
Developing a Pathway for the Adoption of Machine Learning Systems in Organizations: An Analysis of Drivers, Barriers, and Impacts with a Focus on the Healthcare Sector



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Department of Law and Economics
at the Technical University of Darmstadt

Dissertation

by

Luisa Pumplun
born in Dinslaken

submitted in partial 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 2022

Luisa Pumplun

Developing a Pathway for the Adoption of Machine Learning Systems in Organizations:
An Analysis of Drivers, Barriers, and Impacts with a Focus on the Healthcare Sector

Darmstadt, Technical University of Darmstadt

Dissertation published on TUpriints in 2022

Day of the viva voce: 14 July 2022

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 for 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.

Luisa Pumplun

Darmstadt, 17 February 2022

Abstract

The potential of machine learning (ML) and systems based thereon has grown steadily in recent years. The ability of ML systems to rapidly and systematically identify relationships in large volumes of data, which can be used to analyze new data to make meaningful predictions, enables organizations of all industries to make their processes more effective and efficient. Healthcare in particular may benefit greatly from ML systems in the future, as these systems' capabilities could help to ensure adequate patient care despite many pressing issues, such as the acute shortage of specialists (e.g., through diagnostic support).

However, many organizations are currently still failing to harness the potential of ML systems to their advantage, as implementing these systems is not a trivial task. Rather, the integration of ML systems requires the organization to identify and meet novel, multi-faceted preconditions that are unfamiliar as compared with previous, conventional technologies. This is mainly because ML systems exhibit unique characteristics. In particular, ML systems possess probabilistic properties due to their data-based learning approach, implying that their application can lead to erroneous results and that their functioning is often opaque. Particularly in healthcare, in which patients' lives depend on proper diagnoses and treatment, these characteristics result in ML systems not only being helpful, but – if introduced improperly – can also lead to severe detrimental consequences. Since previous research on the adoption of conventional technologies has not yet considered the characteristic properties of ML systems, the aim of this dissertation is to better understand the complex requirements for the successful adoption of ML systems in organizations in order for them to sustainably realize ML systems' potential. The three qualitative, two experimental, and one simulation study included in this cumulative dissertation have been published in peer-reviewed journals and conference proceedings and are divided into three distinct parts with different focuses:

The first part of this dissertation identifies the drivers of and barriers to the adoption of ML systems in organizations in general, and in healthcare organizations specifically. Drawing on an interview study with 14 experts from a variety of industries, an integrative overview of the factors influencing the adoption of ML systems is provided, structured according to technical, organizational, and environmental aspects. The interviews further reveal several problem areas where ML provider and ML user organizations' perceptions diverge, which can lead to the flawed design of ML systems and thus delayed integration. In a second qualitative study, specific factors affecting the integration of ML systems in healthcare organizations are derived based on 22 expert interviews with physicians with ML expertise, and with health information

technology providers. In a following step, these interviews are used to establish an operationalized maturity model, which allows for the analysis of the status quo in the adoption process of ML systems in healthcare organizations.

How the identified requirements for the organizational introduction of ML systems can be fulfilled is subject of the second part of this dissertation. First, the concept of data donation is introduced as a potential mechanism for organizations, particularly in the healthcare sector, to achieve a valid database. More specifically, individuals' donation behavior along with its antecedents, such as privacy risks and trust, and under different emotional states, is investigated based on an experimental study among 445 Internet users. Next, a design for rendering ML systems more transparent is proposed and evaluated using a questionnaire and an experiment among 223 Internet users. Thereby, the relevance of transparency for building trust among potential users and the resulting willingness to pay for transparent designs is highlighted. A qualitative study is further employed to reveal what motivates potential users, and especially the elderly, to accept health-related ML systems.

The third part of this work includes a simulation study that presents the potential impact of adopting ML systems for organizational learning. The results suggest that an organization's employees can be relieved of some of their learning burden through the application of ML systems, but the systems must be reconfigured appropriately over time. This holds especially true in case of rapid environmental changes, such as those caused by the COVID-19 pandemic.

In summary, this dissertation assumes a socio-technical perspective to shed light on the integration of ML systems in organizations. It helps organizations better understand the complex interplay of technical, organizational, human, and environmental factors that are critical to the successful adoption of ML systems, enabling decision makers to target scarce corporate resources more effectively. Moreover, this work enables IS researchers to better grasp the specifics of ML systems, provide required adjustments to theoretical foundations, and sharpen their understanding of the contextual factors involved in the adoption of ML systems in organizations.

Abstract (German Version)

Das Potenzial des Maschinellen Lernens (ML) und darauf basierender Systeme ist in den letzten Jahren stetig gewachsen. Die Fähigkeit von ML-Systemen, aus großen Mengen an Daten schnell und systematisch Zusammenhänge zu erlernen, die zur Analyse neuer Daten genutzt werden können, um aussagekräftige Vorhersagen zu treffen, ermöglicht es Organisationen aller Branchen, ihre Prozesse effektiver und effizienter zu gestalten. Insbesondere das Gesundheitswesen könnte zukünftig stark von ML-Systemen profitieren, da die Fähigkeiten dieser Systeme dazu beitragen könnten, trotz vieler dringlicher Problemstellungen wie dem akuten Fachkräftemangel eine angemessene Versorgung von Patienten¹ sicherzustellen (z. B. durch Diagnoseunterstützung).

Allerdings scheitern viele Organisationen derzeit noch daran, das Potenzial von ML-Systemen für sich zu nutzen, da die Einführung dieser Systeme keine triviale Aufgabe darstellt. Vielmehr erfordert die Integration von ML-Systemen, dass die Organisation neuartige, vielschichtige Anforderungen identifiziert und erfüllt, die von früheren konventionellen Technologien nicht bekannt waren. Dies liegt vorwiegend daran, dass ML-Systeme spezifische Merkmale aufweisen. Insbesondere besitzen ML-Systeme aufgrund ihres datenbasierten Lernansatzes probabilistische Eigenschaften, sodass ihre Anwendung zu fehlerhaften Ergebnissen führen kann und ihre Funktionsweise oftmals intransparent ist. Gerade im Gesundheitswesen, in dem das Leben der Patienten von der korrekten Diagnostik und Behandlung abhängt, führen diese Merkmale dazu, dass ML-Systeme nicht nur hilfreich sind, sondern – falsch eingeführt – auch schwerwiegende negative Konsequenzen nach sich ziehen können. Da die bestehende Forschung zur Einführung konventioneller Technologien die charakteristischen Eigenschaften von ML-Systemen bislang nicht berücksichtigt, ist das Ziel dieser Dissertation, die komplexen Anforderungen an eine erfolgreiche Adoption von ML-Systemen in Organisationen besser zu verstehen, um die Potenziale dieser Systeme nachhaltig heben zu können. Die drei qualitativen, zwei experimentellen und eine Simulationsstudie, die in dieser kumulativen Dissertation enthaltenen sind, wurden in von Fachexperten begutachteten Fachzeitschriften und Konferenzberichten veröffentlicht und gliedern sich in drei verschiedene Teilbereiche:

Im ersten Teil dieser Arbeit werden die Triebkräfte und Hemmnisse für die Adoption von ML-Systemen in Organisationen im Allgemeinen und in Gesundheitsorganisationen im

¹ Im Folgenden wird der besseren Lesbarkeit halber das generische Maskulinum verwendet. Weibliche und andere Geschlechteridentitäten sind darin ausdrücklich eingeschlossen.

Spezifischen identifiziert. Auf Grundlage einer Interviewstudie mit 14 Experten aus unterschiedlichen Branchen wird eine integrative Gesamtübersicht der Einflussfaktoren erstellt, die nach technischen, organisatorischen und umweltbezogenen Aspekten gegliedert ist. Darüber hinaus zeigen die Interviews mehrere Problembereiche auf, in denen die Vorstellungen von Anbieter- und Nutzerorganisationen von ML-Systemen voneinander abweichen, was zu einer fehlerhaften Konzeption von ML-Systemen und somit zu einer verzögerten Einführung führen kann. Zusätzlich werden auf der Grundlage von 22 Experteninterviews mit Ärzten, die über ML-Fachwissen verfügen, und Anbietern von diagnostischen Technologien spezifische Faktoren für das Gesundheitswesen abgeleitet, die die Einführung von ML-Systemen beeinflussen. Diese werden in einem weiteren Schritt herangezogen, um ein operationalisiertes Reifegradmodell zu entwickeln, welches erlaubt, den Status Quo im Einführungsprozess von ML-Systemen in Organisationen des Gesundheitswesens zu analysieren.

Wie die Voraussetzungen für die organisationale Adoption von ML-Systemen erfüllt werden können, wird im zweiten Teil dieser Dissertation behandelt. Zunächst wird das Konzept der Datenspende als potenzieller Mechanismus für Organisationen, insbesondere im Gesundheitswesen, vorgestellt, um eine valide Datenbasis zu erhalten. Im Einzelnen werden das Spendenverhalten von Individuen, darauf wirkende Einflussfaktoren wie Datenschutzrisiken und Vertrauen, sowie die Bedeutung verschiedener emotionaler Zustände für die Spende auf der Grundlage einer experimentellen Studie unter 445 Internetnutzern untersucht. Anschließend wird ein Design für eine transparentere Gestaltung von ML-Systemen präsentiert und anhand eines Fragebogens und eines Experiments unter 223 Internetnutzern evaluiert. Hierbei wird hervorgehoben, welche Relevanz die Transparenz für das Vertrauen potenzieller Nutzer hat und welche Zahlungsbereitschaft für transparente Designs daraus entsteht. Mithilfe einer qualitativen Studie wird zudem gezeigt, was potenzielle Nutzer und insbesondere ältere Menschen dazu bewegt, gesundheitsbezogene ML-Systeme zu akzeptieren.

Der dritte Teil dieser Arbeit umfasst eine Simulationsstudie, welche die möglichen Auswirkungen einer Adoption von ML-Systemen für das organisationale Lernen darlegt. Die Ergebnisse legen nahe, dass Mitarbeiter einer Organisation durch den Einsatz von ML-Systemen in ihren Lernanstrengungen entlastet werden können, die Systeme aber regelmäßig gewartet werden müssen. Dies gilt insbesondere im Falle von schnellen Umweltveränderungen, wie sie z.B. durch die COVID-19-Pandemie verursacht werden.

Zusammengefasst nimmt diese Dissertation eine sozio-technische Perspektive ein, um die Adoption von ML-Systemen in Organisationen zu beleuchten. Sie hilft Organisationen, das

komplexe Zusammenspiel aus technischen, organisatorischen und menschlichen Aspekten sowie Umweltfaktoren besser zu verstehen, die für die erfolgreiche Einführung von ML-Systemen ausschlaggebend sind, und somit ihre knappen Ressourcen gezielter einzusetzen. Darüber hinaus ermöglicht diese Arbeit Forschern, die Spezifika von ML-Systemen besser zu erfassen, die erforderlichen Anpassungen der theoretischen Grundlagen vorzunehmen und ihr Verständnis für die kontextuellen Faktoren zu schärfen, die bei der Adoption von ML-Systemen in Organisationen eine Rolle spielen.

