML and account for ML-specific challenges (e.g., Ågerfalk, 2020; Amershi et al., 2019; Benbya et al., 2020a; Lange et al., 2021; Steininger et al., 2022; Sturm et al., 2021b). However, to date, research has left largely unexplored exactly how ML influences the business logic of organizations and how organizations adapt their business models to optimally utilize ML. Thus, this dissertation aims to further the understanding of the nature and unique potential of ML-driven business models for both scholars and practitioners. It expounds on the emerging ML-specific challenges influencing how organizations create, deliver, and capture value with ML systems, delineating the respective avenues by which organizations tackle these difficulties while exploiting the unique capabilities of ML to ensure their ML-driven business models' sustainable success. To this end, this dissertation draws on empirical evidence and several explorative research approaches, which entail some limitations (detailed in each of the papers above), yet provide various important contributions that can inspire practitioners and inform future studies in this underexplored area of research, which appear below.

8.1 Theoretical Contributions

This dissertation illuminates the influence of ML on the business logic of organizations and provides various contributions to research, which are detailed in the three dimensions of value creation, value delivery, and value capture in the following.

First, concerning **value creation**, this dissertation identifies both ML-induced challenges complicating the creation of value in ML-driven business models, and unique ways in which ML systems can provide additional value during their development. Regarding the latter, the results of Paper A show that when organizations conduct ML development projects, they can create value not only upon successful completion through the productively usable ML systems at the end of the development process, but are also able to stimulate learning processes among the organization's domain experts during all phases of the development. In other words, ML development itself can be of value to the organization by fostering organizational learning, a process critical to maintaining and increasing its knowledge base and thus vital to its long-term performance (e.g., Argote et al., 2021; March, 1991). To date, few studies have investigated the learning processes involving ML systems and human learners, yet with a focus on productive ML

systems and the organizational learning they enable (e.g., Balasubramanian et al., 2022; Lyytinen et al., 2021; Ransbotham et al., 2020; Seidel et al., 2019; Sturm et al., 2021b). Paper A thus contributes to the current discussion of how ML impacts organizational learning, and highlights how organizations can create valuable knowledge as a byproduct of their value creation efforts to develop productively usable ML systems. In particular, the findings show that organizations can employ the development of ML systems as a novel tool to initiate the revision of suboptimal or outdated knowledge, which is notoriously difficult in organizational learning (e.g., Argote et al., 2021; Argyris, 1976; Levinthal & March, 1993), and as a mechanism to facilitate the similarly challenging management of tacit knowledge (e.g., Argote et al., 2021; Nonaka, 1994). In this regard, recent literature indicates that productive generative AI may play a more central role in future knowledge management (Alavi et al., 2024) and merit a reexamination in further studies of potential learning processes during projects aimed at developing or adapting generative AI systems for the organization. With respect to the challenges in the way of value creation with ML, papers A, B, and E emphasize the need for extensive experimentation in ML development. This need arises from both domain experts' attempts to integrate their own knowledge with insights and solutions from the development process, to simultaneously evaluate the system and their knowledge base (paper A), and the uncertainty inherent in ML development, in which software and data are inextricably linked due to the data-based learning approach of ML systems (paper E). However, the demand for experimentation runs counter to the demands of many conventional business processes for exploiting existing strengths (e.g., Gerbert et al., 2020; Pumplun et al., 2019), creating tensions during ML development (paper B). Therefore, Paper B identifies four distinct tensions arising when organizations create value through ML developments. They concern, for instance, the type of knowledge that ML systems are to generate, between granting the ML system the freedom to learn autonomously (generating independent, potentially novel knowledge) and more strictly guiding its learning process through domain experts (generating knowledge relevant to the domain; e.g., van den Broek et al., 2021; Seidel et al., 2020). Moreover, the paper presents exploration and exploitation tactics that organizations use to alleviate the problematic poles of the tensions, depending on the chosen organizational structure for ML development. An example is helping organizations with inherently exploitation-oriented decentralized structures to satisfy the demand for experimentation through workshops or open time during working hours, motivating business units to think "outside of the box". Furthermore, these alleviating tactics cause organizations to shift their organizational structure toward hybrid approaches, underscoring the importance of satisfying the ML demand for extensive experimentation as part of the exploitation-exploration duality during ML development, to ensure long-term performance through ambidexterity (see, e.g., Raisch & Birkinshaw, 2008; Smith & Lewis, 2011; Smith et al., 2011). Collectively, the conceptualizations within these papers can be used in future research on the management of ML-driven value creation efforts, which I hope to inspire to consider both the challenges and the unique avenues for learning that ML's data-based learning and the associated need for experimentation present.

Second, on the topic of **value delivery**, the results reported in this dissertation show that ML-driven business models must tailor their offerings to their customers, both concerning how the value generated by the ML system is communicated to its users through XAI approaches (paper C), and concerning the necessary adaptation of the provided ML system to the customer's data and business problem at hand (paper E). Regarding the latter, organizations employing ML-driven business models must gain a detailed understanding of both data and business problems of their clients. Accordingly, they often initially employ a consulting-based approach, to learn alongside their clients, or enter close development partnerships to enable data sharing and co-development,

as paper E details. While the positive effects of inter-organizational collaboration in networks on innovation appear in the literature (e.g., Schilling & Phelps, 2007), paper E corroborates initial research suggesting the heightened importance of ecosystem play for delivering value through ML-driven business models (Burström et al., 2021; Weber et al., 2022). In addition, paper C furthers our understanding of how organizations can ensure that the potential value their ML system can provide is also receivable by and of benefit to the system's users through user-specific explanations. Only recently has research on XAI begun considering the needs of lay users who must understand the ML system's output to incorporate it into their decision-making, such as doctors making diagnostic decisions in the healthcare sector (e.g., Ellenrieder et al., 2023; Gaube et al., 2023). By carefully delineating the user groups of an ML system for a forecasting use case in the corporate performance management domain and considering their requirements before developing an approach to designing XAI, the paper thus makes an important contribution to an ongoing scholarly effort to provide explainability specifically tailored to the various users of ML systems. In particular, the defined user groups include such lay users as managers who urgently need explanations to consider the ML system's output in their decision-making (e.g., Barredo Arrieta et al., 2020; Martin, 2019b). The design principles the paper presents can thus inform future XAI research on the design of user-centric explanations provided by ML systems, both in the field of corporate performance management specifically and in further domains with comparable forecasting use cases.

Third, with respect to **value capture**, this dissertation offers insight into the anatomy of ML-driven business models through which organizations in practice aim to capture value sustainably (paper D) and fosters our understanding of how organizations build the dynamic capabilities to successfully realize these business models (paper E). Regarding the former, paper D develops a taxonomy of ML-driven business models, answering calls from the literature for taxonomic research related to data-driven business models (e.g., Veit et al., 2014; Müller & Buliga, 2019) and complementing the discussion with ML-focused observations. The taxonomy provides researchers with a standardized vocabulary to describe and differentiate ML-driven business models found in practice in a common language. Thus, it facilitates the future scholarly exchange of new ideas on the topic and contributes to a unified understanding of ML-driven business models. In addition, this dissertation supplements insights into frequently employed ML-driven business models and their constituent components, through the identified archetypes in paper D, as well as information on the characteristics and particularities distinctly inherent in ML-driven business models and their environment, which complicate their realization (paper E). Building on these findings, paper E presents an in-depth examination of how organizations can nevertheless successfully realize ML-driven business models. In particular, it theorizes on eleven microfoundations (i.e., practices and processes) with which organizations build the dynamic capabilities that allow them to sustainably capture value through ML-driven business model realization, despite ML-specific challenges. The study finds that while all organizations benefit from experimenting with business model ideas in the market to reduce uncertainty and ensure their suitability for eventual value capture (e.g., Andries et al., 2013), the data-based learning approach of ML systems adds an additional layer of uncertainty inherent in the technology. In addition, while other digital products and services can be experimented with at relatively low costs (e.g., Huang et al., 2017; Lange et al., 2021), the development of ML entails high up-front investments. The results of the paper thus suggest that organizations are well advised to interleave and simultaneously conduct both the ML development processes and the business model development process to reduce their respective uncertainties in tandem. The paper thereby contributes to a thriving research stream on the learning processes that foster successful

business model realization by reducing uncertainty around markets, products, and activities (e.g., experimentation, trial-and-error learning, vicarious learning; see Snihur & Eisenhardt, 2022), underscoring the importance of addressing ML-induced uncertainties in the case of ML-driven business models. In this context, future studies may build on these findings by examining in depth how organizations realizing ML-driven business models account for these uncertainties in their learning processes. In addition, drawing on the findings of paper A, studies that focus on the interplay between these learning processes about markets, products, and activities during business model development and the simultaneous learning processes during ML development might also prove interesting. Furthermore, the results of paper E show that the capacity to tackle environmental dynamism is particularly important for ML-driven business models because their business environments are commonly inherently complex and fast-moving (e.g., Chalmers et al., 2021; Steininger et al., 2022), and additionally because their ML systems can also be affected by said dynamism when learning from data from the environment. The identified microfoundation data-driven sensing can support in this endeavor and emphasizes the need for organizations to listen to their data when determining which opportunities to target in the market. Recent research finds this beneficial for all organizations in fast-moving environments (e.g., Nambisan et al., 2017; Warner & Wäger, 2019). Yet, the findings of paper E deem it crucial for ML-driven business models, due to their reliance on data as a value driver. Moreover, realizing a novel ML-driven business model out of an existing business model can fundamentally change how an established organization does business. In such cases, the results show that the ML-driven business model has an increased capacity to cannibalize the organization's core business, potentially meriting realization outside of existing organizational boundaries. The results further indicate that scaling ML-driven business models can prove difficult, e.g., due to the required training efforts for new (groups of) customers, corroborating recent research on the topic (Haefner et al., 2023; Zebhauser, 2024). They suggest that new ventures are thus incentivized to start with a consulting-oriented approach, as mentioned above, which limits their ability to capture value, and have difficulty transitioning to a more scalable product-oriented approach, meriting additional studies on the avenues through which ML-driven business models can perform the desirable transformation. Taken together, the findings of papers D and E thus provide to future studies insights and a common language for examining how new ventures powered by ML can succeed and for investigating the factors explicating the long-term success of ML-driven organizations.

Overall, this dissertation thus contributes to a growing understanding of ML-driven business models, both by verifying that the ML-specific characteristics and challenges indeed change the logic underlying the value creation, value delivery, and value capture of organizations (e.g., Benbya et al., 2020a; Berente et al., 2021; Burström et al., 2021; Steininger et al., 2022) and particularly by elaborating on concrete ways they do so in each of the three dimensions. It is my hope that the results this dissertation presents form a foundation that inspires scholars to further investigate the demand of ML systems for unique consideration in the business logic of organizations, as well as their great potential to enable and power novel types of business models.

8.2 Practical Contributions

In addition to its theoretical contributions, this dissertation yields practical guidance on various aspects of the creation, delivery, and capture of value with ML systems. It also holds particularly valuable insights for practitioners aiming to realize their own ML-driven business models. Two major contributions to practice are discussed below.