Table of Contents

List of Figures	XII
List of Tables.....	XIII
List of Abbreviations.....	XIV
1 Introduction	1
1.1 Overarching Motivation and Problem Description.....	1
1.2 Overarching Research Questions and Contributions.....	2
1.3 Structure of the Thesis.....	5
2 Related Work.....	10
2.1 Machine Learning Systems in Healthcare Organizations.....	10
2.2 Organizational Adoption of Machine Learning Systems	13
3 Paper 1.A: A New Organizational Chassis for Artificial Intelligence – Exploring Organizational Readiness Factors	20
3.1 Introduction	21
3.2 Theoretical and Conceptual Background	22
3.2.1 Artificial Intelligence and Adoption	22
3.2.2 TOE Framework and Diffusion of Innovation.....	24
3.3 Qualitative Research Methodology	25
3.3.1 Research Design	26
3.3.2 Sample and Data Collection.....	26
3.3.3 Coding Concept	28
3.4 Results and Discussion.....	29
3.5 Conclusion, Limitations, and Future Research.....	40
4 Paper 1.B: Adoption of Machine Learning Systems for Medical Diagnostics in Clinics: Qualitative Interview Study	42
4.1 Introduction	44
4.1.1 Machine Learning Systems for Medical Diagnostics	44
4.1.2 Challenges of Adopting Machine Learning Systems in Clinics	45
4.1.3 Objectives and Research Approach	46
4.2 Methods	48
4.2.1 Overview.....	48
4.2.2 Data Collection and Sample.....	50
4.2.3 Directed Content Analysis	51
4.2.4 Maturity Model Development.....	53
4.3 Results	55
4.3.1 Factors Influencing the Adoption Process of ML Systems in Clinics	55
4.3.2 A Maturity Model for ML Systems in Clinics.....	66

4.4	Discussion.....	70
4.4.1	Principal Findings.....	70
4.4.2	Limitations and Future Research Opportunities	70
4.4.3	Theoretical Contributions	71
4.4.4	Practical Contributions	72
5	Paper 2.A: Acting Egoistically in a Crisis: How Emotions Shape Data Donations	73
5.1	Introduction	74
5.2	Theoretical Background	76
5.3	Hypotheses Development.....	78
5.4	Methodology.....	80
5.5	Results	83
5.6	Discussion.....	85
5.7	Conclusion.....	88
6	Paper 2.B: Opening the Black Box: Consumers' Willingness to Pay for Transparency of Intelligent Systems	89
6.1	Introduction	90
6.2	Related Work.....	92
6.2.1	Intelligent Systems.....	92
6.2.2	Computer Science Perspective on Transparency	93
6.2.3	Information Systems Perspective on Transparency	95
6.3	Theoretical Framework	96
6.3.1	Willingness to Pay for Transparency Features	96
6.3.2	Mechanisms Behind Purchase Decision	97
6.4	Research Design and Data Collection	99
6.4.1	Study A: Measuring Willingness to Pay.....	99
6.4.2	Study B: Explaining Purchase Decision	101
6.4.3	Sample	102
6.5	Data Analysis and Results	102
6.5.1	Results of Study A	102
6.5.2	Results of Study B	103
6.6	Discussion.....	105
6.7	Conclusion.....	108
7	Paper 2.C: “Hello, I’m Here to Help You” – Medical Care Where It Is Needed Most: Seniors’ Acceptance of Health Chatbots.....	109
7.1	Introduction	110
7.2	Related Research	111
7.3	Qualitative Research Methodology	114

7.3.1 Research Design, Sample, and Data Collection.....	114
7.3.2 Content Analysis.....	115
7.4 Results and Discussion.....	116
7.4.1 Known UTAUT2 Factors	117
7.4.2 Redefined UTAUT2 Factors.....	118
7.4.3 Newly Added Factors	120
7.5 Conclusion, Limitations, and Future Research.....	127
8 Paper 3.A: Coordinating Human and Machine Learning for Effective Organizational Learning	130
9 Discussion of Contributions and Concluding Remarks.....	132
9.1 Overarching Theoretical Contributions.....	132
9.2 Overarching Practical Contributions	135
9.3 Concluding Remarks	138
References	140

List of Figures

Figure 1: Outline of the Thesis.....	8
Figure 2: Overview of the Combined TOE Framework and Diffusion of Innovations Theory (based on DePietro et al. 1990; Rogers 1995)	15
Figure 3: Overview of the NASSS Framework (based on Greenhalgh et al. 2017)	16
Figure 4: TOE Framework as Conceptual Base (based on DePietro et al. 1990; Rogers 1995).....	25
Figure 5: Content Analysis Process (based on Hsieh and Shannon 2005).....	26
Figure 6: Extended and Deepened Framework for AI Adoption	38
Figure 7: Overview of Research Procedure (Illustration based on Jöhnk et al. 2020).....	49
Figure 8: Integrative Framework for the Adoption Process of Machine Learning Systems in Clinics	56
Figure 9: Determine Your Clinic's Readiness for Machine Learning-Supported Diagnostics (Screenshot 1 of the Web Application)	69
Figure 10: Thank You for Using the Maturity Model (Screenshot 2 of the Web Application).....	69
Figure 11: Research Model for Mechanisms Behind Purchase Decision	98
Figure 12: Experiment Procedure	100
Figure 13: SEM Results	105
Figure 14: UTAUT2 Model (based on Venkatesh et al. 2012).....	113
Figure 15: Seniors' Acceptance of HC – An Extended UTAUT2 Model	116

List of Tables

Table 1: Overview of Publications Included in This Thesis	5
Table 2: Participant Overview.....	27
Table 3: Findings: Examination of Proposed Factors in TOE Framework.....	30
Table 4: Overview of Interviewed Experts	51
Table 5: Maturity Model for Machine Learning Systems in Clinics	67
Table 6: Measurement Scales.....	82
Table 7: Characteristics of Sample.....	82
Table 8: Validity Measurements	84
Table 9: Multi-Group-Analysis (* $P < .05$).....	84
Table 10: Overview of Pre-Studies and Experiment Parameters	99
Table 11: Examples of Construct Operationalization	101
Table 12: Descriptive Statistics Regarding Measured WTP	102
Table 13: Factor Loadings and Reliability	103
Table 14: Factor Correlations, Reliability and Validity Measures.....	104
Table 15: Overview of Tested Hypotheses	105
Table 16: Summary of Results Regarding Organizational Learning Effectiveness.....	131

List of Abbreviations

AI	Artificial intelligence
AVE	Average variance extracted
B2B	Business-to-business
B2C	Business-to-consumer
CFA	Confirmatory factor analysis
CFI	Comparative fit index
CIO	Chief information officer
CMB	Common method bias
COVID-19	Coronavirus disease 2019
CR	Composite reliability
CRM	Customer relationship management
CS	Computer science
CT	Computed tomography
DICOM	Digital imaging and communications in medicine
DOI	Digital object identifier
ECIS	European Conference on Information Systems
EFA	Exploratory factor analysis
ERP	Enterprise resource planning
FHIR	Fast healthcare interoperability resources
GDPR	General Data Protection Regulation
H	Hypothesis
HCI	Human-computer interaction
HICSS	Hawaii International Conference on System Sciences
HIT	Health information technology
HR	Human resources
ICHMS	International Conference on Human-Machine Systems
IEEE	Institute of Electrical and Electronics Engineers
IS	Information systems
IT	Information technology
KBS	Knowledge-based systems
KPI	Key performance indicator
MGA	Multi-group analysis
mHealth	Mobile health
ML	Machine learning
MRI	Magnetic resonance imaging
NASSS	Nonadoption, abandonment, scale-up, spread, and sustainability
NLP	Natural language processing
NYOP	Name-your-own-price

P	Proposition
PACIS	Pacific Asia Conference on Information Systems
R&D	Research and development
RMSEA	Root mean squared error of approximation
ROI	Return on investment
RQ	Research question
SD	Standard deviation
SRMR	Standardized root mean square residual
TAM	Technology acceptance model
TOE	Technology-organization-environment
TPB	Theory of planned behavior
UMLS	Unified medical language system
UTAUT	Uniefied theory of acceptance and use of technology
VHB	Verband der Hochschullehrer für Betriebswirtschaft e.V.
VIF	Variable inflation factor
WI	Internationale Tagung Wirtschaftsinformatik
WTP	Willingness to pay

1 Introduction

“Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don’t think AI will transform in the next several years.” Andrew Ng (published in Lynch 2017)

1.1 Overarching Motivation and Problem Description

Machine learning (ML) as a general-purpose technology has applications in a variety of different use cases, such as intelligent assistants, facial recognition, or fraud detection (Brynjolfsson and McAfee 2017a). Thereby, ML holds the potential to greatly influence how organizations create value (Brynjolfsson and McAfee 2017a). As Andrew Ng, a computer scientist, entrepreneur, and adjunct professor who has been influential in the field of ML states above, the technology is just beginning to spread across a wide range of industries (Lynch 2017). In this context, ML enables a new form of systems that are able to analyze data, recognize underlying patterns from it, and make predictions based on the learned correlations (Jordan and Mitchell 2015; Mitchell 1997). ML systems thus represent a branch of artificial intelligence (AI), because, like humans, they are able to learn by themselves (McCarthy 2007). This capability makes ML systems compelling for organizations that seek to augment or automate their processes in order to make them more effective and efficient. To date, however, most organizations are still in the early stages of this effort (Brynjolfsson and McAfee 2017a). A 2020 global survey by Google’s online ML community Kaggle shows that 61.5% of respondents indicated that their organization has not yet adopted ML systems, or is in the early stages of adoption (7.8 % responded “I do not know”; Kaggle 2020). Since only organizations that have already engaged with the Kaggle ML community were considered in this survey, it can be assumed that the actual adoption rate is even lower. This is mainly due to the special characteristics of ML systems, which strongly differ from previous technologies, such as rule-based enterprise resource planning (ERP) or customer relationship management (CRM) systems. In particular, ML systems often exhibit high levels of inscrutability due to their data-based learning approach (Rudin 2019). Moreover, unlike earlier systems, they do not always lead to predictable results, but may suggest erroneous strategies (Domingos 2012). Their integration into the system landscape therefore differs significantly from the integration of other systems (Amershi et al. 2019). As a result, the adoption of ML systems requires fundamental organizational change. However, if this change fails, it can have far-reaching consequences for the organization, which may then fall behind its competitors (Rana et al. 2021). Therefore,

researchers such as Benbya et al. (2021) call for a closer look at the characteristics of ML systems that might have an impact on the adoption of these systems. In addition, Kane et al. (2021) suggest investigating how organizations can meaningfully prepare for a future with ML systems. Such research could help organizations reap the benefits of ML systems while avoiding the risks associated with unsuccessful system adoption.

While ML systems are driving major changes in all industries, healthcare in particular could benefit from the capabilities of these systems. The healthcare system is currently facing multiple major challenges, such as the coronavirus disease 2019 (COVID-19) pandemic, ongoing demographic change, and an increasing shortage of skilled workers (Lewnard et al. 2020; Shrestha 2000; World Health Organization 2016). Against this backdrop, ML systems could improve processes such as diagnostics in healthcare organizations in order to address the aforementioned challenges (Lebovitz et al. 2021). However, healthcare organizations in particular differ vastly from other organizations in that they are responsible for the well-being of their patients or, in the case of for-profit healthcare organizations, their customers (Golden 2006; Thompson et al. 2017), which adds to the challenge of adopting opaque, potentially flawed ML systems. To examine the specific conditions that enable the successful integration of ML systems into healthcare organizations, the specific context of healthcare must thus be considered (Davison and Martinsons 2015). In this vein, Shaw et al. (2019) call for a dedicated investigation of ML system adoption in healthcare organizations.

Overall, the success of ML systems and the unleashing of their immense potential will, above all, depend on whether and how well organizations are able to adopt them (Rana et al. 2021). This work helps to identify and fulfill the factors necessary for the adoption of ML systems and provides an outlook on what impact this may have on the organization's performance.