First, organizations can utilize the findings and artifacts contained in this dissertation to inform a multitude of design decisions arising in the context of their ML development process (papers B and C) and their ML-driven business model realization process (papers D and E). Regarding the former, paper B outlines the effects of structural decisions on ML development and empowers practitioners to identify the tensions that arise when developing and operating ML systems, depending on their chosen structure. After their identification, the paper further provides organizations with specific tactics they can use to alleviate these tensions. Concerning the practical contributions to ML development, paper C focuses on the design of the ML system itself. Accordingly, it offers design principles in three decomposition layers, which organizations that aim to provide explainability through their ML system can use as concrete guidelines for designing their XAI approaches, particularly for tailoring them to their system's user groups. Organizations can further utilize the evaluation framework that paper C presents to evaluate their ML systems for their capacity to provide suitable explanations to their user groups and, in turn, for their ability to address the lack of explainability standing in the way of increased use of ML systems in practice (see Chui et al., 2021). With respect to guidance on the ML-driven business model realization process, paper D first offers an artifact, in the form of the developed taxonomy, that organizations can use to create blueprints of ML-driven business models. Thus, the taxonomy can function as a supporting tool in their business model innovation efforts to create novel ML-driven business models (Bouwman et al., 2020) and, in conjunction with the identified archetypes, as an evaluation instrument for their own or others' ML-driven business models. When taking the next step of realizing ML-driven business models, organizations can then apply the dynamic capabilities microfoundations that paper E presents as templates for practices and processes to implement, setting up the realization process for success. Equipped with these insights, decision-makers can thus identify and make the required changes in their organization to maximize its capacity to successfully realize ML-driven business models.

Second, this dissertation highlights the value of experimentation during multiple periods in an organization's journey toward an ML-driven business model. Beginning with the development of the ML systems, paper B emphasizes that ML development itself necessitates extensive experimentation in fast and exploratory processes, a demand that organizations are well advised to meet. For instance, tactical solutions, such as the organization's leadership creating the necessary space for experimentation during working hours and for their employees' personal development, can serve in this regard. This dovetails nicely with the findings of paper A, which support granting the domain experts involved in ML development additional time and resources for exploratory thinking, analyses, and discussions. These can yield immense value in the form of knowledge, generated through organizational learning processes, which the experts returning to their domains disseminate in the organization. However, as paper E shows, the demand for experimentation extends beyond ML development, as organizations that realize business models with ML at their core must simultaneously reduce uncertainties from the side of the market (e.g., whether intended customers will even find their planned offering worthwhile). Therefore, papers B and E inform organizations of avenues along the business model realization processes for fostering experimentation, e.g., by letting data guide both ML development and business model experimentation, by having the top management designate an appropriate budget allowing the freedom to experiment, or by creating an organizational culture commending risk-taking and fail-fast approaches. Complementing these results with those of paper A, this dissertation thus encourages organizations to both consider the important contribution of experimentation to the success of ML-driven business model realization, and view failed attempts as generators of

valuable knowledge through organizational learning, a powerful motor to help organizations achieve performance gains in the long term.

8.3 Concluding Remarks

Organizations incorporating ML systems into their business models can benefit immensely from the distinct power of these learning machines, that have long transcended their use in organizations alone and now permeate most aspects of human experience, for instance, steering autonomous vehicles around pedestrians or helping us select which movies to watch (e.g., Benbya et al., 2020b; Berente et al., 2021; Lyytinen et al., 2021; Yoo, 2010). However, intertwined with their marvelous capabilities come unique particularities that demand consideration in the business logic of ML-infused organizations (e.g., Benbya et al., 2020a; Burström et al., 2021; Steininger et al., 2022). To further our still sparse understanding of ML-driven business models, this dissertation illuminates their distinguishing aspects and examines in depth how organizations utilize ML to create and deliver value for their customers and capture value for themselves. While this dissertation thus constitutes an important step in shedding light on factors explicating the success of ML-infused organizations for both practitioners and scholars, further research is required to comprehensively reveal changes to the business logic of all kinds of organizations in the era of ML.

References

- Abiteboul, S., Buneman, P., & Suciu, D. (1999). *Data on the web: From relations to semistructured data and XML*. Morgan Kaufmann.
- Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- Adya, M. P., & Collopy, F. L. (1998). How effective are neural networks at forecasting and prediction? A review and evaluation. *Journal of Forecasting*, 17, 481–495. [https://doi.org/10.1002/\(SICI\)1099-131X\(1998090\)17:5/6%3C481::AID-FOR709%3E3.0.CO;2-Q](https://doi.org/10.1002/(SICI)1099-131X(1998090)17:5/6%3C481::AID-FOR709%3E3.0.CO;2-Q)
- Afuah, A. (2003). Redefining firm boundaries in the face of the internet: Are firms really shrinking? *Academy of Management Review*, 28(1), 34–53. <https://doi.org/10.2307/30040688>
- Agarwal, R., & Dhar, V. (2014). Editorial—big data, data science, and analytics: The opportunity and challenge for IS research. *Information Systems Research*, 25(3), 443–448. <http://doi.org/10.1287/isre.2014.0546>
- Ågerfalk, P. J. (2020). Artificial intelligence as digital agency. *European Journal of Information Systems*, 29(1), 1–8. <https://doi.org/10.1080/0960085X.2020.1721947>
- Alavi, M., & Leidner, D. E. (2001). Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS Quarterly*, 25(1), 107–136. <http://doi.org/10.2307/3250961>
- Alavi, M., Leidner, D. E., & Mousavi, R. (2024). A knowledge management perspective of generative artificial intelligence. *Journal of the Association for Information Systems*, 25(1), 1–12. <http://doi.org/10.17705/1jais.00859>
- Al-Debei, M. M., & Avison, D. (2010). Developing a unified framework of the business model concept. *European Journal of Information Systems*, 19(3), 359–376. <https://doi.org/10.1057/ejis.2010.21>
- AltaML. (2021, November 16). *Elevating business through AI*. <https://www.altaml.com>
- Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., & Zimmermann, T. (2019). Software engineering for machine learning: A case study. *Proceedings of the 41st IEEE/ACM International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, 291–300. <https://doi.org/10.1109/ICSE-SEIP.2019.00042>
- Anderberg, M. R. (1973). *Cluster analysis for applications: Probability and mathematical statistics: A series of monographs and textbooks*. Academic Press. <https://doi.org/10.1016/C2013-0-06161-0>

- Andries, P., Debackere, K., & van Looy, B. (2013). Simultaneous experimentation as a learning strategy: Business model development under uncertainty. *Strategic Entrepreneurship Journal*, 7(4), 288–310. <https://doi.org/10.1002/sej.1170>
- Anton, E., Oesterreich, T. D., Schuir, J., Protz, L., & Teuteberg, F. (2021). A business model taxonomy for start-ups in the electric power industry — the electrifying effect of artificial intelligence on business model innovation. *International Journal of Innovation and Technology Management*, 18(03), Article 2150004. <https://doi.org/10.1142/S0219877021500048>
- Arend, R. J. (2013). The business model: Present and future—beyond a skeumorph. *Strategic Organization*, 11(4), 390–402. <https://doi.org/10.1177/1476127013499636>
- Argote, L., Lee, S., & Park, J. (2021). Organizational learning processes and outcomes: Major findings and future research directions. *Management Science*, 67(9), 5399–5429. <https://doi.org/10.1287/mnsc.2020.3693>
- Argote, L., & Miron-Spektor, E. (2011). Organizational learning: From experience to knowledge. *Organization Science*, 22(5), 1123–1137. <https://doi.org/10.1287/orsc.1100.0621>
- Argyris, C. (1976). Single-loop and double-loop models in research on decision making. *Administrative Science Quarterly*, 21(3), 363–375. <https://doi.org/10.2307/2391848>
- Armour, J., & Sako, M. (2020). AI-enabled business models in legal services: From traditional law firms to next-generation law companies? *Journal of Professions and Organization*, 7(1), 27–46. <https://doi.org/10.1093/jpo/joaa001>
- Asatiani, A., Malo, P., Nagbøl, P., Penttinen, E., Rinta-Kahila, T., & Salovaara, A. (2021). Sociotechnical envelopment of artificial intelligence: An approach to organizational deployment of inscrutable artificial intelligence systems. *Journal of the Association for Information Systems*, 22(2), 325–352. <https://doi.org/10.17705/1jais.00664>
- Aubert, B., & Chan, Y. (2024). Evolving strategic IS themes. *The Journal of Strategic Information Systems*, 33(1), Article 101821. <https://doi.org/10.1016/j.jsis.2024.101821>
- Azevedo, A., & Santos, M. F. (2008). KDD, SEMMA and CRISP-DM: A parallel overview. *Proceedings of the IADIS European Conference on Data Mining*, 182–185. <https://www.iadisportal.org/digital-library/kdd-semma-and-crisp-dm-a-parallel-overview>
- Azkan, C., Iggena, L., Gür, I., Möller, F., & Otto, B. (2020). A taxonomy for data driven services in manufacturing industries. *Proceedings of the 24th Pacific Asia Conference on Information Systems (PACIS)*, Paper 184. <https://aisel.aisnet.org/pacis2020/184>
- Bačić, D., & Fadlalla, A. (2016). Business information visualization intellectual contributions: An integrative framework of visualization capabilities and dimensions of visual intelligence. *Decision Support Systems*, 89, 77–86. <http://doi.org/10.1016/j.dss.2016.06.011>
- Baecker, J., Böttcher, T. P., & Weking, J. (2021). How companies create value from data – a taxonomy on data, approaches, and resulting business value. *Proceedings of the 29th European Conference on Information Systems (ECIS)*, Paper 124. https://aisel.aisnet.org/ecis2021_rp/124
- Baier, L., Jöhren, F., & Seebacher, S. (2019). Challenges in the deployment and operation of machine learning in practice. *Proceedings of the 27th European Conference on Information Systems (ECIS)*, Paper 163. <https://doi.org/10.5445/IR/1000095028>

- Bailey, K. D. (1994). *Typologies and taxonomies: An introduction to classification techniques*. SAGE Publications.
- Baird, A., & Maruping, L. M. (2021). The next generation of research on IS use: A theoretical framework of delegation to and from agentic IS artifacts. *MIS Quarterly*, 45(1), 315–341. <https://doi.org/10.25300/MISQ/2021/15882>
- Balasubramanian, N., Ye, Y., & Xu, M. (2022). Substituting human decision-making with machine learning: Implications for organizational learning. *Academy of Management Review*, 47(3), 448–465. <https://doi.org/10.5465/amr.2019.0470>
- Barker, J. (2020). Machine learning in M4: What makes a good unstructured model? *International Journal of Forecasting*, 36(1), 150–155. <https://doi.org/10.1016/j.ijforecast.2019.06.001>
- Barney, J. B. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120. <https://doi.org/10.1177/014920639101700108>
- Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Barreto, I. (2010). Dynamic capabilities: A review of past research and an agenda for the future. *Journal of Management*, 36(1), 256–280. <http://doi.org/10.1177/0149206309350776>
- Basadur, M., Graen, G. B., & Green, S. G. (1982). Training in creative problem solving: Effects on ideation and problem finding and solving in an industrial research organization. *Organizational Behavior and Human Performance*, 30, 41–70. [https://doi.org/10.1016/0030-5073\(82\)90233-1](https://doi.org/10.1016/0030-5073(82)90233-1)
- Baumer, E. P. S., Mimno, D., Guha, S., Quan, E., & Gay, G. K. (2017). Comparing grounded theory and topic modeling: Extreme divergence or unlikely convergence? *Journal of the Association for Information Systems and Technology*, 68(6), 1397–1410. <https://doi.org/10.1002/asi.23786>
- Benbya, H., Davenport, T. H., & Pachidi, S. (2020a). Artificial intelligence in organizations: Current state and future opportunities. *MIS Quarterly Executive*, 19(4), 9–21. <https://doi.org/10.2139/ssrn.3741983>
- Benbya, H., Ning, N., Tanriverdi, H., & Youngjin, Y. (2020b). Complexity and information systems research in the emerging digital world. *MIS Quarterly*, 44(1), 1–17. <http://doi.org/10.25300/MISQ/2020/13304>
- Benbya, H., Pachidi, S., & Jarvenpaa, S. L. (2021). Special issue editorial: Artificial intelligence in organizations: Implications for information systems research. *Journal of the Association for Information Systems*, 22(2), 281–303. <https://doi.org/10.17705/1jais.00662>
- Benbya, H., Strich, F., & Tamm, T. (2024). Navigating generative artificial intelligence promises and perils for knowledge and creative work. *Journal of the Association for Information Systems*, 25(1), 23–36. <https://doi.org/10.17705/1jais.00861>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Special issue editor's comments: Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450. <http://doi.org/10.25300/MISQ/2021/16274>