1.2 Overarching Research Questions and Contributions

The potential of ML systems has been growing considerably for several years due to the availability of ever increasing data, the rise in cheaply available computing capacity, and the existence of more powerful algorithms for implementing ML systems (Brynjolfsson and McAfee 2017b; Jordan and Mitchell 2015). Due to these new technological achievements, many organizations are eager to adopt ML systems to improve their organizational performance. However, now it is not the technology itself that sets the limits, but the adoption and management of ML systems in organizations (Brynjolfsson and McAfee 2017a). Given the unique characteristics of ML systems, organizations face a challenging new landscape that is

vastly different from the adoption of previous technologies (Coombs et al. 2020). This is especially the case as the adoption of ML systems, by virtue of their characteristics, not only requires the renewal of existing legacy technology, but has also far-reaching implications for an organization's structural, cultural, and political legacy (Marabelli et al. 2021; Willcocks 2020). In fact, the adoption of ML systems cannot be viewed from a solely technical angle, but needs to be approached from a socio-technical, multidisciplinary perspective, as it significantly involves a broad range of diverse stakeholders (e.g., customers, workforce) and their needs (Asatiani et al. 2021; Coombs et al. 2020; Lebovitz et al. 2021; Li et al. 2021; Marabelli et al. 2021; Willcocks 2020). After all, employees will, for example, be heavily affected by the new technology, e.g., through the increased automation enabled by ML systems (Willcocks 2020). Failure to adequately consider these facets when adopting ML systems can have serious consequences for the organization in question. Not only are potential opportunities offered by the technology being missed, but there could be negative implications for the organization's operational efficiency, employee satisfaction, and competitiveness as well (Rana et al. 2021). Actively shaping the adoption process of ML systems is therefore inevitable for the ultimate success of the technology and for the organization as a whole (Willcocks 2020).

While successfully adopting ML systems could benefit organizations in all industries, the upsides and downsides of integrating ML systems are particularly glaring in healthcare (Lebovitz et al. 2021). Healthcare organizations are high-stakes contexts, which can severely complicate the adoption of ML systems (Golden 2006; Lebovitz et al. 2021). This is also why current research cautions against rushing to adopt ML systems in healthcare organizations without due consideration (Lebovitz et al. 2021). As a result, the drivers and barriers to ML system adoption should be identified to ensure that organizations in general, and healthcare organizations in particular, have an overview of what factors influence a successful adoption, as reflected in the first research question (RQ) of this thesis:

RQ1: *What are the major factors that influence the successful adoption of ML systems in organizations, specifically in the healthcare sector?*

While it is valuable to gain an overview of the key factors influencing the organizational adoption of ML systems, it is even more important to explore *how* to overcome the identified barriers and leverage the driving forces for adoption. Although many companies are already beginning to be aware of the problems associated with ML systems, they are still struggling to solve these problems in order to effectively integrate ML systems into their processes (Brynjolfsson and McAfee 2017a; Kaggle 2020). Therefore, the second RQ is aimed at gaining

deeper insight into factors influencing the successful adoption of ML systems in order to help organizations fulfill these requirements. Among the most relevant factors are the data that forms the basis for ML systems, the handling of ML systems' unique characteristics, especially their opacity and error-proneness, and the reluctance of potential users due to these characteristics (e.g., Benbya et al. 2020; Berente et al. 2021; Jöhnk et al. 2020). In particular, many organizations do not yet own sufficient high-quality data and have difficulty finding entry points at which to collect it (Giermindl et al. 2021; Jöhnk et al. 2020; Marabelli et al. 2021). Many organizations also face the challenge of how to make ML systems more transparent to identify potential errors in the systems (e.g., Asatiani et al. 2021; Benbya et al. 2021; Berente et al. 2021; Rudin 2019). Especially at the beginning, many ML systems still perform poorly, which can lead to a defensive attitude among users, above all in high-stakes environments such as healthcare (e.g., Benbya et al. 2021; Lebovitz et al. 2021; Rudin 2019). However, in order to optimize the ML systems over time using newly collected data, acceptance among potential users such as an organization's customers must be ensured (Benbya et al. 2021). The second RQ therefore explores how these barriers to adoption can be overcome:

RQ2: How can organizations address selected factors that facilitate the successful adoption of ML systems?

Once an organization has successfully implemented ML systems, it can have a far-reaching impact on various business performance metrics. One of the most vital performance indicators in an organization is the level of knowledge obtained (e.g., Bushee 1998; Cohen and Levinthal 1989). ML systems that are able to learn autonomously from data are now contributing to this knowledge (Jordan and Mitchell 2015; Mitchell 1997). Therefore, organizational learning is significantly influenced by the adoption of ML systems. In particular, ML systems can help explore new perspectives on reality and introduce this fresh knowledge into an organization (e.g., Choudhury et al. 2021; Ransbotham et al. 2020). This not only affects the level of knowledge of organizational members, but can also significantly alter an organization's ability to deal with unforeseen changes in an organization's environment (e.g., Ransbotham et al. 2020). However, the organizational learning process is complex, as it involves many different actors interacting with each other (e.g., Argote et al. 2021; Dodgson 1993). In line with the call of Argote et al. (2021) to further investigate ML systems' impact on organizational learning, the third RQ thus asks how the adoption of ML systems affects this process:

RQ3: What is the impact of the adoption of ML systems on organizational learning?

1.3 Structure of the Thesis

In response to the three RQs formulated, this thesis presents six research papers² that have been published in a range of peer-reviewed outlets, including different journals and conference proceedings (see Table 1 for an overview).

Table 1: Overview of Publications Included in This Thesis

Research Question	Paper	Citation
RQ1	1.A	Pumplun, Luisa; Tauchert, Christoph; Heidt, Margareta (2019): A New Organizational Chassis for Artificial Intelligence – Exploring Organizational Readiness Factors . In: European Conference on Information Systems (ECIS), Stockholm-Uppsala, Sweden. VHB-Ranking: B.
	1.B	Pumplun, Luisa; Fecho, Mariska; Wahl, Nihal; Peters, Felix; Buxmann, Peter (2021): Adoption of Machine Learning Systems for Medical Diagnostics in Clinics: Qualitative Interview Study . In: Journal of Medical Internet Research, DOI: 10.2196/29301. Impact Factor: 5.43.
RQ2	2.A	Pumplun, Luisa; Wagner, Amina; Olt, Christian; Zöll, Anne; Buxmann, Peter (2022): Acting Egoistically in a Crisis: How Emotions Shape Data Donations . In: Hawaii International Conference on System Sciences (HICSS), Maui, Hawaii, United States. VHB-Ranking: C.
	2.B	Peters, Felix; Pumplun, Luisa; Buxmann, Peter (2020): Opening the Black Box: Consumers’ Willingness to Pay for Transparency of Intelligent Systems . In: European Conference on Information Systems (ECIS), A Virtual AIS Conference. VHB-Ranking: B.
	2.C	Mesbah, Neda; Pumplun, Luisa (2020): “Hello, I’m Here to Help You”: Medical Care Where It Is Needed Most: Senior’s Acceptance of Health Chatbots . In: European Conference on Information Systems (ECIS), A Virtual AIS Conference. VHB-Ranking: B.
RQ3	3.A	Sturm, Timo; Gerlach, Jin; Pumplun, Luisa; Mesbah, Neda; Peters, Felix; Tauchert, Christoph; Nan, Ning; Buxmann, Peter (2021): Coordinating Human and Machine Learning for Effective Organizational Learning . In: MIS Quarterly, DOI: 10.25300/MISQ/2021/16543. VHB-Ranking: A+.

To answer RQ1, papers 1.A and 1.B explore the factors that influence the adoption of ML systems in organizations. Paper 1.A draws on qualitative expert interviews to constitute a comprehensive framework on the drivers of and barriers to organizational adoption of ML systems, and to elaborate on the potential mismatches between the requirements of provider

² Research papers have been marginally adjusted from the original publication to ensure a unified layout. The contributions are formulated in the first person plural, as several authors collaborated on each publication.

and user organizations of ML systems that could inhibit the integration of these systems. The second paper, 1.B, transfers this research goal to the healthcare context. Building on qualitative expert interviews with physicians, hospital managers, and managers of health information technology (HIT) providers, the factors specific to the adoption of ML systems in healthcare organizations were identified. In addition, the interviews served as a basis for the methodological development of a maturity model for the introduction of ML systems in clinics. While both 1.A and 1.B provide the starting point to dive deeper into the prerequisites for the adoption of ML systems in organizations, especially in the healthcare sector, papers 2.A, 2.B, and 2.C analyze more in-depth how organizations can meet requirements to adopt ML systems. In this regard, paper 2.A shows a feasible avenue for achieving the necessary database for training ML systems. Specifically, this research paper investigates by means of an experimental survey study how much (health) data users are willing to donate and how this eagerness to donate is formed. This allows organizations to actively incentivize the donation behavior of users in order to improve their database. Apart from the necessary data, the transparency of ML systems plays an essential role in the success of ML systems in organizations and for their customers. In this regard, paper 2.B, an experimental survey study, provides an exemplary implementation of transparency features that can be applied to explain a black box ML system. In addition, the research examines customers' attitudes towards ML systems' transparency and how these are shaped. Apart from the transparency of ML systems, other factors contribute to an organization's customers' acceptance of the technology. Therefore, paper 2.C takes a closer look at customers' attitudes toward ML systems and employs a specific ML system from the healthcare domain to examine the factors that motivate a particular user group to utilize these systems. Healthcare organizations may use the results of paper 2.C to guide the design of ML systems offered. Once ML systems are successfully adopted in and around the organization, there are far-reaching consequences for the organization as a whole. A particularly relevant organizational dimension on which ML systems have an impact is the organizational knowledge level, given the pivotal nature of learning for the survival of the organization (e.g., Bushee 1998; Cohen and Levinthal 1989). Paper 3.A thus employs a series of agent-based simulations to explore ML systems' effects on organizational learning under different conditions, such as active reconfiguration of ML systems or environmental turbulence as caused by COVID-19, for example. The potential impact of ML systems on an organization's knowledge base is illustrated by the example of drug development, which will be strongly affected by ML systems due to the high importance of scientific knowledge for pharmaceutical discoveries.

The research papers draw on different theoretical backgrounds to explain the factors that influence the adoption of ML systems and the impacts that follow from their integration. This involves the application of organizational theoretical approaches as well as the use of individually focused theories to study, for instance, the behavior of an organization's customers (see Chapter 2.2 for a more detailed description). In addition, a variety of research methods are employed to examine the adoption of ML systems from different perspectives. While paper 1.A and 1.B are based on semi-structured qualitative expert interviews, paper 2.C leverages walk-throughs and exploratory interviews with users as a foundation. On top of the interviews, contribution 1.B applies an established method for maturity model development. Both papers 2.A and 2.B are experimental survey studies with an initial online experiment followed by a questionnaire. The last work, 3.A, relies on agent-based simulation to predict possible future effects of ML systems on organizational learning.

The referenced research papers can be found in the Chapters 3 to 8. Chapter 2 leads toward these contributions by describing the research context in more detail and identifying the relevant theoretical foundations. This thesis concludes by situating the findings of the papers within existing research, clarifying their contributions to theory and practice, and suggesting possibilities for further research in Chapter 9 (see Figure 1 for an outline of the thesis).

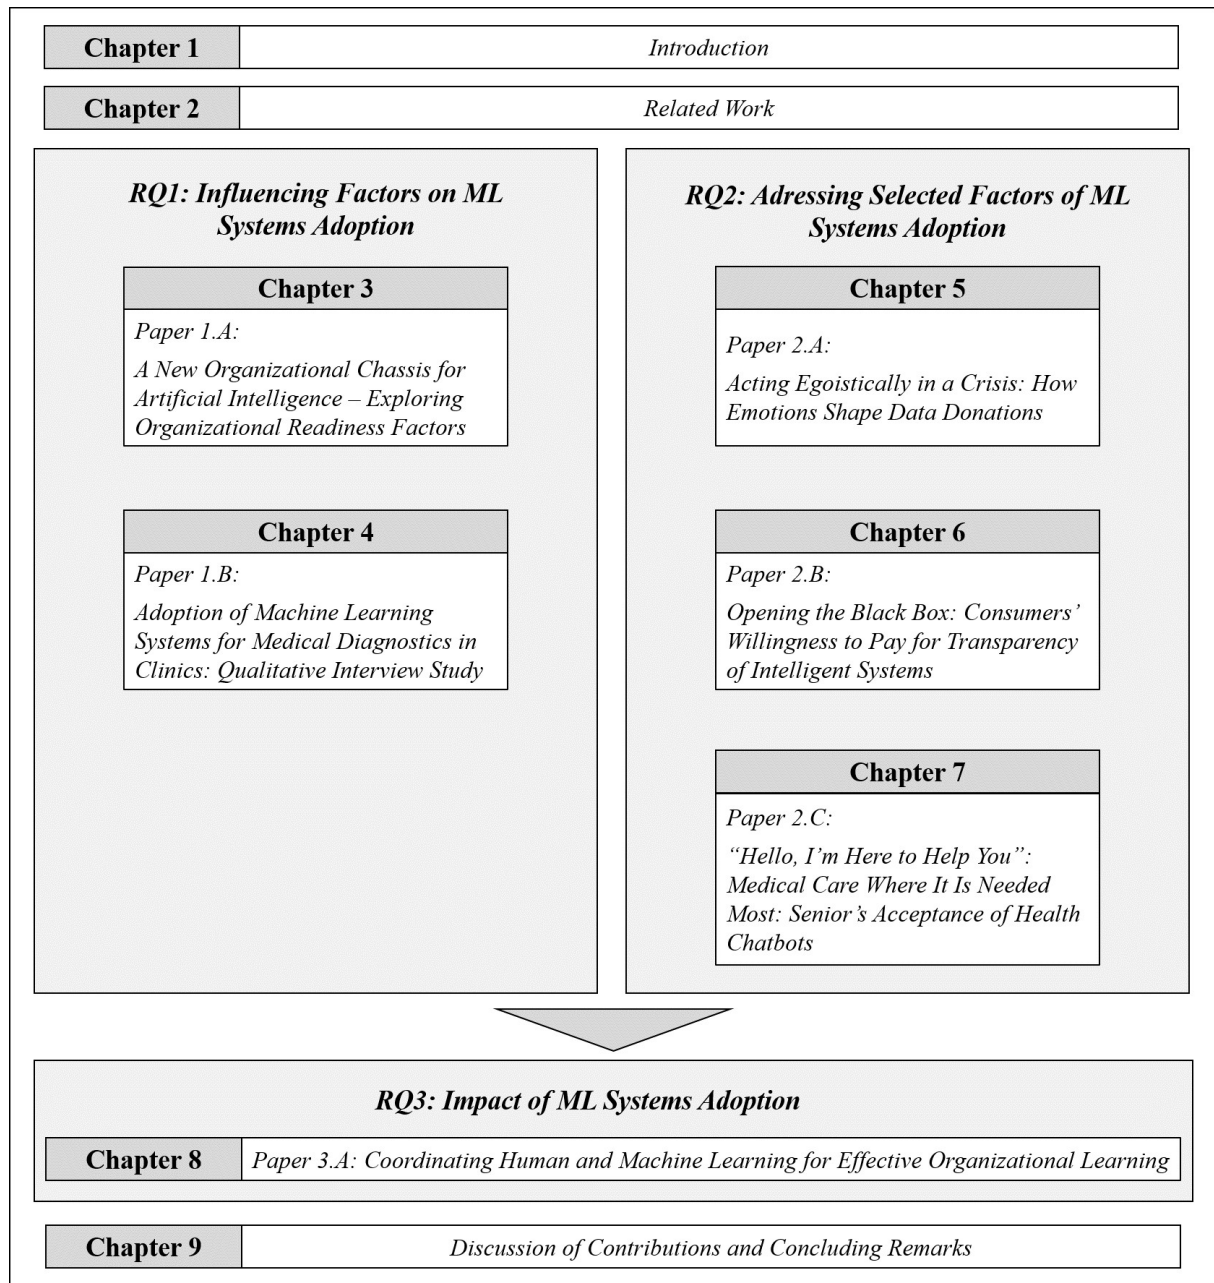


Figure 1: Outline of the Thesis

Apart from the publications included in this dissertation, I contributed to the following peer-reviewed publications during my doctoral studies at the Technical University of Darmstadt, Germany:

- Pumplun, Luisa; Wiefel, Jennifer; Wächter, Katharina; Barth, Niklas; Buxmann, Peter (2021): **Smart Car Service Adoption: Investigating the Role of Information Privacy**. In: Pacific Asia Conference on Information Systems (PACIS), Dubai, United Arab Emirates. VHB-Ranking: C.

- Pumplun, Luisa; Fecho, Mariska; Wahl, Nihal; Buxmann, Peter (2021): **Machine Learning Systems in Clinics – How Mature Is the Adoption Process in Medical Diagnostics?** In: Hawaii International Conference on System Sciences (HICSS), Maui, Hawaii, United States. VHB-Ranking: C.
- Reuter-Oppermann, Melanie; Wolff, Clemens; Pumplun, Luisa (2021): **Next Frontiers in Emergency Medical Services in Germany: Identifying Gaps between Academia and Practice.** In: Hawaii International Conference on System Sciences (HICSS), Maui, Hawaii, United States. VHB-Ranking: C.
- Roth, Elisa; Möncks, Mirco; Bohné, Thomas; Pumplun, Luisa (2020): **Context-Aware Cyber-Physical Worker Assistance in Industrial Systems: A Human Activity Recognition Approach.** In: IEEE International Conference on Human-Machine Systems (ICHMS), Rome, Italy.
- Pumplun, Luisa; Buxmann, Peter (2020): **Intelligent Systems and Hospitals: Joint Forces in the Name of Health?** In: Internationale Tagung Wirtschaftsinformatik (WI), Potsdam, Germany. VHB-Ranking: C.

2 Related Work

In the following, the research context of this dissertation is elaborated upon. First, the characteristic features of healthcare organizations and ML systems that influence the organizational deployment of these systems are discussed. Subsequently, different frameworks that can be used to study the adoption of ML systems in organizations are described in order to lay a foundation for the following research.

2.1 Machine Learning Systems in Healthcare Organizations

Today's population healthcare is threatened by several societal, pandemic, and workforce issues. Demographic change is causing the average age in many regions to increase. Simultaneously, the number of people with chronic, degenerative diseases who are in need of medical care is on the rise (Shrestha 2000). Along with the change in age structure, the COVID-19 pandemic has contributed to citizens requiring longer and more complex medical care when hospitalized with symptoms (Lewnard et al. 2020). The impact of the COVID-19 pandemic is hitting healthcare facilities particularly hard, as there is a global shortfall of skilled health workers (World Health Organization 2016). While these current developments threaten available capacity in medical facilities, many healthcare organizations are working on digitalizing their processes (Tresp et al. 2016; Weissman et al. 2020). In the past, many workflows in healthcare organizations, such as the handling of patient records, were carried out in paper format. With the introduction of electronic patient records in many countries, more and more of these processes are being moved to the digital realm (Tresp et al. 2016). The resulting volume of digital data offers the opportunity to deploy data-driven HIT that could significantly improve the effectiveness and efficiency of healthcare delivery to counteract the aforementioned challenges (e.g., Thrall et al. 2018; Tresp et al. 2016).

Recently, the concept of ML in particular has come into focus as a way to harness the volumes of data in healthcare (e.g., Shaw et al. 2019; Thrall et al. 2018; Tresp et al. 2016). ML is a subfield of AI, a research area that deals with the “science and engineering of making intelligent machines, especially intelligent computer programs” (McCarthy 2007, p. 2). Systems based on ML (i.e., ML systems) are capable of learning a problem's solution autonomously, through experience, without the need for explicit human instructions. Instead, ML systems rely on large amounts of data, which they examine for patterns using algorithms to derive their own rules for solving the problem, which they store in models (Jordan and Mitchell 2015; Mitchell 1997). ML systems are particularly useful, compared to conventional systems, when the rules that help

solve a delimited problem are difficult to infer or express by humans. This is the case, for instance, with image data, which is frequently analyzed in the healthcare sector as well (Domingos 2012; Meskó and Görög 2020); it is not straightforward to describe exactly what a malignant nodule on computed tomography (CT) scans might look like in terms of pixels. ML systems can essentially be divided into three different subtypes (Meskó and Görög 2020; Mitchell 1997): while supervised learning is based on annotated data containing an input-output relationship (e.g. CT scans and related condition), unsupervised learning involves data without a corresponding label (e.g., CT scans). Thus, while supervised learning seeks to establish generalizable rules for the relationship between input and output values, unsupervised learning examines the data for patterns in the input values or anomalies that deviate from the norm. The third subtype of ML systems is built on reinforcement learning, which relies on a reward function to either incentivize or penalize the system's actions in its environment to learn the best possible solution policy for a problem (Meskó and Görög 2020; Mitchell 1997).

ML systems of all subtypes are becoming increasingly common in our daily lives. We use them to plan our driving routes, discover new movies that might be of interest to us, or communicate with our mobile phones (Brynjolfsson and McAfee 2017a). Besides their everyday applications, ML systems also have the potential to support many healthcare processes and thus to contribute considerably to their effectiveness and efficiency (e.g., Thrall et al. 2018; Tresp et al. 2016). Therefore, ML systems are not only being researched intensively in terms of their applicability in healthcare, but are expected to become increasingly prevalent along the entire healthcare process in practice (e.g., Brinati et al. 2020; Shaw et al. 2019; Thrall et al. 2018; Wu et al. 2019). Looking at the example of the COVID-19 pandemic alone, it is evident how broad the possible spectrum of applications for ML systems is. In the pharmaceutical industry, for instance, ML systems offer great potential to simplify the often lengthy and complex development of drugs and vaccines by suggesting relevant molecules (Lou and Wu 2021; Vamathevan et al. 2019). In analyzing pandemic progression, ML systems can help predict outbreak hotspots in advance and initiate appropriate countermeasures (Ardabili et al. 2020; Cho 2020). ML systems can also provide early clarity in the diagnosis of COVID-19. Zoabi et al. (2021) show, for example, how an ML system can diagnose COVID-19 with a high degree of certainty through the answering of simple questions. Such systems could, among others, allow us to learn more about effective drugs and vaccines, minimize disease outbreaks, and prevent the progression of severe disease through early diagnosis, thereby addressing many of today's pressing healthcare challenges.

However, apart from these potentials, ML systems exhibit specific characteristics that are different from those of previous HIT and that have prevented their widespread use in the healthcare sector to date (Kuan 2019). Most importantly, ML systems are, at their core, learning based on data rather than being programmed with explicitly coded rule sets (Jordan and Mitchell 2015; Meskó and Görög 2020; Mitchell 1997). Since complex statistical methods are used instead, many modern ML systems constitute black boxes, implying that the mechanisms behind the solution process are opaque to humans. This is especially the case if state-of-the-art algorithms, for instance for training neural networks, are applied (Rudin 2019). Moreover, the data-based learning approach results in ML systems hardly ever being able to solve tasks perfectly. Rather, the statistical approaches attempt to generalize the patterns in the data, which can result in a specific edge case being incorrectly predicted by the ML system (Domingos 2012). Integrating ML systems into the existing information technology (IT) landscape is challenging as well, since complex relationships between different artifacts (e.g., data sets, models, and source codes) have to be defined, tracked, and versioned (Amershi et al. 2019). Finally, ML systems differ from existing HIT in that they are increasingly capable of performing knowledge work tasks that could not previously be captured by technologies, which might increase peoples' skepticism toward these systems (e.g., Jammal et al. 2020; Wang et al. 2017).