- Bi, X., Adomavicius, G., Li, W., & Qu, A. (2022). Improving sales forecasting accuracy: A tensor factorization approach with demand awareness. *INFORMS Journal on Computing*, 34(3), 1644–1660. <https://doi.org/10.1287/ijoc.2021.1147>
- Bidnamic. (2022, March 22). *Google shopping management software*. <https://www.bidnamic.com>
- Bigelow, L. S., & Barney, J. B. (2021). What can strategy learn from the business model approach? *Journal of Management Studies*, 58(2), 528–539. <https://doi.org/10.1111/joms.12579>
- Birkinshaw, J., & Ansari, S. (2015). Understanding management models: Going beyond “what” and “why” to “how” work gets done in organizations. In N. J. Foss & T. Saebi (Eds.), *Business model innovation: The organizational dimension* (pp. 85–103). Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780198701873.003.0005>
- Birks, D. F., Fernandez, W., Levina, N., & Nasirin, S. (2013). Grounded theory method in information systems research: Its nature, diversity and opportunities. *European Journal of Information Systems*, 22(1), 1–8. <https://doi.org/10.1057/ejis.2012.48>
- Björkdahl, J. (2020). Strategies for digitalization in manufacturing firms. *California Management Review*, 62(4), 17–36. <https://doi.org/10.1177/0008125620920349>
- Bock, M., & Wiener, M. (2017). Towards a taxonomy of digital business models – conceptual dimensions and empirical illustrations. *Proceedings of the 38th International Conference on Information Systems (ICIS)*. <https://aisel.aisnet.org/icis2017/Strategy/Presentations/19>
- Bogner, A., Littig, B., & Menz, W. (2009). *Interviewing experts*. Palgrave Macmillan. <https://doi.org/10.1057/9780230244276>
- Bontempi, G., Taieb, S. B., & Le Borgne, Y. A. (2013). Machine learning strategies for time series forecasting. In M.-A. Aufaure & E. Zimányi (Eds.), *Lecture notes in business information processing* (Vol. 138, pp. 62–77). Springer. https://doi.org/10.1007/978-3-642-36318-4_3
- Bostrom, N. (2017). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.
- Böttcher, T. P., Weking, J., Hein, A., Böhm, M., & Krcmar, H. (2022). Pathways to digital business models: The connection of sensing and seizing in business model innovation. *The Journal of Strategic Information Systems*, 31(4), Article 101742. <https://doi.org/10.1016/j.jsis.2022.101742>
- Bouwman, H., de Reuver, M., Heikkilä, M., & Fiel, E. (2020). Business model tooling: Where research and practice meet. *Electronic Markets*, 30(3), 413–419. <https://doi.org/10.1007/s12525-020-00424-5>
- Brennen, A. (2020). What do people really want when they say they want “explainable AI?” we asked 60 stakeholders. *Extended Abstracts of the 2020 ACM Conference on Human Factors in Computing Systems (CHI EA)*, 1–7. <https://doi.org/10.1145/3334480.3383047>
- van den Broek, E., Sergeeva, A., & Huysman, M. (2021). When the machine meets the expert: An ethnography of developing AI for hiring. *MIS Quarterly* 45(3), 1557–1580. <https://doi.org/10.25300/MISQ/2021/16559>
- Brynjolfsson, E., & McAfee, A. (2017, July 18). The business of artificial intelligence. *Harvard Business Review*. <https://hbr.org/2017/07/the-business-of-artificial-intelligence>
- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370), 1530–1534. <https://doi.org/10.1126/science.aap8062>

- Burkhart, T., Krumeich, J., Werth, D., & Loos, P. (2011). Analyzing the business model concept — a comprehensive classification of literature. *Proceedings of the 32nd International Conference on Information Systems (ICIS)*. <https://aisel.aisnet.org/icis2011/proceedings/generaltopics/12>
- Burnell, D., Stevenson, R., & Fisher, G. (2023). Early-stage business model experimentation and pivoting. *Journal of Business Venturing*, 38(4), Article 106314. <https://doi.org/10.1016/j.jbusvent.2023.106314>
- Burr, J. (2023, January 19). *9 ways we use AI in our products*. Google. <https://blog.google/technology/ai/9-ways-we-use-ai-in-our-products>
- Burström, T., Parida, V., Lahti, T., & Wincent, J. (2021). AI-enabled business-model innovation and transformation in industrial ecosystems: A framework, model and outline for further research. *Journal of Business Research*, 127, 85–95. <https://doi.org/10.1016/j.jbusres.2021.01.016>
- Butts, M. M., Becker, W. J., & Boswell, W. R. (2015). Hot buttons and time sinks: The effects of electronic communication during nonwork time on emotions and work-nonwork conflict. *Academy of Management Journal*, 58(3), 763–788. <http://doi.org/10.5465/amj.2014.0170>
- Carter, N., Bryant-Lukosius, D., Dicenso, A., Blythe, J., & Neville, A. J. (2014). The use of triangulation in qualitative research. *Oncology Nursing Forum*, 41(5), 545–547. <https://doi.org/10.1188/14.onf.545-547>
- Casadesus-Masanell, R., & Ricart, J. E. (2010). From strategy to business models and onto tactics. *Long Range Planning*, 43(2-3), 195–215. <https://doi.org/10.1016/j.lrp.2010.01.004>
- Casadesus-Masanell, R., & Ricart, J. E. (2011). How to design a winning business model. *Harvard Business Review*, 89(1/2), 100–107.
- Castelvecchi, D. (2016). Can we open the black box of AI? *Nature*, 538, 20–23. <https://doi.org/10.1038/538020a>
- Chalmers, D., MacKenzie, N. G., & Carter, S. (2021). Artificial intelligence and entrepreneurship: Implications for venture creation in the fourth industrial revolution. *Entrepreneurship Theory and Practice*, 45(5), 1028–1053. <https://doi.org/10.1177/1042258720934581>
- Charmaz, K., & Belgrave, L. L. (2012). Qualitative interviewing and grounded theory analysis. In J. F. Gubrium, J. A. Holstein, A. B. Marvasti, and K. D. McKinney (Eds.), *The SAGE handbook of interview research: The complexity of the craft* (2nd ed., pp. 347–365). SAGE Publications. <http://doi.org/10.4135/9781452218403.n25>
- Choudhury, P., Allen, R. T., & Endres, M. G. (2021). Machine learning for pattern discovery in management research. *Strategic Management Journal*, 42(1), 30–57. <https://doi.org/10.1002/smj.3215>
- Chui, M., Hall, B., Singla, A., & Sukharevsky, A. (2021, December 8). *The state of AI in 2021*. McKinsey. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/global-survey-the-state-of-ai-in-2021>
- Circuit Mind. (2021, November 16). *AI electronics design software*. <https://www.circuitmind.io>
- CognitOps. (2021, November 16). *Automate warehouse management*. <https://cognitops.com>

- Corbin, J., & Strauss, A. (2015). *Basics of qualitative research: Techniques and procedures for developing grounded theory* (4th ed.). SAGE Publications.
- Correa-Baena, J. P., Hippalgaonkar, K., van Duren, J., Jaffer, S., Chandrasekhar, V. R., Stevanovic, V., Wadia, C., Guha, S., & Buonassisi, T. (2018). Accelerating materials development via automation, machine learning, and high-performance computing. *Joule*, 2(8), 1410–1420. <https://doi.org/10.1016/j.joule.2018.05.009>
- Crunchbase. (2021, March 13). *Search less. Close more.* <https://www.crunchbase.com>
- DaSilva, C. M., & Trkman, P. (2014). Business model: What it is and what it is not. *Long range planning*, 47(6), 379–389. <https://doi.org/10.1016/j.lrp.2013.08.004>
- Davenport, T. H. (2018). From analytics to artificial intelligence. *Journal of Business Analytics*, 1(2), 73–80. <https://doi.org/10.1080/2573234X.2018.1543535>
- Davenport, T. H., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
- Davies, D. L., & Bouldin, D. W. (1979). A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-1(2), 224–227. <https://doi.org/10.1109/TPAMI.1979.4766909>
- Davison, R. M., & Martinsons, M. G. (2016). Context is king! Considering particularism in research design and reporting. *Journal of Information Technology*, 31(3), 241–249. <https://doi.org/10.1057/jit.2015.19>
- DeepL. (2023, January 13). *Partnering for progress: DeepL's latest funding round will fuel AI research and development.* <https://www.deepl.com/en/blog/latest-funding-round>
- Dehnert, M., Gleiss, A., & Reiss, F. (2021). What makes a data-driven business model? A consolidated taxonomy. *Proceedings of the 29th European Conference on Information Systems (ECIS)*, Paper 139. https://aisel.aisnet.org/ecis2021_rp/139
- Dempster, M. A. H., & Leemans, V. (2006). An automated FX trading system using adaptive reinforcement learning. *Expert Systems with Applications*, 30(3), 543–552. <https://doi.org/10.1016/j.eswa.2005.10.012>
- Dennett, D. C. (2006). Cognitive wheels: The frame problem of AI. In J. L. Bermúdez (Ed.), *Philosophy of psychology: Contemporary readings* (pp. 433–454). Routledge/Taylor & Francis Group.
- Di Stefano, G., Peteraf, M., & Verona, G. (2014). The organizational drivetrain: A road to integration of dynamic capabilities research. *Academy of Management Perspectives*, 28(4), 307–327. <https://doi.org/10.5465/amp.2013.0100>
- Dingli, A., Haddod, F., & Klüver, C. (Eds.). (2021). *Artificial intelligence in industry 4.0. A collection of innovative research case-studies that are reworking the way we look at industry 4.0 thanks to artificial intelligence.* Springer. <https://doi.org/10.1007/978-3-030-61045-6>
- Dorfer, L. (2016). Datenzentrische Geschäftsmodelle als neuer Geschäftsmodelltypus in der Electronic-Business-Forschung: Konzeptionelle Bezugspunkte, Klassifikation und Geschäftsmodellarchitektur. *Schmalenbachs Zeitschrift Für Betriebswirtschaftliche Forschung*, 68(3), 307–369. <https://doi.org/10.1007/s41471-016-0014-9>