An organization is an association of members who assume responsibilities in order to achieve a common goal (Stogdill 1950). If ML systems are now integrated into this social construct, challenges may arise due to their characteristics. Not only do extensive technical and procedural adjustments have to be made, but the people in and around the organization (e.g., their customers, their patients) need to be convinced of the advantages ML systems offer (Jöhnk et al. 2020; Marabelli et al. 2021). Healthcare organizations in particular exhibit unique attributes that further complicate their ability to adopt ML systems (Golden 2006). These organizations can be divided into two subgroups that are either directly involved in the care of patients or only implicitly contribute to the provision of care to people. While in direct care facilities (e.g., clinics, physician practices) the organization delivers healthcare services to patients, residents, or customers first-hand, healthcare organizations involving non-direct care settings (e.g., pharmaceutical companies, medical equipment companies) provide well-being products and services in support of direct care organizations (Hilaris 2021; Thompson et al. 2017). All healthcare organizations share the common goal of enabling the most adequate, timely, and effective care for people possible (Thompson et al. 2017). As a consequence of this objective, healthcare organizations are among the most complex and dynamic forms of human

organizations (Golden 2006; Thompson et al. 2017). Indeed, the provision of adequate healthcare cannot be accomplished by one person alone; healthcare organizations often encompass many different highly specialized disciplines and diverse stakeholder groups, such as physicians, pharmacists, patients, or the responsible government, who are all intertwined in intricate processes with each other (Golden 2006; Thompson et al. 2017). In these processes, conflicts of interests may arise since, despite the ambitious aim of providing the best possible healthcare, financial resources in healthcare organizations are limited (Golden 2006). Complicating matters further is the fact that many healthcare professionals rely on being able to make quick, autonomous decisions in their daily work (Golden 2006). In these often complex, human-centric processes, the introduction and beneficial use of ML systems is particularly challenging. With much at stake, healthcare organizations must ensure that ML systems are used only in the best interest of people – a task that is demanding due to the probabilistic properties of these systems. As a result, the adoption of ML systems in healthcare organizations is not yet very advanced (Kuan 2019). To drive the adoption of ML systems in healthcare organizations and seize the opportunities that these systems offer, it is thus imperative to improve our understanding of the drivers of and barriers to as well as the potential impacts of ML systems in healthcare organizations.

2.2 Organizational Adoption of Machine Learning Systems

The adoption of technological innovations in organizations is a topic of interest within information systems (IS) research. In the past, the organizational requirements and the impacts of adopting a variety of technologies and IT-based processes such as cloud computing, ERP systems, and e-business have been studied extensively (e.g., Law and Ngai 2007; Lin and Chen 2012; Zhu et al. 2003, 2006). This also applies to healthcare research, which is equally concerned with the introduction of technological innovations into organizations. Healthcare organizations, for instance, have been examined to understand what requirements they need to fulfill in order to adopt different HITs (e.g., Cresswell and Sheikh 2013; Hikmet et al. 2008).

The adoption of technological innovations in organizations adheres to different, consecutive phases. Research has presented various divisions of these phases, but three have emerged that describe the organizational adoption process in its entirety (Damanpour and Schneider 2006). These are (1) the initiation phase, which takes place before the actual adoption decision and consists of activities such as familiarization with the innovation and assessing its suitability for the organization; (2) the second phase, in which the organization decides on adoption and allocates corresponding resources; and (3) the final implementation phase, in which the

organization promotes user acceptance and seeks routine use of the innovation (Damanpour and Schneider 2006). From a socio-technical perspective, organizational adoption of a technological innovation is a complex process that requires consideration of not only the technical characteristics of the innovation, but also its potential impact on people and organizational procedures (e.g., Damanpour and Schneider 2009; Marabelli et al. 2021; Wischnevsky et al. 2011). Investigating the underlying mechanisms that influence the adoption of technological innovations in organizations should therefore not be done from a unidimensional standpoint, but rather follow an integrative approach.

The technology-organization-environment (TOE) framework is a holistic framework that is frequently applied in IS research (Awa and Ojiabo 2016; Cruz-Jesus et al. 2019; Furneaux and Wade 2011) to investigate the factors that “influence the process by which [an organization] adopts and implements technological innovations” (DePietro et al. 1990, p. 152). Thereby, the TOE framework developed by DePietro et al. (1990) integrates three organizational-level dimensions in one overview: the *technological*, *organizational*, and the *environmental contexts*. While the technological context accounts for the properties of the technology and the technical requirements in the organization, the organizational context covers the structures and processes of an organization that promote or inhibit innovation adoption. The environmental context, on the other hand, deals with the influences of institutions and the market in which the organization operates on the adoption process (DePietro et al. 1990). These dimensions can in turn be subdivided into factors that specify the respective dimension. For example, the characteristics of a technology being studied (technological context), the size of an organization, e.g., measured in terms of the workforce or yearly revenue (organizational context), or the governmental regulations that may have an impact on the technology adoption (environmental context; DePietro et al. 1990). As proposed by Zhu et al. (2003), the TOE framework is frequently complemented by Rogers’ (1995) diffusion of innovations theory, which identifies further technological factors influencing the adoption of innovations such as relative advantage and compatibility. Thereby, relative advantage refers to the extent to which an organization considers a technological innovation to be superior to the former solution, while compatibility is the degree to which the innovation is consistent with existing values and satisfies the actual technical requirements of the organization (Rogers 1995). Combining the TOE framework and the diffusion of innovations theory results in the following overview (see Figure 2):

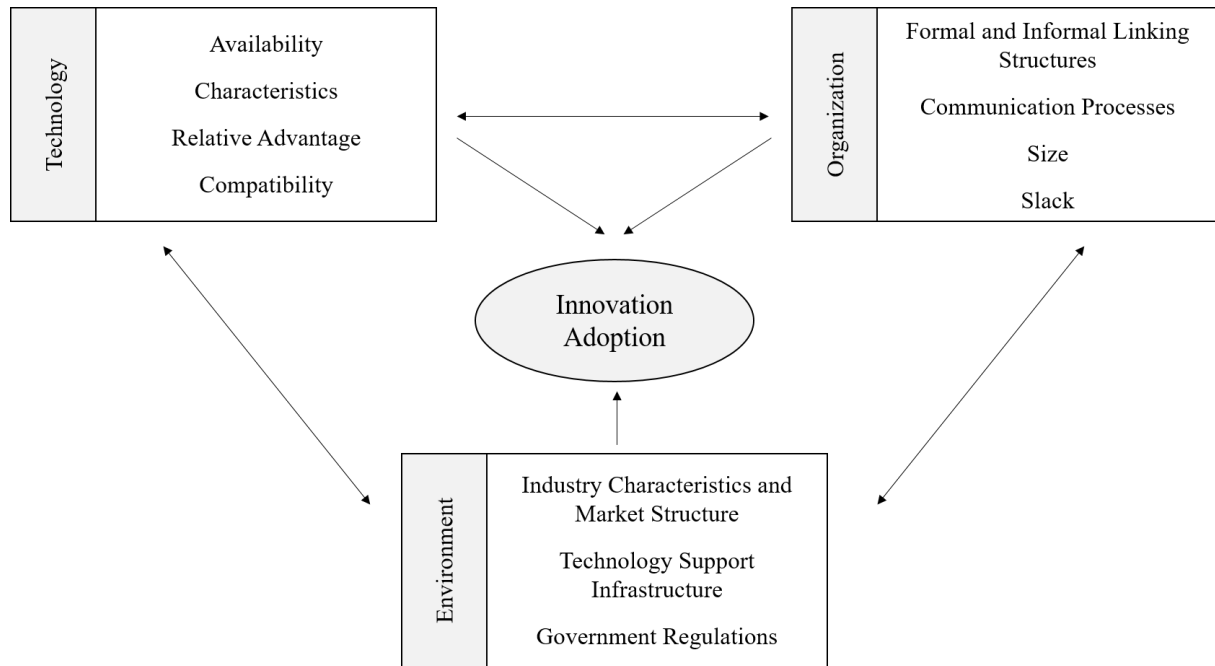


Figure 2: Overview of the Combined TOE Framework and Diffusion of Innovations Theory (based on DePietro et al. 1990; Rogers 1995)

Rapid and large-scale technological developments have led to the TOE framework and the diffusion of innovations theory being frequently used to investigate the factors influencing the adoption of technological innovations in organizations. For instance, researchers have used these theoretical foundations to study the prerequisites of adopting e-business (Zhu et al. 2003), broadband mobile applications (Chiu et al. 2017), and cloud computing (Wulf et al. 2021). Although the TOE framework and the diffusion of innovations theory do not originate from healthcare research, they have also been applied to examine the determinants of innovation adoption in healthcare organizations (e.g., Chong and Chan 2012; Sulaiman and Wickramasinghe 2014).

Nevertheless, healthcare organizations exhibit highly specific processes and structures that are not covered by the TOE framework or the diffusion of innovations theory (Golden 2006). For this reason, specific applicable approaches have been developed to investigate the adoption of technological innovations in healthcare organizations. One particularly pervasive theoretical foundation is the framework of nonadoption, abandonment, scale-up, spread, and sustainability (NASSS) by Greenhalgh et al. (2017), which among others has been applied to investigate the adoption of video consulting in healthcare, care organizing software, and patient case management via data sharing (Greenhalgh et al. 2017; James et al. 2021). Drawing on other established health and social care frameworks, Greenhalgh et al. (2017) identified seven dimensions that influence the adoption of technological innovation in healthcare organizations: these are the *adopter system* comprising the patients, their dependents, and medical staff; the

condition of the patient being diagnosed and treated; the *technology* itself and its technical features; the *value proposition* emanating from using the technology; the *healthcare organization* and its capacity to adopt the technology; the *wider system* surrounding the healthcare organization; and the continuous *embedding and adaptation over time*, which depicts the evolution and interaction of the above dimension in the future (Greenhalgh et al. 2017). The NASSS framework places a particular emphasis on the temporal adjustments that occur during the implementation phase of a technological innovation. This is the case because complex healthcare organizations, their political environments, and the conditions of patients are rapidly changing and evolving. Long-term adoption success can therefore only be ensured if an adaptive, targeted interaction of these dimensions is achieved across time (Greenhalgh et al. 2017). As with the TOE framework, the dimensions are divided into corresponding factors. For example, while the dimension “condition” distinguishes between an illness of a patient and possible co-occurring comorbidities as an influence on technology adoption, the dimension adopter system can be divided according to the changes in procedures and identities that the staff must accept, the acceptance of the patients encountering the technology, and the demands on the patients’ caregiver networks (Greenhalgh et al. 2017). This results in the following overview of the NASSS framework (see Figure 3):

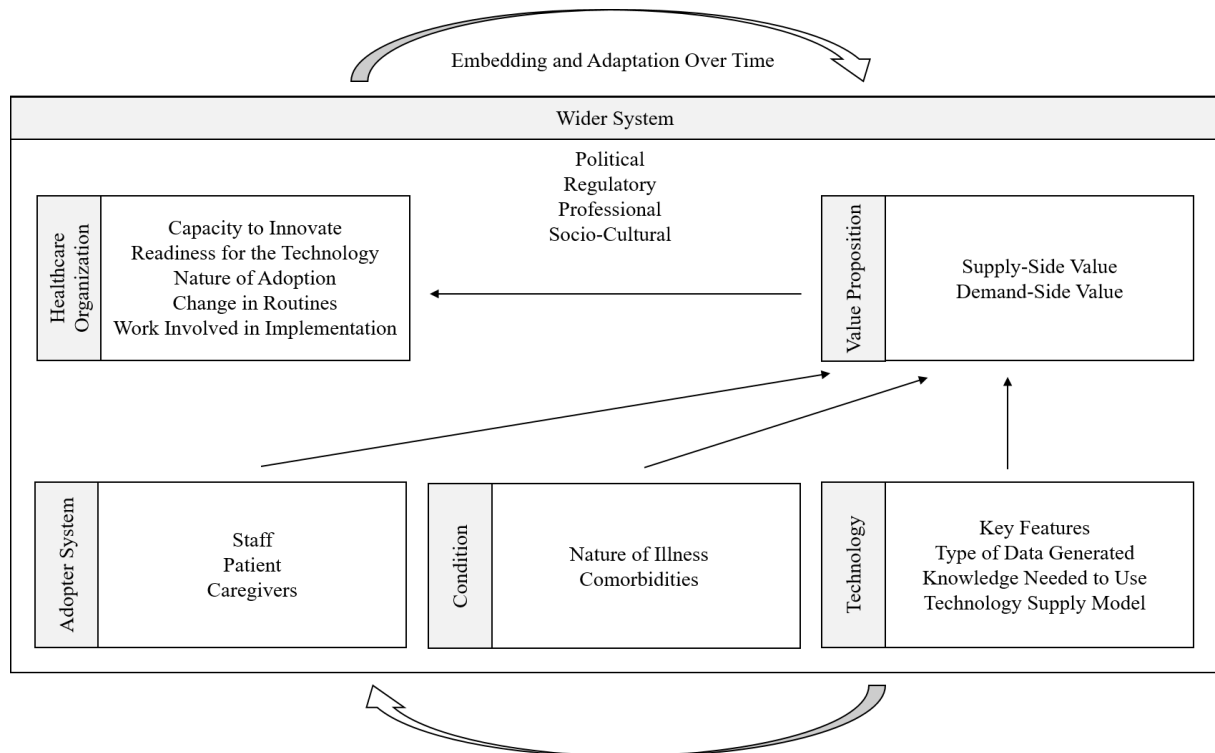


Figure 3: Overview of the NASSS Framework (based on Greenhalgh et al. 2017)

As regards content, the TOE and the NASSS frameworks partially intersect. The dimensions *technological*, *organizational*, and *environmental context* of the TOE framework are analogous

to *technology*, *healthcare organization*, and the *wider system*, even though the dimensions differ in terms of factors. The NASSS framework additionally includes further dimensions that are particularly relevant to healthcare organizations, such as the patient's condition. Despite these differences, which arise from the varying scope of the frameworks, both take a comprehensive perspective and consider the influence of different dimensions on the adoption of innovations in organizations (DePietro et al. 1990; Greenhalgh et al. 2017).

In line with the socio-technical approach, IS research does not only look at adoption from an organizational standpoint, but frequently studies individual intentions, behavior, and the decision-making of potential users of a technological innovation within and around the organization (e.g., Damanpour and Schneider 2009; Marabelli et al. 2021; Wischnevsky et al. 2011). This provides a new level of analysis that can shed further light on the adoption of technologies in organizations. For example, by analyzing individual intentions and behavior, the dimension adopter system of the NASSS framework can be examined in more detail to explain why a particular technological innovation may or may not be accepted by the medical staff, patients, or carers (Greenhalgh et al. 2017). Typical theoretical approaches that help explain the intentions and behavior of individuals are for instance the *theory of planned behavior* (TPB; Ajzen 1991), the related *unified theory of acceptance and use of technology* (UTAUT), which has been developed to predict the behavioral intention of an organization's employees to use a technology in a professional context (Venkatesh et al. 2003), and UTAUT2, an extension of UTAUT that moves the prior theory from the organizational setting to the consumer context (Venkatesh et al. 2012). Thus, these theories can be used to predict technology acceptance among organizational employees within the organization, as well as acceptance among customers (or patients) who obtain products and services from an organization. Aside from theories that investigate the acceptance of potential users, the privacy calculus, for example, can be applied to study individuals' intention to disclose personal information. The privacy calculus states that individuals weigh contrary beliefs in a process to decide whether to disclose their information (Culnan and Armstrong 1999; Dinev and Hart 2006; Krasnova et al. 2010). Such a theoretical perspective can be applied to exploring, at a detailed level, how organizations need to approach individuals to influence their decisions about disclosing data, thereby expanding an organization's database. Such a capability could be crucial for the introduction of data-driven technological innovations.

By applying the described theoretical approaches at the organizational and individual levels, the factors that influence successful organizational adoption of technological innovations can

be better understood. Harnessing the drivers and overcoming the barriers, technological innovations become deeply embedded within the organization. As a result, workflows may be performed more effectively and/or efficiently than without the technological innovation, for example, because cost advantages can be achieved or decision errors may be reduced (Porter 1985). Moreover, the introduction of strategically relevant technological innovations can also lead to central processes or even the entire business of the organization being transformed (Zhu and Kraemer 2005).

ML systems are among the technologies that have the potential to affect the core of an organization. Their adoption and organizational impact differ significantly from those of previous conventional technologies, given their specific characteristics. ML systems may recommend erroneous strategies, are in many cases inscrutable even to their developers, and are making more and more inroads into domains that were previously reserved for humans (e.g., Amershi et al. 2019; Chen et al. 2018; Domingos 2012; Jammal et al. 2020). As a result, IS researchers like Benbya et al. (2021) and Kane et al. (2021), and healthcare scientists such as Shaw et al. (2019) are calling for a more detailed investigation of the influence of ML systems on organizations in general, and on healthcare organizations specifically. Frameworks such as the TOE or NASSS can provide a first starting point to identifying the factors influencing the adoption of ML systems in organizations, but need to be adapted and extended to address the specific characteristics of ML systems (Davison and Martinsons 2015; DePietro et al. 1990; Greenhalgh et al. 2017). In addition to *what* needs to be done, the implementation of the identified factors necessitates a more detailed investigation of *how* the prerequisites can be met – a task that requires the involvement of further theories such as UTAUT2 (Venkatesh et al. 2012). Beyond examining the drivers of and barriers to the adoption of ML systems, the impact of adopting these systems into organizations is worth investigating as well. This is because, alongside humans, ML systems are now capable of learning on their own, leading to far-reaching consequences for organizations (Argote et al. 2021; Mitchell 1997; Ransbotham et al. 2020). For instance, Argote et al. (2021) suggest further investigation into ML systems' impact on organizational learning, a process vital to organizational survival that has previously been performed exclusively by humans (e.g., Bushee 1998; Cohen and Levinthal 1989). This is because ML systems, now able to deliver novel insights into the organization, could significantly affect the fragile balance between exploring new knowledge and leveraging existing expertise within the organization (e.g., Lavie et al. 2010; March 1991). Especially since humans tend to rely on pre-existing expertise instead of constantly acquiring new, cutting-edge

knowledge – a notion known as learning myopia – ML systems could hold promise for organizational learning (Levinthal and March 1993).

IS researchers have recently begun to discuss the characteristics of ML systems and their influence on the adoption of these systems in organizations and among their members. While Giermndl et al. (2021) identify the lack of a consistent, integrated, and cleansed database as an obstacle to ML systems' adoption, Asatiani et al. (2021), for example, emphasize the influence that the intransparency of ML systems can have on the organization and how potential negative effects can be mitigated. Lebovitz et al. (2021), on the other hand, investigate the impact of ML systems' fault-proneness on the integration of these systems into organizations. However, IS research on the adoption of ML systems is still scarce. Similarly, there are only a few articles in healthcare research that contribute to understanding the adoption of ML in healthcare organizations, many of which are merely viewpoints, such as Alami et al. (2021), Briganti and Le Moine (2020), or He et al. (2020). However, especially in healthcare organizations, the adoption of ML systems – despite their promises – can be challenging due to the high cost of misjudgments, for example in diagnostics (e.g., Lebovitz 2019; Lebovitz et al. 2021). Therefore, an in-depth analysis of the factors influencing the adoption of ML systems in organizations, and especially in healthcare organizations, as well as the resulting outcomes, is urgently needed.

3 Paper 1.A: A New Organizational Chassis for Artificial Intelligence – Exploring Organizational Readiness Factors

Title

A New Organizational Chassis for Artificial Intelligence – Exploring Organizational Readiness Factors

Authors

Pumplun, Luisa; Tauchert, Christoph; Heidt, Margareta

Publication Outlet

European Conference on Information Systems (ECIS)

Abstract

In 2018, investments in AI rapidly increased by over 50 percent compared to the previous year and reached 19.1 billion USD. However, little is known about the necessary AI-specific requirements or readiness factors to ensure a successful organizational implementation of this technological innovation. Additionally, extant IS research has largely overlooked the possible strategic impact on processes, structures, and management of AI investments. Drawing on TOE framework, different factors are identified and then validated conducting 12 expert interviews with 14 interviewees regarding their applicability on the adoption process of artificial intelligence. The results strongly suggest that the general TOE framework, which has been applied to other technologies such as cloud computing, needs to be revisited and extended to be used in this specific context. Exemplary, new factors emerged which include data – in particular, availability, quality and protection of data – as well as regulatory issues arising from the newly introduced GDPR. Our study thus provides an expanded TOE framework adapted to the specific requirements of artificial intelligence adoption as well as 12 propositions regarding the particular effects of the suggested factors, which could serve as a basis for future AI adoption research and guide managerial decision-making.

Keywords

Artificial Intelligence, Adoption, TOE Framework, Organizational Readiness

3.1 Introduction

“The world’s most valuable resource is no longer oil, but data” – proclaimed by The Economist (2017) and a plethora of other articles, the business value of data is widely accepted. If data is the new oil of our economy and artificial intelligence (AI) is fuelled by data, then AI can analogously be referred to as the engine (Agrawal et al. 2018). Thanks to improved algorithms in deep learning and ample access to historical datasets as well as cost-effective computing power and storage space, AI applications are on the rise and receive increasing attention from both technology companies and more ‘traditional’ companies that anticipate competitive advantages (MSV 2018). Despite inconspicuous short term impact, long term commitment is important since AI represents a paradigm shift for organizations (Hosanagar and Saxena 2017). According to Gartner, “85 percent of CIOs will be piloting AI programs through a combination of buy, build, and outsource efforts” by 2020 (Brant et al. 2017, p. 2) – however, just like a new engine for electric vehicles requires a new chassis, approaching an organizational AI project requires an assessment whether the focal organization possesses the necessary prerequisites and framework to enable successful AI initiatives.