- Dotzel, T., & Shankar, V. (2019). The relative effects of business-to-business (vs. business-to-consumer) service innovations on firm value and firm risk: An empirical analysis. *Journal of Marketing*, *83*(5), 133–152. <https://doi.org/10.1177/0022242919847221>
- EAGLE. (2021, November 16). *Visitor management | wellness screening*. <https://www.eagle.com>
- Eisenhardt, K. M., & Martin, J. A. (2000). Dynamic capabilities: What are they? *Strategic Management Journal*, *21*(10–11), 1105–1121. [https://doi.org/10.1002/1097-0266\(200010/11\)21:10/11%3C1105::AID-SMJ133%3E3.0.CO;2-E](https://doi.org/10.1002/1097-0266(200010/11)21:10/11%3C1105::AID-SMJ133%3E3.0.CO;2-E)
- Ellenrieder, S., Kallina, E. M., Pumplun, L., Gawlitza, J. F., Ziegelmayr, S., & Buxmann, P. (2023). Promoting learning through explainable artificial intelligence: An experimental study in radiology. *Proceedings of the 44th International Conference on Information Systems (ICIS)*. <https://aisel.aisnet.org/icis2023/learnandiscurrecula/learnandiscurrecula/3>
- Ellström, D., Holtström, J., Berg, E., & Josefsson, C. (2022). Dynamic capabilities for digital transformation. *Journal of Strategy and Management*, *15*(2), 272–286. <http://doi.org/10.1108/JSMA-04-2021-0089>
- Engelbrecht, A., Gerlach, J., & Widjaja, T. (2016). Understanding the anatomy of data-driven business models - towards an empirical taxonomy. *Proceedings of the 24th European Conference on Information Systems (ECIS)*, Paper 128. https://aisel.aisnet.org/ecis2016_rp/128
- Eszergár-Kiss, D., & Caesar, B. (2017). Definition of user groups applying Ward's method. *Transportation Research Procedia*, *22*, 25–34. <https://doi.org/10.1016/j.trpro.2017.03.004>
- Fang, C., Lee, J., & Schilling, M. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science*, *21*(3), 625–642. <https://doi.org/10.1287/orsc.1090.0468>
- Faraj, S., Pachidi, S., & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization*, *28*(1), 62–70. <https://doi.org/10.1016/j.infoandorg.2018.02.005>
- Farjoun, M. (2010). Beyond dualism: Stability and change as a duality. *Academy of Management Review*, *35*(2), 202–225. <https://doi.org/10.5465/amr.35.2.zok202>
- Fiol, C., & Lyles, M. (1985). Organizational learning. *Academy of Management Review*, *10*(4), 803–813. <https://doi.org/10.5465/amr.1985.4279103>
- Flick, U. (2004a). *Triangulation: Eine Einführung*. VS Verlag für Sozialwissenschaften. <https://doi.org/10.1007/978-3-322-97512-6>
- Flick, U. (2004b). Triangulation in qualitative research. In U. Flick, E. von Kardoff, & I. Steinke (Eds.), *A companion to qualitative research* (pp. 178–183). SAGE Publications.
- Foss, N. J., & Saebi, T. (Eds.). (2015). *Business model innovation: The organizational dimension*. Oxford University Press.
- Frolick, M. N., & Ariyachandra, T. R. (2006). Business performance management: One truth. *Information Systems Management*, *23*(1), 41–48. <http://doi.org/10.1201/1078.10580530/45769.23.1.20061201/91771.5>
- Gartner. (2024, February 20). *Get AI ready — what IT leaders need to know and do*. <https://www.gartner.com/en/information-technology/topics/ai-readiness>

- Gaube, S., Suresh, H., Raue, M., Lermer, E., Koch, T. K., Hudecek, M. F. C., Ackery, A. D., Grover, S. C., Coughlin, J. F., Frey, D., Kitamura, F. C., Ghassemi, M., & Colak, E. (2023). Non-task expert physicians benefit from correct explainable AI advice when reviewing X-rays. *Scientific Reports*, *13*(1), Article 1383. <https://doi.org/10.1038/s41598-023-28633-w>
- George, G., & Bock, A. J. (2011). The business model in practice and its implications for entrepreneurship research. *Entrepreneurship Theory and Practice*, *35*(1), 83–111. <https://doi.org/10.1111/j.1540-6520.2010.00424.x>
- Gerbert, P., Hartmann, P., Liebl, A., Schamberger, M., & Waldmann, A. (2020). *Building the organization for scaling AI*. appliedAI Initiative. <https://www.appliedai.de/hub/building-the-organization-for-scaling-ai>
- Ghoddusi, H., Creamer, G. G., & Rafizadeh, N. (2019). Machine learning in energy economics and finance: A review. *Energy Economics*, *81*, 709–727. <https://doi.org/10.1016/j.eneco.2019.05.006>
- Gibson, C. B., & Birkinshaw, J. (2004). The antecedents, consequences, and mediating role of organizational ambidexterity. *Academy of Management Journal*, *47*(2), 209–226. <http://doi.org/10.5465/20159573>
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, *16*(1), 15–31. <http://doi.org/10.1177/1094428112452151>
- Glaser, B. G. (1978). *Theoretical sensitivity: Advances in the methodology of grounded theory*. Sociology Press.
- Glaser, B. G., & Strauss, A. L. (1967). *The discovery of grounded theory: Strategies for qualitative research*. Aldine Publishing Company. <https://doi.org/10.4324/9780203793206>
- Gläser, J., & Laudel, G. (2004). *Experteninterviews und qualitative Inhaltsanalyse*. VS Verlag.
- Gregor, S., Chandra Kruse, L., & Seidel, S. (2020). Research perspectives: The anatomy of a design principle. *Journal of the Association for Information Systems*, *21*(6), 1622–1652. <http://doi.org/10.17705/1jais.00649>
- Gregor, S., & Hevner, R. A. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, *37*(2), 337–355. <http://doi.org/10.25300/MISQ/2013/37.2.01>
- Grewal, D., Guha, A., Satornino, C. B., & Schweiger, E. B. (2021). Artificial intelligence: The light and the darkness. *Journal of Business Research*, *136*, 229–236. <https://doi.org/10.1016/j.jbusres.2021.07.043>
- Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2019). A survey of methods for explaining black box models. *ACM Computing Surveys*, *51*(5), 1–42. <https://doi.org/10.1145/3236009>
- Gunning, D., Vorm, E., Wang, J. Y., & Turek, M. (2021). DARPA's explainable AI (XAI) program: A retrospective. *Applied AI Letters*, *2*(4), 1–11. <https://doi.org/10.1002/ail2.61>
- Günther, W. A., Rezazade Mehrizi, M. H., Huysman, M., & Feldberg, F. (2017). Debating big data: A literature review on realizing value from big data. *The Journal of Strategic Information Systems*, *26*(3), 191–209. <https://doi.org/10.1016/j.jsis.2017.07.003>

- Hadjimichael, D., & Tsoukas, H. (2019). Toward a better understanding of tacit knowledge in organizations: Taking stock and moving forward. *Academy of Management Annals*, 13(2), 672–703. <http://doi.org/10.5465/annals.2017.0084>
- Haefner, N., Parida, V., Gassmann, O., & Wincent, J. (2023). Implementing and scaling artificial intelligence: A review, framework, and research agenda. *Technological Forecasting and Social Change*, 197, Article 122878. <https://doi.org/10.1016/j.techfore.2023.122878>
- Haffke, I., Kalgovas, B., & Benlian, A. (2017a). Options for transforming the IT function using bimodal IT. *MIS Quarterly Executive*, 16(2), 101–120. <https://aisel.aisnet.org/misqe/vol16/iss2/2>
- Haffke, I., Kalgovas, B., & Benlian, A. (2017b). The transformative role of bimodal IT in an era of digital business. *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*, 5460–5469. <http://doi.org/10.24251/HICSS.2017.660>
- Hahn, C., Traunecker, T., Niever, M., & Basedow, G. N. (2020). Exploring AI-driven business models: Conceptualization and expectations in the machinery industry. *Proceedings of the 2020 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, 567–570. <https://doi.org/10.1109/IEEM45057.2020.9309824>
- Hanelt, A., Hildebrandt, B., & Polier, J. (2015). Uncovering the role of IS in business model innovation - a taxonomy-driven approach to structure the field. *Proceedings of the 23rd European Conference on Information Systems (ECIS)*, Paper 71. <https://doi.org/10.18151/7217345>
- Hartmann, P. M., Zaki, M., Feldmann, N., & Neely, A. (2016). Capturing value from big data – a taxonomy of data-driven business models used by start-up firms. *International Journal of Operations & Production Management*, 36(10), 1382–1406. <https://doi.org/10.1108/IJOPM-02-2014-0098>
- Hedman, J., & Kalling, T. (2003). The business model concept: Theoretical underpinnings and empirical illustrations. *European Journal of Information Systems*, 12(1), 49–59. <https://doi.org/10.1057/palgrave.ejis.3000446>
- Helfat, C. E., Finkelstein, S., Mitchell, W., Peteraf, M. A., Singh, H., & Winter, S. G. (2007). *Dynamic capabilities: Understanding strategic change in organizations*. Blackwell Publishing.
- Helfat, C. E., & Peteraf, M. A. (2003). The dynamic resource-based view: Capability lifecycles. *Strategic Management Journal*, 24(10), 997–1010. <https://doi.org/10.1002/smj.332>
- Helfat, C. E., & Peteraf, M. A. (2015). Managerial cognitive capabilities and the microfoundations of dynamic capabilities. *Strategic Management Journal*, 36(6), 831–850. <https://doi.org/10.1002/smj.2247>
- Hewamalage, H., Bergmeir, C., & Bandara, K. (2021). Recurrent neural networks for time series forecasting: Current status and future directions. *International Journal of Forecasting*, 37(1), 388–427. <https://doi.org/10.1016/j.ijforecast.2020.06.008>
- Hsieh, H.-F., & Shannon, S. E. (2005). Three approaches to qualitative content analysis. *Qualitative Health Research*, 15(9), 1277–1288. <https://doi.org/10.1177/1049732305276687>
- Hu, P. J.-H., Ma, P.-C., & Chau, P. Y. (1999). Evaluation of user interface designs for information retrieval systems: A computer-based experiment. *Decision Support Systems*, 27(1), 125–143. [https://doi.org/10.1016/S0167-9236\(99\)00040-8](https://doi.org/10.1016/S0167-9236(99)00040-8)

- Huang, J., Henfridsson, O., & Liu, M. (2017). Growing on steroids: Rapidly scaling the user base of digital ventures through digital innovation. *MIS Quarterly*, 41(1), 301–314. <https://doi.org/10.25300/MISQ/2017/41.1.16>
- Huber, G. P. (1991). Organizational learning: The contributing processes and the literatures. *Organization Science*, 2(1), 88–115. <http://doi.org/10.1287/orsc.2.1.88>
- Hunke, F., Engel, C. T., Schüritz, R., & Ebel, P. (2019). Understanding the anatomy of analytics-based services - a taxonomy to conceptualize the use of data and analytics in services. *Proceedings of the 27th European Conference on Information Systems*, Paper 25. https://aisel.aisnet.org/ecis2019_rp/25
- Hunke, F., Seebacher, S., Schüritz, R., & Illi, A. (2017). Towards a process model for data-driven business innovation. *Proceedings of the 19th IEEE Conference on Business Informatics (CBI)*, 150–157. <http://doi.org/10.1109/CBI.2017.43>
- Iankova, S., Davies, I., Archer-Brown, C., Marder, B., & Yau, A. (2019). A comparison of social media marketing between B2B, B2C and mixed business models. *Industrial Marketing Management*, 81, 169–179. <https://doi.org/10.1016/j.indmarman.2018.01.001>
- IDC. (2021). *Data creation and replication will grow at a faster rate than installed storage capacity, according to the IDC global DataSphere and StorageSphere forecasts*. Retrieved November 2, 2021, from <https://www.idc.com/getdoc.jsp?containerId=prUS47560321>
- Jain, A. K. (2010). Data clustering: 50 years beyond K-means. *Pattern Recognition Letters*, 31(8), 651–666. <https://doi.org/10.1016/j.patrec.2009.09.011>
- Jöhnk, J., Weißert, M., & Wyrтки, K. (2020). Ready or not, AI comes— an interview study of organizational AI readiness factors. *Business & Information Systems Engineering*, 63(1), 5–20. <https://doi.org/10.1007/s12599-020-00676-7>
- Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Kaggle. (2021, November 16). *Your machine learning and data science community*. <https://www.kaggle.com>
- Kane, G. C., & Alavi, M. (2007). Information technology and organizational learning: An investigation of exploration and exploitation processes. *Organization Science*, 18(5), 796–812. <https://doi.org/10.1287/orsc.1070.0286>
- Karimi, J., & Walter, Z. (2015). The role of dynamic capabilities in responding to digital disruption: A factor-based study of the newspaper industry. *Journal of Management Information Systems*, 32(1), 39–81. <https://doi.org/10.1080/07421222.2015.1029380>
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410. <https://doi.org/10.5465/annals.2018.0174>
- Kenessey, Z. (1987). The primary, secondary, tertiary and quaternary sectors of the economy. *Review of Income and Wealth*, 33(4), 359–385. <https://doi.org/10.1111/j.1475-4991.1987.tb00680.x>
- Kowalczyk, M., & Buxmann, P. (2015). An ambidextrous perspective on business intelligence and analytics support in decision processes: Insights from a multiple case study. *Decision Support Systems*, 80, 1–13. <https://doi.org/10.1016/j.dss.2015.08.010>

- Laato, S., Tiainen, M., Najmul Islam, A.K.M., & Mäntymäki, M. (2022). How to explain AI systems to end users: A systematic literature review and research agenda. *Internet Research*, 32(7), 1–31. <https://doi.org/10.1108/INTR-08-2021-0600>
- Lange, H. E., Drews, P., & Höft, M. (2021). Realization of data-driven business models in incumbent companies : An exploratory study based on the resource-based view. *Proceedings of the 42nd International Conference on Information Systems (ICIS)*. https://aisel.aisnet.org/icis2021/dig_innov/dig_innov/2
- Langer, M., Oster, D., Speith, T., Hermanns, H., Kästner, L., Schmidt, E., Sesing, A., & Baum, K. (2021). What do we want from explainable artificial intelligence (XAI)? -- A stakeholder perspective on XAI and a conceptual model guiding interdisciplinary XAI research. *Artificial Intelligence*, 296, Article 103473. <https://doi.org/10.1016/j.artint.2021.103473>
- Lanzolla, G., & Markides, C. (2021). A business model view of strategy. *Journal of Management Studies*, 58(2), 540–553. <https://doi.org/10.1111/joms.12580>
- Lebovitz, S. (2019). Diagnostic doubt and artificial intelligence: An inductive field study of radiology work. *Proceedings of the 40th International Conference on Information Systems (ICIS)*. https://aisel.aisnet.org/icis2019/future_of_work/future_work/11
- Lebovitz, S., Levina, N., & Lifshitz-Assaf, H. (2021). Is AI ground truth really true? The dangers of training and evaluating AI tools based on experts' know-what. *MIS Quarterly*, 45(3b), 1501–1525. <http://doi.org/10.25300/MISQ/2021/16564>
- Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science*, 33(1), 126–148. <http://doi.org/10.1287/orsc.2021.1549>
- Lee, G., DeLone, W. H., & Espinosa, J. A. (2010). The main and interaction effects of process rigor, process standardization, and process agility on system performance in distributed IS development: An ambidexterity perspective. *Proceedings of the 31st International Conference on Information Systems (ICIS)*, Paper 34. https://aisel.aisnet.org/icis2010_submissions/34
- Lee, J., Suh, T., Roy, D., & Baucus, M. (2019). Emerging technology and business model innovation: The case of artificial intelligence. *Journal of Open Innovation: Technology, Market, and Complexity*, 5(3), Article 44. <https://doi.org/10.3390/joitmc5030044>
- Leemann, N., & Kanbach, D. K. (2022). Toward a taxonomy of dynamic capabilities – a systematic literature review. *Management Research Review*, 45(4), 486–501. <http://doi.org/10.1108/MRR-01-2021-0066>
- Leppänen, P., George, G., & Alexy, O. (2023). When do novel business models lead to high performance? A configurational approach to value drivers, competitive strategy, and firm environment. *Academy of Management Journal*, 66(1), 164–194. <https://doi.org/10.5465/amj.2020.0969>
- Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal*, 14(S2), 95–112. <https://doi.org/10.1002/smj.4250141009>
- Levitt, B., & March, J. G. (1988). Organizational learning. *Annual Review of Sociology*, 14, 319–338. <https://doi.org/10.1146/annurev.so.14.080188.001535>

- Liao, Q. V., Gruen, D., & Miller, S. (2020). Questioning the AI: Informing design practices for explainable AI user experiences. *Proceedings of the 2020 ACM Conference on Human Factors in Computing Systems (CHI)*, Paper 463. <https://doi.org/10.1145/3313831.3376590>
- Likert, R (1932). A technique for the measurement of attitudes. *Archives of Psychology*, 140, 5–55.
- Lindebaum, D., Vesa, M., & den Hond, F. (2020). Insights from “the machine stops” to better understand rational assumptions in algorithmic decision making and its implications for organizations. *Academy of Management Review*, 45(1), 247–263. <http://doi.org/10.5465/amr.2018.0181>
- Lyytinen, K., Nickerson, J. V., & King, J. L. (2021). Metahuman systems = humans + machines that learn. *Journal of Information Technology*, 36(4), 427–445. <https://doi.org/10.1177/0268396220915917>
- Ma, S., & Fildes, R. (2021). Retail sales forecasting with meta-learning. *European Journal of Operational Research*, 288(1), 111–128. <https://doi.org/10.1016/j.ejor.2020.05.038>
- Makridakis, S. (2017). The forthcoming artificial intelligence (AI) revolution: Its impact on society and firms. *Futures*, 90, 46–60. <http://doi.org/10.1016/j.futures.2017.03.006>
- Makridakis, S., Spiliotis, E., Assimakopoulos, V., Semenoglou, A.-A., Mulder, G., & Nikolopoulos, K. (2023). Statistical, machine learning and deep learning forecasting methods: Comparisons and ways forward. *Journal of the Operational Research Society*, 74(3), 840–859. <http://doi.org/10.1080/01605682.2022.2118629>
- Mangalindan, J. P. (2010, July 9). *Twitter’s business model: A visionary experiment*. Fortune. https://money.cnn.com/2010/07/09/magazines/fortune/Twitter_business_model.fortune/
- Marabelli, M., Newell, S., & Handunge, V. (2021). The lifecycle of algorithmic decision-making systems: Organizational choices and ethical challenges. *The Journal of Strategic Information Systems*, 30(3), 1–15. <https://doi.org/10.1016/j.jsis.2021.101683>
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87. <https://doi.org/10.1287/orsc.2.1.71>
- March, J. G. (2006). Rationality, foolishness, and adaptive intelligence. *Strategic Management Journal*, 27(3), 201–214. <https://doi.org/10.1002/smj.515>
- March, J. G. (2010). *The Ambiguities of experience*. Cornell University Press.
- Martin, K. (2019a). Designing ethical algorithms. *MIS Quarterly Executive*, 18(2), 129–142. <http://doi.org/10.17705/2msqe.00012>
- Martin, K. (2019b). Ethical implications and accountability of algorithms. *Journal of Business Ethics*, 160(4), 835–850. <https://doi.org/10.1007/s10551-018-3921-3>
- Martinez, I., Viles, E., & G. Olazola, I. (2021). Data science methodologies: Current challenges and future approaches. *Big Data Research*, 24. <http://doi.org/10.1016/j.bdr.2020.100183>
- Martínez-Plumed, F., Contreras-Ochando, L., Ferri, C., Hernandez-Orallo, J., Kull, M., Lachiche, N., Ramirez-Quintana, M. J., & Flach, P. (2021). CRISP-DM twenty years later: From data mining processes to data science trajectories. *IEEE Transactions on Knowledge and Data Engineering*, 33(8), 3048–3061. <https://doi.org/10.1109/TKDE.2019.2962680>

- Massa, L., & Tucci, C. L. (2014). Business model innovation. In M. Dodgson, D. M. Gann, & N. Phillips (Eds.), *The Oxford handbook of innovation management*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199694945.013.002>
- Massa, L., Tucci, C. L., & Afuah, A. (2017). A critical assessment of business model research. *Academy of Management Annals*, 11(1), 73–104. <https://doi.org/10.5465/annals.2014.0072>
- Mayring, P. (2007). On generalization in qualitatively oriented research. *Forum: Qualitative Social Research*, 8(3), 1–11. <https://doi.org/10.17169/fqs-8.3.291>
- McCarthy, J. (2007, November 12). *What is artificial intelligence?* Stanford University. <http://www-formal.stanford.edu/jmc/whatisai.pdf>
- McCarthy, J. L., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955, August 31). *A proposal for the dartmouth summer research project on artificial intelligence*. <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>
- McLoughlin, S., Puvvala, A., Maccani, G., & Donnellan, B. (2019). A framework for understanding & classifying urban data business models. *Proceedings of the 52nd Hawaii International Conference on System Sciences (HICSS)*, 3295–3304. <http://doi.org/10.24251/HICSS.2019.398>
- Meske, C., Bunde, E., Schneider, J., & Gersch, M. (2022). Explainable artificial intelligence: Objectives, stakeholders, and future research opportunities. *Information Systems Management*, 39(1), 53–63. <https://doi.org/10.1080/10580530.2020.1849465>
- van der Meulen, R., & McCall, T. (2018). *Gartner says nearly half of CIOs are planning to deploy artificial intelligence*. Gartner. <https://www.gartner.com/en/newsroom/press-releases/2018-02-13-gartner-says-nearly-half-of-cios-are-planning-to-deploy-artificial-intelligence>
- Mikalef, P., Krogstie, J., Pappas, I. O., & Pavlou, P. (2020). Exploring the relationship between big data analytics capability and competitive performance: The mediating roles of dynamic and operational capabilities. *Information & Management*, 57(2), Article 103169. <http://doi.org/10.1016/j.im.2019.05.004>
- Miranda, S. (2004). Beyond BI: Benefiting from corporate performance management solutions. *Financial Executive*, 2(20), 58–61.
- MIT Technology Review Insights. (2018, November 2). *Professional services firms see huge potential in machine learning*. <https://www.technologyreview.com/2018/11/02/139216/professional-services-firms-see-huge-potential-in-machine-learning>
- Mitchell, T. M. (1997). *Machine learning*. McGraw-Hill. [https://doi.org/10.1002/\(SICI\)1099-1689\(199909\)9:3%3C191::AID-STVR184%3E3.0.CO;2-E](https://doi.org/10.1002/(SICI)1099-1689(199909)9:3%3C191::AID-STVR184%3E3.0.CO;2-E)
- Möhlmann, M., Zalmanson, L., Henfridsson, O., & Gregory, R. W. (2021). Algorithmic management of work on online labor platforms: When matching meets control. *MIS Quarterly*, 45(4), 1999–2022. <https://doi.org/10.25300/MISQ/2021/15333>
- Mojena, R. (1977). Hierarchical grouping methods and stopping rules: An evaluation. *The Computer Journal*, 20(4), 359–363. <https://doi.org/10.1093/comjnl/20.4.359>
- Möller, F., Bauhaus, H., Hoffmann, C., Niess, C., & Otto, B. (2019). Archetypes of digital business models in logistics start-ups. *Proceedings of the 27th European Conference on Information Systems (ECIS)*, Paper 17. https://aisel.aisnet.org/ecis2019_rp/17