Despite ever increasing organizational (and governmental) investments in AI (Bughin et al. 2017), less than 39 percent of all companies have an AI strategy in place, only 20 percent of companies have actually incorporated AI in some offerings or processes, and merely 5 percent have extensively incorporated AI (Ransbotham et al. 2017). The easiest explanation for this apparent hesitance are prominent examples of AI projects gone awry, like the Microsoft Chatbot Tay tweeting racist slurs (Reese 2016) or IBM’s Watson failing to diagnose cancer as promised in their advertising campaign (Flam 2018). However, most so-called AI failures cannot be attributed to AI itself but rather to the underlying processes and the involved people. Current AI research has focused predominantly on technical advancements (e.g., Lu et al. 2018; Monroe 2018) but largely factored out the necessity to analyse the readiness of the ‘organizational chassis’ to successfully support AI initiatives. In this regard, AI initiatives cannot be approached like yet another new technology trend since several aspects distinguish these projects from previous technology initiatives, e.g., cloud computing adoption or social media marketing: in its essence, AI refers to a broad and complex set of approaches that do not have to confine themselves to methods that are observable and have thus been often compared to a black box (McCarthy 2007). In accordance with McCarthy (2007, p. 2) we understand AI as a “science and engineering of making intelligent machines, especially intelligent computer programs”, which tries but is not limited to simulate human intelligence and which includes underlying technologies like machine learning, deep learning and natural language processing

(Elliot and Andrews 2017). AI differs from non-AI technology as it learns to make decision based on incoming data, rather than being based on an explicitly defined set of rules (Crowston and Bolici 2019). This self-adaptive property allows AI to learn from user behaviour, react to its environment, and make complex decisions automatically. These properties result in human attributes being assigned to AI (Rzepka and Berger 2018). However, the technology is also perceived as a threat because the algorithm's decision is not transparent (i.e., black box behaviour) and is likely exceeding human capabilities in a particular task due to its efficiency and scalability (Brundage et al. 2018).

In an information systems (IS) context, researchers have only recently begun to examine organizational readiness factors for AI (e.g., Alsheibani et al. 2018) but have as of now not yet expanded frameworks like TOE (technological-organizational-environmental) to cover the specific characteristics AI initiatives entail across industries and adoption stages. Due to the scarce extant literature, this study explores organizational readiness factors through a qualitative interview approach with 14 experts from both user and provider firms at various adoption stages. Building on TOE as conceptual framework, our approach thus aims to identify:

RQ1: Which factors influence the decision and the ability to adopt AI in organizations?

And sets out to shed further light onto

RQ2: What explicitly distinguishes the introduction of AI from other technologies?

The remainder of this manuscript is structured as follows: To begin with, we provide a brief overview of the related work and theoretical background (TOE) to mark off the research area before the qualitative study design is presented. After introducing our study sample comprising 14 interviewees, we derive empirical results which are integrated to expand the TOE framework. The results of our paper are a first step in providing a holistic view of the factors that are relevant for adoption of AI in the nascent research landscape. Thereby, the discussion of our key findings illustrates contributions to research and practice and an approach to future work. Finally, we conclude the manuscript by pointing out the limitations of our study and providing specific avenues for future research.

3.2 Theoretical and Conceptual Background

3.2.1 Artificial Intelligence and Adoption

The nascent ubiquitous adoption of AI in companies is currently omnipresent in research and practice which indicates the potential attributed to AI. However, only few studies have dealt

with the organizational aspects of AI adoption like the implementation of the technology into organizational processes and governance structures (e.g., Ransbotham et al. 2017). Extant published studies rather focus on the improvement of this technology and its underlying algorithms (e.g., Monroe 2018; Yan et al. 2016) or the impact of AI on specific industries and departments (e.g., Huang and Rust 2018; Kruse et al. 2019; Moncrief 2017) – whereas overarching aspects like the influence on AI applications exerted by an organization's strategy or the macro-environment, have scarcely been taken into account in information systems (IS) literature (Nascimento et al. 2018).

Indeed, a literature review by Nascimento et al. (2018) demonstrates possible avenues for future studies by identifying specific aspects which should be considered when adopting AI technologies (i.e., high commitment to the area, human requirements to deal with the techniques), but they do not integrate their findings into a theoretical framework. Similarly, Rzepka and Berger (2018) focus on the interaction of AI systems and users and address important factors (e.g., the fit between the user, system and task), but do not apply a distinct adoption framework. There are some further, rather practice-oriented contributions analysing or discussing the adoption of AI. For example, vom Brocke et al. (2018) state that new job profiles have to be created, resulting in the necessity of adequate skill development of employees and the adjustment of corporate strategies.

However, the aforementioned findings are still rather disparate and do not provide a concise framework that could guide future organizational studies regarding AI and the actual implementation of AI in companies. To the best of our knowledge, there are only two contributions that consider the adoption of AI in organizations from a more theoretical perspective and across various industries (Alsheibani et al. 2018; Rana et al. 2014). Alsheibani et al. (2018), a research-in-progress publication, draw on the TOE framework (DePietro et al. 1990) to explain an organization's readiness to introduce AI into their organization. In line with the existing theory, they constitute technological (T), organizational (O), and environmental (E) factors which influence AI adoption and propose a quantitative, thus confirmative, approach. Accordingly, influencing factors are selected on the basis of assumptions from past studies, which are not specified in more detail, and on the basis of previous technologies, which do not have the same specific characteristics as AI. Rana et al. (2014), on the other hand, use the Technology Acceptance Model (TAM) to explain the organizational adoption of machine learning techniques in the specific context of software defect prediction. Again, the unique characteristics of AI are not sufficiently addressed. Instead, existing concepts (e.g., perceived

benefits) are examined based on a sample of only four interviewees from two companies. Given that AI differs from previous technologies in several ways, an all-embracing framework needs to take these differences into account (Zhu and Kraemer 2005): AI is considered both efficient and scalable, is able to exceed human capabilities and comprehension (Brundage et al. 2018), derives its own rules from added data (Crowston and Bolici 2019) and shows a distinctive black box behaviour (Adadi and Berrada 2018). In addition, recent developments affect the organizational use of AI (e.g., improvement of deep learning algorithms) making it necessary to collect comprehensive, up-to-date data.

Since no current exploratory study investigates the adoption of AI across various industries, an explorative approach is necessary to provide further insights that potentially deepen and extend the proposed TOE framework to account for the novelty regarding the organizational implementation and adoption of AI.

3.2.2 TOE Framework and Diffusion of Innovation

In general, the TOE framework represents a useful and somewhat flexible starting point to study innovations as it provides a generic theory for the diffusion of technologies (Zhu and Kraemer 2005). Therefore, it has been widely applied to other contexts and technologies like cloud computing (e.g., Lian et al. 2014), big data (e.g., Bremser 2018) and business intelligence systems (e.g., Hatta et al. 2017). In essence, the TOE framework comprises three main elements that influence the adoption process of technological innovations: (a) the technological context describing the internal and external relevant technologies available, (b) the organizational context that depends on internal structures and processes measured by various factors such as company size and free resources and (c) the environmental context, which describes the business related field of action, taking into account industry, competitors, government, and suppliers (DePietro et al. 1990). Following Zhu and Kraemer (2005) the TOE framework can be extended by using the innovation diffusion theory of Rogers (1995), which states different technological factors including relative advantage and compatibility. Relative advantage is described as the degree to which an organization perceives an innovation better compared to the previous solution. The second factor, compatibility, is the degree to which an innovation matches the actual needs of the potential user organization. Both factors are positively related to its rate of adoption (Rogers 1995). Looking at the organizational readiness, DePietro et al. (1990) postulate a positive influence of the strategic behaviour of management, organization's size and slack resources. They also point out the relevance of the intensity of competition as a positive factor on adoption as well as governmental regulations, which can have both, negative

and positive effects on innovation implementation. Since there is only little research on AI adoption, a general TOE framework as described above is used as an initial conceptual starting point (see Figure 4), which will be expanded in the course of the study.

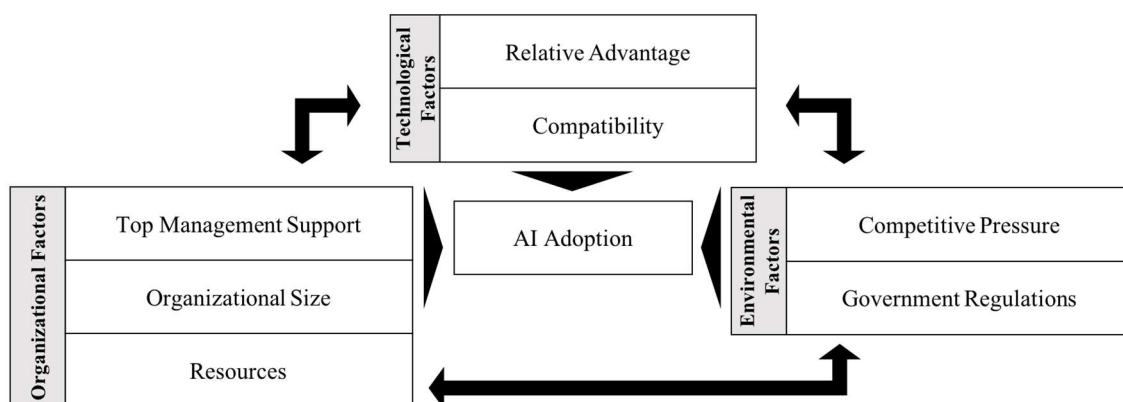


Figure 4: TOE Framework as Conceptual Base (based on DePietro et al. 1990; Rogers 1995)

3.3 Qualitative Research Methodology

The aim of the study is to expand the current state of IS research concerning AI application in organizations by questioning experts who work on managerial and operational levels for AI provider and user firms. Organizational AI adoption is a complex topic and has not yet been fully explored. Therefore, an explorative approach using interviews with experts seems appropriate to investigate the problems occurring in this particular context (Flick et al. 2004). According to Weber (1990), content analysis can be used to assess open-ended questions, making the approach suitable for evaluation of the collected qualitative data. Thus, in order to develop an organizational adoption framework, this paper follows the steps of content analysis (see Figure 5): Based on the TOE framework, which serves as a conceptual framework, seven initial categories were derived from relevant literature (e.g., factors “compatibility” or “top management support” in Figure 4). By analysing the interviews, these categories are examined and extended gradually, resulting in 23 categories and subcategories of the final framework for AI adoption. The interviews are transcribed, coded and analysed taking into account relevant practice-oriented studies through triangulation (Hsieh and Shannon 2005). In particular, we use a combination of directed and conventional analysis, where the directed approach uses codes derived from theory (i.e., TOE framework) and the conventional analysis takes into account information obtained directly from the data since the applied theory is not specifically adjusted to AI technology and therefore should be supplemented and deepened inductively (Hsieh and Shannon 2005).

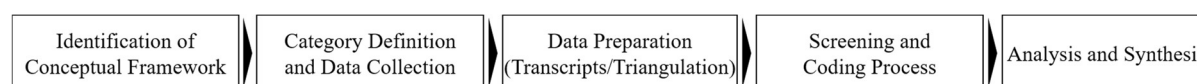


Figure 5: Content Analysis Process (based on Hsieh and Shannon 2005)

3.3.1 *Research Design*

Our main information source were in-depth expert interviews, which were conducted in a semi-structured way. Thereby, the guiding principles of Sarker et al. (2013) were considered by preparing an interview protocol and questioning key informants in different companies. In order to avoid typical pitfalls of semi-structured qualitative interviews, contact was established with the interview partners via e-mail and telephone before the interviews were carried out (Hermanns 2004). While conducting the interviews we kept our questions open in order to enable participants to speak freely.

The interview guide comprises three different sections. The first section comprised general questions about the position and responsibility of the interviewee and their previous experience in the field of AI and related technologies in an operational or managerial context. The second and most comprehensive section considered advantages and risks of using AI (i.e., the possible results of AI initiatives) and the triggers, prerequisites and limitations of using this technology in organizations. In addition, we inquired the criteria used by the companies to assess the general potential of AI. The last set of questions dealt with the actual use of AI and the strategic and tactical challenges it poses. For example, we asked the interview partners which AI-based applications are currently being used and which specific actions were associated with the introduction and implementation of these projects. Due to the semi-structured approach, initial questions were subject to a gradual adjustment in order to account for the individual expertise and position of the participants and to develop the focus during the interviewing process.