- Möller, F., Stachon, M., Hoffmann, C., Bauhaus, H., & Otto, B. (2020). Data-driven business models in logistics: A taxonomy of optimization and visibility services. *Proceedings of the 53rd Hawaii International Conference on System Sciences (HICSS)*, 5379–5388. <http://doi.org/10.24251/HICSS.2020.661>
- Morris, M., Schindehutte, M., & Allen, J. (2005). The entrepreneur's business model: Toward a unified perspective. *Journal of Business Research*, 58(6), 726–735. <https://doi.org/10.1016/j.jbusres.2003.11.001>
- Mousavi, S., Bossink, B., & van Vliet, M. (2019). Microfoundations of companies' dynamic capabilities for environmentally sustainable innovation: Case study insights from high-tech innovation in science-based companies. *Business Strategy and the Environment*, 28(2), 366–387. <https://doi.org/10.1002/bse.2255>
- Müller, J., & Buliga, O. (2019). Archetypes for data-driven business models for manufacturing companies in industry 4.0. *Proceedings of the 40th International Conference on Information Systems (ICIS)*. <https://aisel.aisnet.org/sigbd2019/2>
- Müller, O., Junglas, I., vom Brocke, J., & Debortoli, S. (2016). Utilizing big data analytics for information systems research: Challenges, promises and guidelines. *European Journal of Information Systems*, 25(4), 289–302. <http://doi.org/10.1057/ejis.2016.2>
- Murdoch, W. J., Singh, C., Kumbier, K., Abbasi-Asl, R., & Yu, B. (2019). Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences of the United States of America*, 116(44), 22071–22080. <https://doi.org/10.1073/pnas.1900654116>
- Murray, A., Rhymer, J., & Sirmon, D. G. (2021). Humans and technology: Forms of conjoined agency in organizations. *Academy of Management Review*, 46(3), 552–571. <https://doi.org/10.5465/amr.2019.0186>
- Myers, M. (1997). Qualitative research in information systems. *MIS Quarterly*, 21(2), 241–242. <http://doi.org/10.2307/249422>
- Myers, M. D., & Newman, M. (2007). The qualitative interview in IS research: Examining the craft. *Information and Organization*, 17(1), 2–26. <https://doi.org/10.1016/j.infoandorg.2006.11.001>
- Nambisan, S., Lyytinen, K., Majchrzak, A., & Song, M. (2017). Digital innovation management: Reinventing innovation management research in a digital world. *MIS Quarterly*, 41(1), 223–238. <http://doi.org/10.25300/MISQ/2017/41:1.03>
- Naous, D., Schwarz, J., & Legner, C. (2017). Analytics as a service: Cloud computing and the transformation of business analytics business models and ecosystems. *Proceedings of the 25th European Conference on Information Systems (ECIS)*, Paper 32. https://aisel.aisnet.org/ecis2017_rp/32
- Nelson, L. K. (2020). Computational grounded theory: A methodological framework. *Sociological Methods and Research*, 49(1), 3–42. <https://doi.org/10.1177/0049124117729703>
- Nickerson, R. C., Varshney, U., & Muntermann, J. (2013). A method for taxonomy development and its application in information systems. *European Journal of Information Systems*, 22(3), 336–359. <https://doi.org/10.1057/ejis.2012.26>

- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37. <https://doi.org/10.1287/orsc.5.1.14>
- Obschonka, M., & Audretsch, D. B. (2020). Artificial intelligence and big data in entrepreneurship: A new era has begun. *Small Business Economics*, 55(3), 529–539. <https://doi.org/10.1007/s11187-019-00202-4>
- Oh, C., Song, J., Choi, J., Kim, S., Lee, S., & Suh, B. (2018). I lead, you help but only with enough details: Understanding user experience of co-creation with artificial intelligence. *Proceedings of the 2018 ACM Conference on Human Factors in Computing Systems (CHI)*, Paper 649. <https://doi.org/10.1145/3173574.3174223>
- Omerovic, M., Islam, N., & Buxmann, P. (2020). Unlashing the next wave of business models in the internet of things era: New directions for a research agenda based on a systematic literature review. *Proceedings of the 53rd Hawaii International Conference on System Sciences (HICSS)*, 4569–4578. <http://doi.org/10.24251/HICSS.2020.559>
- OpenAI. (2022, November 30). *Introducing ChatGPT*. <https://openai.com/index/chatgpt>
- O'Reilly, C. A., & Tushman, M. L. (2008). Ambidexterity as a dynamic capability: Resolving the innovator's dilemma. *Research in Organizational Behavior*, 28, 185–206. <http://doi.org/10.1016/j.riob.2008.06.002>
- O'Reilly, C. A., & Tushman, M. L. (2013). Organizational ambidexterity: Past, present, and future. *Academy of Management Perspectives*, 27(4), 324–338. <https://doi.org/10.5465/amp.2013.0025>
- Ossenbrink, J., Hoppmann, J., & Hoffmann, V. H. (2019). Hybrid ambidexterity: How the environment shapes incumbents' use of structural and contextual approaches. *Organization Science*, 30(6), 1319–1348. <http://doi.org/10.1287/orsc.2019.1286>
- Osterwalder, A., & Pigneur, Y. (2010). *Business model generation: A handbook for visionaries, game changers, and challengers*. John Wiley & Sons.
- Passlick, J., Dreyer, S., Olivotti, D., Grützner, L., Eilers, D., & Breitner, M. H. (2021). Predictive maintenance as an internet of things enabled business model: A taxonomy. *Electronic Markets*, 31, 67–87. <https://doi.org/10.1007/s12525-020-00440-5>
- Pavlyshenko, B. (2019). Machine-learning models for sales time series forecasting. *Data*, 4(1), 1–15. <http://doi.org/10.3390/data4010015>
- Peffer, K., Tuunanen, T., Gengler, C., Rossi, M., Hui, W., Virtanen, V., & Bragge, J. (2006). The design science research process: A model for producing and presenting information systems research. *Proceedings of the 1st International Conference on Design Science Research in Information Systems and Technology (DESRIST)*, 83–106.
- Peteraf, M., Di Stefano, G., & Verona, G. (2013). The elephant in the room of dynamic capabilities: Bringing two diverging conversations together. *Strategic Management Journal*, 34(12), 1389–1410. <https://doi.org/10.1002/smj.2078>
- Porter, M. E. (2001). Strategy and the internet. *Harvard Business Review*, 79(3), 62–78.
- Prescott, J. E., & Filatotchev, I. (2021). The business model phenomenon: Towards theoretical relevance. *Journal of Management Studies*, 58(2), 517–527. <https://doi.org/10.1111/joms.12610>

- Pumplun, L., Peters, F., Gawlitza, J. F., & Buxmann, P. (2023). Bringing machine learning systems into clinical practice: A design science approach to explainable machine learning-based clinical decision support systems. *Journal of the Association for Information Systems*, 24(4), 953–979. <https://doi.org/10.17705/1jais.00820>
- Pumplun, L., Tauchert, C., & Heidt, M. (2019). A new organizational chassis for artificial intelligence - exploring organizational readiness factors. *Proceedings of 27th European Conference on Information Systems (ECIS)*, Paper 106. https://aisel.aisnet.org/ecis2019_rp/106
- Punj, G., & Stewart, D. W. (1983). Cluster analysis in marketing research: Review and suggestions for application. *Journal of Marketing Research*, 20(2), 134–148. <https://doi.org/10.1177/002224378302000204>
- Putnam, L. L., Fairhurst, G. T., & Banghart, S. (2016). Contradictions, dialectics, and paradoxes in organizations: A constitutive approach. *Academy of Management Annals*, 10(1), 65–171. <https://doi.org/10.1080/19416520.2016.1162421>
- Rai, A. (2020). Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science*, 48(1), 137–141. <https://doi.org/10.1007/s11747-019-00710-5>
- Raisch, S., & Birkinshaw, J. (2008). Organizational ambidexterity: Antecedents, outcomes, and moderators. *Journal of Management*, 34(3), 375–409. <https://doi.org/10.1177/0149206308316058>
- Raisch, S., & Fomina, K. (2024). Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Academy of Management Review*. Advance online publication. <https://doi.org/10.5465/amr.2021.0421>
- Ransbotham, S., Khodabandeh, S., Fehling, R., LaFountain, B., & Kiron, D. (2019, October 15). Winning with AI. *MIT Sloan Management Review*. <https://sloanreview.mit.edu/projects/winning-with-ai>
- Ransbotham, S., Khodabandeh, S., Kiron, D., Candelon, F., Chu, M., & LaFountain, B. (2020, October 20). Expanding AI's impact with organizational learning. *MIT Sloan Management Review*. <https://sloanreview.mit.edu/projects/expanding-ais-impact-with-organizational-learning>
- Ransbotham, S., Kiron, D., Gerbert, P., & Reeves, M. (2017, September 6). Reshaping business with artificial intelligence. *MIT Sloan Management Review*. <https://sloanreview.mit.edu/projects/reshaping-business-with-artificial-intelligence>
- Rashed, F., & Drews, P. (2021). Pathways of data-driven business model design and realization: A qualitative research study. *Proceedings of the 54th Hawaii International Conference on System Sciences (HICSS)*, 5676–5685. <http://doi.org/10.24251/HICSS.2021.689>
- Remane, G., Nickerson, R. C., Hanelt, A., Tesch, J. F., & Kolbe, L. M. (2016). A taxonomy of carsharing business models. *Proceedings of the 37th International Conference on Information Systems (ICIS)*. <https://aisel.aisnet.org/icis2016/Crowdsourcing/Presentations/18>
- Rencher, A. C. (2002). *Methods of multivariate analysis* (2nd ed.). Wiley. <http://doi.org/10.1002/0471271357>
- Ricciardi, F., Zardini, A., & Rossignoli, C. (2016). Organizational dynamism and adaptive business model innovation: The triple paradox configuration. *Journal of Business Research*, 69(11), 5487–5493. <http://doi.org/10.1016/j.jbusres.2016.04.154>

- Richards, G., Yeoh, W., Chong, A. Y. L., & Popovič, A. (2019). Business intelligence effectiveness and corporate performance management: An empirical analysis. *Journal of Computer Information Systems*, *59*(2), 188–196. <http://doi.org/10.1080/08874417.2017.1334244>
- Rindova, V., Martins, L., & Yeow, A. (2016). The hare and the fast tortoise: Dynamic resource reconfiguration and the pursuit of new growth opportunities by Yahoo and Google (1995–2007). In T. Folta, C. E. Helfat, & S. Karim (Eds.), *Resource redeployment and corporate strategy (advances in strategic management)* (Vol. 35, pp. 253–284). <http://doi.org/10.1108/S0742-332220160000035009>
- Ritter, T., & Lettl, C. (2018). The wider implications of business-model research. *Long Range Planning*, *51*(1), 1–8. <https://doi.org/10.1016/j.lrp.2017.07.005>
- Rogers, P. R., Miller, A., & Judge, W. Q. (1999). Using information-processing theory to understand planning/performance relationships in the context of strategy. *Strategic Management Journal*, *20*(6), 567–577. [https://doi.org/10.1002/\(SICI\)1097-0266\(199906\)20:6%3C567::AID-SMJ36%3E3.0.CO;2-K](https://doi.org/10.1002/(SICI)1097-0266(199906)20:6%3C567::AID-SMJ36%3E3.0.CO;2-K)
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, *1*(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: A modern approach* (3rd ed.). Pearson.
- Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- Sabatier, V., Mangematin, V., & Rousselle, T. (2010). From recipe to dinner: Business model portfolios in the european biopharmaceutical industry. *Long Range Planning*, *43*(2-3), 431–447. <https://doi.org/10.1016/j.lrp.2010.02.001>
- Saldaña, J. (2009). *The coding manual for qualitative researchers* (1st ed.). SAGE Publications.
- Saldaña, J. (2015). *The coding manual for qualitative researchers* (3rd ed.). SAGE Publications.
- Salesforce. (2022, March 11). *Artificial intelligence technology and resources: Salesforce Einstein*. <https://www.salesforce.com/products/einstein/overview/?d=cta-body-promo-90>
- Salovaara, A., Lyytinen, K., & Penttinen, E. (2019). High reliability in digital organizing: Mindlessness, the frame problem, and digital operations. *MIS Quarterly*, *43*(2), 555–578. <http://doi.org/10.25300/MISQ/2019/14577>
- Sanctify Financial Technologies. (2021, November 16). *Taking financial forecasting into the future*. <https://sanctify.ai>
- Santos, F., Pereira, R., & Vasconcelos, J. B. (2020). Toward robotic process automation implementation: An end-to-end perspective. *Business Process Management Journal*, *26*(2), 405–420. <https://doi.org/10.1108/BPMJ-12-2018-0380>
- Sarker, S., Sarker, S., Sahaym, A., & Bjørn-Andersen, N. (2012). Exploring value cocreation in relationships between an ERP vendor and its partners: A revelatory case study. *MIS Quarterly*, *36*(1), 317–338. <http://doi.org/10.2307/41410419>
- Schad, J., Lewis, M. W., Raisch, S., & Smith, W. K. (2017). Paradox research in management science: Looking back to move forward. *Academy of Management Annals*, *10*(1), 5–64. <https://doi.org/10.1080/19416520.2016.1162422>

- Schaltegger, S., Lüdeke-Freund, F., & Hansen, E. G. (2012). Business cases for sustainability: The role of business model innovation for corporate sustainability. *International Journal of Innovation and Sustainable Development*, 6(2), 95–119. <https://doi.org/10.1504/IJISD.2012.046944>
- Schilling, M. A., & Phelps, C. C. (2007). Interfirm collaboration networks: The impact of large-scale network structure on firm innovation. *Management Science*, 53(7), 1113–1126. <https://doi.org/10.1287/mnsc.1060.0624>
- Schneider, S., & Spieth, P. (2013). Business model innovation: Towards an integrated future research agenda. *International Journal of Innovation Management*, 17(01), Article 1340001. <https://doi.org/10.1142/S136391961340001X>
- Schoemaker, P. J. H., Heaton, S., & Teece, D. (2018). Innovation, dynamic capabilities, and leadership. *California Management Review*, 61(1), 15–42. <https://doi.org/10.1177/0008125618790246>
- Schroeder, R. (2016). Big data business models: Challenges and opportunities. *Cogent Social Sciences*, 2(1), Article 1166924. <https://doi.org/10.1080/23311886.2016.1166924>
- Schuetz, S., & Venkatesh, V. (2020). The rise of human machines: How cognitive computing systems challenge assumptions of user-system interaction. *Journal of the Association for Information Systems*, 21(2), 460–482. <https://doi.org/10.17705/1jais.00608>
- Schüritz, R., & Satzger, G. (2016). Patterns of data-infused business model innovation. *Proceedings of the 18th IEEE Conference on Business Informatics (CBI)*, 133–142. <http://doi.org/10.1109/CBI.2016.23>
- Schüritz, R., Seebacher, S., & Dorner, R. (2017). Capturing value from data: Revenue models for data-driven services. *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*, 5348–5357. <http://doi.org/10.24251/HICSS.2017.648>
- Scissero. (2021, November 16). *Unlock data, save time and cut legal costs with Scissero*. <https://www.scissero.com>
- Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511. https://papers.nips.cc/paper_files/paper/2015/hash/86df7dcfd896fc2674f757a2463eba-Abstract.html
- Seelos, C., & Mair, J. (2007). Profitable business models and market creation in the context of deep poverty: A strategic view. *Academy of Management Perspectives*, 21(4), 49–63. <http://doi.org/10.5465/AMP.2007.27895339>
- Seidel, S., Berente, N., Lindberg, A., Lyytinen, K., Martinez, B., & Nickerson, J. V. (2020). Artificial intelligence and video game creation: A framework for the new logic of autonomous design. *Journal of Digital Social Research*, 2(3), 126–157. <http://doi.org/10.33621/jdsr.v2i3.46>
- Seidel, S., Berente, N., Lindberg, A., Lyytinen, K., & Nickerson, J. (2019). Autonomous tools and design: A triple-loop approach to human-machine learning. *Communications of the ACM*, 62(1), 50–57. <https://doi.org/10.1145/3210753>
- Seidel, S., & Urquhart, C. (2013). On emergence and forcing in information systems grounded theory studies: The case of Strauss and Corbin. *Journal of Information Technology*, 28, 237–260. <https://doi.org/10.1057/jit.2013.17>

- Shafer, S. M., Smith, H. J., & Linder, J. C. (2005). The power of business models. *Business Horizons*, 48(3), 199–207. <https://doi.org/10.1016/j.bushor.2004.10.014>
- Sharma, R., Mithas, S., & Kankanhalli, A. (2014). Transforming decision-making processes: A research agenda for understanding the impact of business analytics on organisations. *European Journal of Information Systems*, 23(4), 433–441. <https://doi.org/10.1057/ejis.2014.17>
- Shearer, C. (2000). The CRISP-DM model: The new blueprint for data mining. *Journal of Data Warehousing*, 5(4), 13–22.
- Shollo, A., Hopf, K., Thiess, T., & Müller, O. (2022). Shifting ML value creation mechanisms: A process model of ML value creation. *The Journal of Strategic Information Systems*, 31(3), Article 101734. <https://doi.org/10.1016/j.jsis.2022.101734>
- Simon, H. (1991). Bounded rationality and organizational learning. *Organization Science*, 2(1), 125–134. <https://doi.org/10.1287/orsc.2.1.125>
- Singh, A., & Hess, T. (2017). How chief digital officers promote the digital transformation of their companies. *MIS Quarterly Executive*, 16(1), 1–17. <https://aisel.aisnet.org/misqe/vol16/iss1/5>
- Sint, R., Schaffert, S., Stroka, S., & Ferstl, R. (2009). Combining unstructured, fully structured and semi-structured information in semantic wikis. *Proceedings of the 4th Semantic Wiki Workshop (SemWiki) at the 6th European Semantic Web Conference (ESWC)*, 73–87.
- Smith, W. K., & Lewis, M. W. (2011). Toward a theory of paradox: A dynamic equilibrium model of organizing. *Academy of Management Review*, 36(2), 381–403. <https://doi.org/10.5465/AMR.2011.59330958>
- Smith, W. K., Lewis, M. W., & Tushman, M. L. (2011). Organizational sustainability: Organization design and senior leadership to enable strategic paradox. In K. Cameron & G. Spreitzer (Eds.), *The Oxford handbook of positive organizational scholarship* (pp. 798–810). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199734610.013.0061>
- Snihur, Y., & Eisenhardt, K. M. (2022). Looking forward, looking back: Strategic organization and the business model concept. *Strategic Organization*, 20(4), 757–770. <https://doi.org/10.1177/14761270221122442>
- Snihur, Y., & Markman, G. (2023). Business model research: Past, present, and future. *Journal of Management Studies*, 60(8), e1–e14. <https://doi.org/10.1111/joms.12928>
- Spiegel, O., Abbassi, P., Zylka, M. P., Schlagwein, D., Fischbach, K., & Schoder, D. (2016). Business model development, founders' social capital and the success of early stage internet start-ups: A mixed-method study. *Information Systems Journal*, 26(5), 421–449. <https://doi.org/10.1111/isj.12073>
- Spieth, P., Schneckenberg, D., & Ricart, J. E. (2014). Business model innovation—state of the art and future challenges for the field. *R&D Management*, 44(3), 237–247. <https://doi.org/10.1111/radm.12071>
- Spiliotis, E., Nikolopoulos, K., & Assimakopoulos, V. (2019). Tales from tails: On the empirical distributions of forecasting errors and their implication to risk. *International Journal of Forecasting*, 35(2), 687–698. <http://doi.org/10.1016/j.ijforecast.2018.10.004>

- Steininger, D. M. (2019). Linking information systems and entrepreneurship: A review and agenda for IT-associated and digital entrepreneurship research. *Information Systems Journal*, 29(2), 363–407. <https://doi.org/10.1111/isj.12206>
- Steininger, D. M., Brohman, K. M., & Block, J. H. (2022). Digital entrepreneurship: What is new if anything? *Business & Information Systems Engineering*, 64(1), 1–14. <https://doi.org/10.1007/s12599-021-00741-9>
- 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. L., Shah, J., Tambe, M., & 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. <http://ai100.stanford.edu/2016-report>
- Studer, S., Bui, T. B., Drescher, C., Hanuschkin, A., Winkler, L., Peters, S., & Müller, K. R. (2021). Towards CRISP-ML(Q): A machine learning process model with quality assurance methodology. *Machine Learning and Knowledge Extraction*, 3(2), 392–413. <http://doi.org/10.3390/make3020020>
- Sturm, T., Koppe, T., Scholz, Y., & Buxmann, P. (2021a). The case of human-machine trading as bilateral organizational learning. *Proceedings of the 42nd International Conference on Information Systems (ICIS)*. https://aisel.aisnet.org/icis2021/ai_business/ai_business/3
- Sturm, T., Gerlach, J. P., Pumplun, L., Mesbah, N., Peters, F., Tauchert, C., Nan, N., & Buxmann, P. (2021b). Coordinating human and machine learning for effective organizational learning. *MIS Quarterly*, 45(3), 1581–1602. <https://doi.org/10.25300/MISQ/2021/16543>
- Sturm, T., Fecho, M., & Buxmann, P. (2021c). To use or not to use artificial intelligence? A framework for the ideation and evaluation of problems to be solved with artificial intelligence. *Proceedings of the 54th Hawaii International Conference on System Sciences (HICSS)*, 206–215. <http://doi.org/10.24251/HICSS.2021.023>
- Tallon, P. P., Ramirez, R. v., & Short, J. E. (2013). The information artifact in IT governance: Toward a theory of information governance. *Journal of Management Information Systems*, 30(3), 141–177. <http://doi.org/10.2753/MIS0742-1222300306>
- Täuscher, K., & Laudien, S. M. (2018). Understanding platform business models: A mixed methods study of marketplaces. *European Management Journal*, 36(3), 319–329. <https://doi.org/10.1016/j.emj.2017.06.005>
- Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319–1350. <https://doi.org/10.1002/smj.640>
- Teece, D. J. (2009). *Dynamic capabilities and strategic management: Organizing for innovation and growth*. Oxford University Press.
- Teece, D. J. (2010). Business models, business strategy and innovation. *Long Range Planning*, 43(2–3), 172–194. <https://doi.org/10.1016/j.lrp.2009.07.003>
- Teece, D. J. (2014). The foundations of enterprise performance: Dynamic and ordinary capabilities in an (economic) theory of firms. *Academy of Management Perspectives*, 28(4), 328–352. <http://doi.org/10.5465/amp.2013.0116>

- Teece, D. J. (2018). Business models and dynamic capabilities. *Long Range Planning*, 51(1), 40–49. <http://doi.org/10.1016/j.lrp.2017.06.007>
- Teece, D. J., & Leih, S. (2016). Uncertainty, innovation, and dynamic capabilities: An introduction. *California Management Review*, 58(4), 5–12. <http://doi.org/10.1525/cmr.2016.58.4.5>
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533. [https://doi.org/10.1002/\(SICI\)1097-0266\(199708\)18:7%3C509::AID-SMJ882%3E3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7%3C509::AID-SMJ882%3E3.0.CO;2-Z)
- Tidhar, R., & Eisenhardt, K. M. (2020). Get rich or die trying... finding revenue model fit using machine learning and multiple cases. *Strategic Management Journal*, 41(7), 1245–1273. <https://doi.org/10.1002/smj.3142>
- Townsend, D. M., & Hunt, R. A. (2019). Entrepreneurial action, creativity, & judgment in the age of artificial intelligence. *Journal of Business Venturing Insights*, 11, Article e00126. <http://doi.org/10.1016/j.jbvi.2019.e00126>
- Uber. (2024, February). *Investor update February 2024*. <https://investor.uber.com/financials>
- Uotila, J., Maula, M., Keil, T., & Zhara, S. A. (2008). Exploration, exploitation and firm performance: An analysis of S&P 500 corporations. *Strategic Management Journal*, 30, 221–231. <https://doi.org/10.1002/smj.738>
- Urquhart, C., Lehmann, H., & Myers, M. D. (2010). Putting the ‘theory’ back into grounded theory: Guidelines for grounded theory studies in information systems. *Information systems journal*, 20(4), 357–381. <https://doi.org/10.1111/j.1365-2575.2009.00328.x>
- Urquhart, C., & Fernández, W. (2013). Using grounded theory method in information systems: The researcher as blank slate and other myths. *Journal of Information Technology*, 28(3), 224–236. <https://doi.org/10.1057/jit.2012.34>
- Vaast, E., & Pinsonneault, A. (2021). When digital technologies enable and threaten occupational identity: The delicate balancing act of data scientists. *MIS Quarterly*, 45(3), 1087–1112. <http://doi.org/10.25300/MISQ/2021/16024>
- Veit, D., Clemons, E., Benlian, A., Buxmann, P., Hess, T., Kundisch, D., Leimeister, J. M., Loos, P., & Spann, M. (2014). Business models: An information systems research agenda. *Business & Information Systems Engineering*, 6(1), 45–53. <https://doi.org/10.1007/s12599-013-0308-y>
- Vetter, O. A., Hoffmann, F. S., Pumplun, L., & Buxmann, P. (2022). What constitutes a machine-learning-driven business model? A taxonomy of B2B start-ups with machine learning at their core. *Proceedings of the 30th European Conference on Information Systems (ECIS)*, Paper 29. https://aisel.aisnet.org/ecis2022_rp/29
- Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *The Journal of Strategic Information Systems*, 28(2), 118–144. <http://doi.org/10.1016/j.jsis.2019.01.003>
- Wamba, S. F., Gunasekaran, A., Akter, S., Ren, S. J., Dubey, R., & Childe, S. J. (2017). Big data analytics and firm performance: Effects of dynamic capabilities. *Journal of Business Research*, 70, 356–365. <https://doi.org/10.1016/j.jbusres.2016.08.009>

- Wamba-Taguimdje, S.-L., Fosso Wamba, S., Kala Kamdjoug, J. R., & Tchatchouang Wanko, C. E. (2020). Influence of artificial intelligence (AI) on firm performance: The business value of AI-based transformation projects. *Business Process Management Journal*, 26(7), 1893–1924. <https://doi.org/10.1108/BPMJ-10-2019-0411>
- Ward, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301), 236–244. <https://doi.org/10.1080/01621459.1963.10500845>
- Warner, K. S. R., & Wäger, M. (2019). Building dynamic capabilities for digital transformation: An ongoing process of strategic renewal. *Long Range Planning*, 52(3), 326–349. <https://doi.org/10.1016/j.lrp.2018.12.001>
- Wasserbacher, H., & Spindler, M. (2022). Machine learning for financial forecasting, planning and analysis: Recent developments and pitfalls. *Digital Finance*, 4(1), 63–88. <https://doi.org/10.1007/s42521-021-00046-2>
- Weber, M., Beutter, M., Weking, J., Böhm, M., & Krcmar, H. (2022). AI startup business models: Key characteristics and directions for entrepreneurship research. *Business & Information Systems Engineering*, 64(1), 91–109. <https://doi.org/10.1007/s12599-021-00732-w>
- Weking, J., Stöcker, M., Kowalkiewicz, M., Böhm, M., & Krcmar, H. (2020). Leveraging industry 4.0 – a business model pattern framework. *International Journal of Production Economics*, 225, Article 107588. <https://doi.org/10.1016/j.ijpe.2019.107588>
- Wiener, M., Cram, W. A., & Benlian, A. (2023). Algorithmic control and gig workers: A legitimacy perspective of Uber drivers. *European Journal of Information Systems*, 32(3), 485–507. <https://doi.org/10.1080/0960085X.2021.1977729>
- Wiener, M., Saunders, C., & Marabelli, M. (2020). Big-data business models: A critical literature review and multiperspective research framework. *Journal of Information Technology*, 35(1), 66–91. <http://doi.org/10.1177/0268396219896811>
- Wiesche, M., Jurisch, M. C., Yetton, P. W., & Krcmar, H. (2017). Grounded theory methodology in information systems research. *MIS Quarterly*, 41(3), 685–701. <http://doi.org/10.25300/MISQ/2017/41.3.02>
- Wilden, R., & Gudergan, S. P. (2015). The impact of dynamic capabilities on operational marketing and technological capabilities: Investigating the role of environmental turbulence. *Journal of the Academy of Marketing Science*, 43(2), 181–199. <https://doi.org/10.1007/s11747-014-0380-y>
- Winter, S. G. (2003). Understanding dynamic capabilities. *Strategic Management Journal*, 24(10), 991–995. <https://doi.org/10.1002/smj.318>
- Wirth, R., & Hipp, J. (2000). CRISP-DM: Towards a standard process model for data mining. *Proceedings of the 4th International Conference on the Practical Applications of Knowledge Discovery and Data Mining*, 1, 29–39.
- Wirtz, B. W., Pistoia, A., Ullrich, S., & Göttel, V. (2016). Business models: Origin, development and future research perspectives. *Long Range Planning*, 49(1), 36–54. <http://doi.org/10.1016/j.lrp.2015.04.001>

- Woroch, R., & Strobel, G. (2021). Understanding value creation in digital companies – a taxonomy of IoT enabled business models. *Proceedings of the 29th European Conference on Information Systems (ECIS)*, Paper 11. https://aisel.aisnet.org/ecis2021_rp/11
- Yang, H.-F., & Chen, Y.-P. P. (2019). Representation learning with extreme learning machines and empirical mode decomposition for wind speed forecasting methods. *Artificial Intelligence*, 277, Article 103176. <https://doi.org/10.1016/j.artint.2019.103176>
- Yeow, A., Soh, C., & Hansen, R. (2018). Aligning with new digital strategy: A dynamic capabilities approach. *The Journal of Strategic Information Systems*, 27(1), 43–58. <http://doi.org/10.1016/j.jsis.2017.09.001>
- Yoo, Y. (2010). Computing in everyday life: A call for research on experiential computing. *MIS Quarterly*, 34(2), 213–231. <https://doi.org/10.2307/20721425>
- Zebhauser, J. (2024). Resource configuration for scaling AI startups — an action design research approach. *Proceedings of the 32nd European Conference on Information Systems (ECIS)*. https://aisel.aisnet.org/ecis2024/track23_designresearch/track23_designresearch/4
- Zhou, J., Gandomi, A. H., Chen, F., & Holzinger, A. (2021). Evaluating the quality of machine learning explanations: A survey on methods and metrics. *Electronics*, 10(5), Article 593. <http://doi.org/10.3390/electronics10050593>
- Zott, C., & Amit, R. (2008). The fit between product market strategy and business model: Implications for firm performance. *Strategic Management Journal*, 29, 1–26. <https://doi.org/10.1002/smj.642>
- Zott, C., & Amit, R. (2010). Business model design: An activity system perspective. *Long Range Planning*, 43(2–3), 216–226. <https://doi.org/10.1016/j.lrp.2009.07.004>
- Zott, C., Amit, R., & Massa, L. (2011). The business model: Recent developments and future research. *Journal of Management*, 37(4), 1019–1042. <https://doi.org/10.1177/0149206311406265>

Appendix

Appendix 1. Excerpt of the Examined Sample Set of B2B ML-Driven Start-Ups (Paper D).

Start-up	Website	Start-up	Website
Acquired Insights	www.aiinc.cloud	Alectio	www.alectio.com
Animatico	animati.co	Arva Intelligence	www.arvaintelligence.com
Arytic	arytic.com	Bevov	www.bevov.com
Blyng	blyng.io	Clinicgram	www.clinicgram.com
Cordian	www.cordian.com	CropSafe	www.cropsafe.io
DeepHow	www.deephow.com	DeepRisk.ai	deeperisk.ai
Donna	donna.legal	Edgematrix	edgematrix.com
edison.ai	edison.ai	Eiffo Analytics	www.eiffo-analytics.com
FACTIC	www.factic-sf.com	Frontier Medicines	frontiermeds.com
Gamyte	www.gamyte.com	Hasty	hasty.ai
Inlet Laboratories	inletlabs.com	Intuition	www.intuition.com
LabVoice	www.labvoice.ai	Logmind	logmind.com
Mapxus	www.mapxus.com	Myst AI	www.myst.ai
Nanochomp	www.nanochomp.com	Nucleus Cyber	nucleuscyber.com
Orbem	orbem.ai	Pixofarm	pixofarm.com
RailVision Analytics	www.railvision.ca	REIGO Investments	www.reigo-inv.com
Salesken	salesken.ai	Swiftlane	swiftlane.com
ThreatLandscape	threatlandscape.com	Traverse	www.traverse.ai
uman.ai	uman.ai	Uservision	www.user.vision
Virtual Facility	www.virtualfacility.ai	Xelera Technologies	xelera.io