3.3.2 *Sample and Data Collection*

We provide an overview of the participants in Table 2 (see below) and further details in the following.

Table 2: Participant Overview

Participants (UF): Participants of firms that are predominantly users of AI products and services					Participants (PF): Participants of firms that are predominantly providers of AI products and services				
ID	Position	Job Exp.	Interview Method	Adoption Stage	ID	Position	Job Exp.	Interview Method	Core/Non-core
P-01	Digital Growth Manager	16 years	Face-to-face	Adoption	P-08	Founder	10 years	Face-to-face	C
P-02	Head of Marketing & Analytics	10 years	Face-to-face	Consideration	P-09	Development Manager	6 years	Face-to-face	C
P-03	Head of Digital Communications	14 years			P-10	Solution Manager	15 years		
P-04	Asset Management Strategist	3 years	Telephone / Face-to-face	Adoption	P-11	Development Manager	7 years	Face-to-face	C
P-05	Chief Product Owner	8 years	Face-to-face	Continued use	P-12	Managing Director	19 years	Written answer	NC
P-06	Product owner	8 years	Face-to-face	Continued use	P-13	Consultant	2 years	Telephone	C
P-07	Account Executive	3 years	Telephone	Adoption	P-14	Managing Director	11 years	Telephone	C
Awareness: Org. becomes aware of AI Consideration: Org. considers to adopt AI Intention: Org. intends to adopt AI Adoption: Org. begins to adopt AI Continued use: Org. continues to use AI					Core (C): AI capabilities and products differentiate company strategically from others Non-Core (NC): AI capabilities and products are no strategic factor for company				

The interview partners were selected on the basis of a key informant approach. Following the rules of data triangulation, both user (UF) and provider firms (PF) were surveyed (Flick 2004). The answers were collected over a six-month period and took place between May and October 2018. In total 12 interviews with 14 highly involved participants were conducted within two European countries (Germany and Ireland), taking into account seven experts from provider firms and seven experts of companies, which mainly purchase AI products. After the 12th interview, data collection was discontinued as a further contribution of additional qualitative data was considered unlikely (i.e., theoretical saturation was assumed; Flick 2004).

Among the 14 interviewees were eleven male and three female participants. The total number of respondents is comparable to other qualitative studies that consider the adoption of similar technologies (e.g., Bremser et al. 2017; Labres Mallmann and Gastaud Maçada 2018). In order

to avoid an elite bias, both IT staff and managers were interviewed (Miles and Huberman 1994). Therefore, three of the participants were managing directors or founders, eight identified as middle managers or heads of departments, while the remaining respondents were either consultants or strategists. For the purpose of potentially achieving more generalizable research results and identifying sector and enterprise size-specific differences (Flick 2004), companies across various industries and of differing sizes were selected, including large (75 %), medium-sized (17 %) and very small enterprises (8 %; European Commission 2003) from industries like electricity, gas, steam and air conditioning supply (D), information and communication (J), manufacturing (C) as well as wholesale and retail trade (G; United Nations 2008). At the time of the interviews, the organizations were in different phases of implementation regarding AI. Based on the classification according to Frambach and Schillewaert (2002), user firms are divided into the following stages of adoption: awareness, consideration, intention, adoption and continued use, while provider firms were classified according whether they offered AI as a core competence or not (Leonard-Barton 1992).

The interviews lasted on average 58 minutes and were mainly held face-to-face because of the complexity, scope, and sensitivity of the topic. Nevertheless, a total of four interviews were conducted using telephone calls and one participant replied in a written form due to geographical distance. An overview of the surveyed participants can be found in the table above (see Table 2).

3.3.3 *Coding Concept*

Most of the interviews were recorded and transcribed after agreement by the interviewees. In a single interview only notes were taken and in another case a written answer was submitted. Subsequently, the transcripts were assessed by using the NVivo 12 software and by conducting two coding cycles as recommended in Saldaña (2009). The first coding cycle comprised a mixture of attribute coding, descriptive coding and hypothesis coding. The former is performed to obtain essential insights about the data and its descriptive information (e.g., UF/PF, size of organizations). In addition, hypothesis coding was carried out to account for the initially conceptualized factors from the TOE framework (see Figure 4). These factors mentioned in the existing theory form the focus of the hypothesis-based approach and are deductively tested (Greener 2008). Finally, descriptive coding is used to extract additional aspects that go beyond the previously identified factors (e.g., relative advantage, competitive pressure, and top management support) and thus potentially extend the existing framework. In a second cycle, the formerly created codes are combined into a smaller number of sets using pattern coding

(Saldaña 2009). By discussing and assessing the coding process with a group of four IS researchers and students, an investigator triangulation helped to ensure rigor and trustworthiness. Furthermore, an ongoing data triangulation process took place while coding the interviews by utilizing multiple sources of evidence (Flick 2004). For example, additional corporate resources as well as current practice-oriented AI studies and reports were considered (e.g., Brant et al. 2017; Ransbotham et al. 2017).

3.4 Results and Discussion

While validating the proposed TOE framework for adoption of AI (see Figure 4), we found evidence that the established factors do not fully reflect the challenges that companies face when they want to introduce AI to their companies. The presented TOE framework merely includes fundamental factors that are also applicable to other technologies such as cloud computing. Therefore, the findings that do not go beyond these basics are summarized in tabular form (Table 3). Aspects that supplement or contextualize the original framework will be examined in more detail below.

In addition to the ‘classic’ TOE assumptions, the experts also mention prerequisites for the implementation of AI that result from the special properties of AI and therefore have only been insufficiently addressed or have not been examined in general TOE literature at all before. These new findings are described comprehensively in the following section.

Table 3: Findings: Examination of Proposed Factors in TOE Framework

El.	Fact.	Results	Statements
Technological Factors	Relative Advantage	With the help of AI it is possible to learn from the data over time. However, AI is not a panacea, but should be compared to the use of robust conventional systems for the specific application. The combination of both approaches should also be considered in order to solve the overall problem. This assumption is strengthened by Rzepka and Berger (2018), who indicate that AI is better suited for particular use cases than others. In addition, it is demanded that the results of AI be made comprehensible and no longer represent a black box. The demand for more transparency of AI based systems is also demanded in the current IS literature (e.g., Crowston and Bolici 2019; Rzepka and Berger 2018).	<p>“But that one adapts, that one learns based on collective knowledge, no matter if one provides it now at the beginning or continuously, that one adapts there then, that is actually the strength of this AI.” – P-11</p> <p>“And it may well be that you get on with workflows or get on with fixed processes. Or that you say, you know what, we just run AI in the background. And we just take a look at which needles the system still brings us. But it’s by no means a panacea [...]” – P-01</p> <p>“We know that we can’t really understand machine learning. [...] And that there must be procedures that show that exactly this one feature was responsible for it.” – P-13</p>
	Compatibility	For the successful use of AI, the work processes must be adapted to the technological requirements. Furthermore, there must be a fit between the desired application and technology. ^a	“If I then ask [...] why do the projects fail? You then realize that the need was not clearly communicated, the use case was not right, that it was too big. That you say you want to do something, but you don’t know what.” – P-01
Organizational Factors	Top Management Support	In principle, the support of top management can facilitate the introduction of AI. However, a certain understanding of the technology and its applications is required. Currently, decision makers in middle management are particularly problematic, as they are very KPI-driven and thus inhibit AI use.	“Someone, a top manager or someone comes from some conference, has picked up something like Big Data or Predictive Maintenance as buzzwords and then says, ‘yes, let’s do it’. Yes? And then you start to code somehow and you start to collect and somehow you notice then hey, actually we don’t know exactly what we are supposed to do now.” – P-14
	Organizational Size	It is unclear whether larger companies have a better chance of adopting AI. Basically, a high budget and a large volume of customer data enables and justifies the use of AI. However, the slow group structures are also hampering further development in this area.	“Now are you going to [...] I’d rather say a niche area. Niche in the sense of, you have maybe only 10,000 users. Then it’s not worth the effort that data scientists, Computational Linguists develop something for five years.” – P-11
	Resources	The resources can be divided into the factors budget, employees and data that affect the use of AI. ^a	“I think obstacles [...] are certainly the initial expenditures. At the beginning, you’d need a small one-off budget, a bit of know-how as a starting point [...]” – P-02
Environmental Factors	Comp. Pressure	Competitive pressure leads companies to increasingly deal with AI in order to gain a competitive advantage.	“They [the costumers] challenge us too. They say, look at the competition, the start-up does that, we’ve already looked with them. Why can’t you do that yet?” – P-10
	Gov. Regulations	Many laws complicate the introduction and use of AI. In this context a renewal of the legal situation is demanded. Especially the GDPR and the employees’ council are a particular hurdle for companies. ^a	“And innovation and law are two words that I think rarely appear in one single sentence.” – P-04

^a: Further details on the subcategories are discussed below

Technological Factors

Technological factors comprise two main aspects: Relative advantage, which was already considered in Table 3 in detail, and compatibility, which can be divided into two subcategories on the basis of expert interviews: business processes and business cases. Therefore, we will revisit the second factor compatibility in the following and explain it in more detail.

Compatibility. According to experts, the *business processes* in the company must be adapted to the new requirements that arise from the use of AI. In the context of AI, it is therefore no longer useful to use existing KPIs of other projects, since AI projects have differing properties. For example, the results that arise from such projects can no longer be planned to an extent that would be necessary regarding traditional, common KPIs (e.g., ROI) as demonstrated by the following quote:

“The interesting thing about how we implement these projects here is that we didn’t define KPIs [...]. That means for us, we learn with the information we get back through the system. That’s a very important point. If you apply old KPIs to new technologies and approaches, you run the risk of only digitizing old KPIs.” – P-01

Instead, it becomes necessary to introduce agile forms of work. Particularly in the field of data science, it is important to continuously evaluate the progress of projects, since the feasibility of ideas in this area cannot be proven from the outset. There are only a few, incomplete criteria to evaluate the existing data at the very beginning. Within the framework of agile, flexible working models for software development, the current status and the data can always be viewed in terms of new findings, thus reducing the risk of investing the wrong amount of time and money. The relevance of agile working methods is underlined by the following statement:

“And in IT you had very, rigid waterfalls, that is classic traditional IT project management. Which is not, how shall I say, very beneficial regarding the uncertainties when using data and artificial intelligence. [...] Because you just plan a concept somehow, that’s actually this classic process, over half a year and then you look into the data and notice ‘oh God, that’s all wrong!’. And you can actually throw the concept away! So half a year, more or less, not as much progress has been made as if one had looked at the data in advance.” – P-14

In addition to the work processes, however, further factors must also be checked for compatibility. Another very frequently mentioned aspect is the formulation of a concrete *business case*. Experts believe that AI can only be used successfully if there is a clear problem.

AI must be seen as a tool for a purpose and cannot be viewed in isolation. The problem of prioritizing possible use cases appropriately is known from literature on big data use (e.g., Bremser 2018) which also deals with an underlying technology that can be used in a variety of ways in organizations.

“But you really need to know, ‘where can you solve a problem with that?’. Just because you can do AI, it doesn’t bring you anything, zero, honestly not. [...] They don’t buy it because it’s AI. So really, also corporate customers, they don’t buy it because there is AI in it now. They buy it because it must have a benefit.” – P-08

In line with these factors influencing AI adoption, we formulate the following propositions:

Proposition 1: *Compatibility between AI technology and business processes (e.g., agile forms of work) as well as the development of a dedicated business case will have a positive effect on adoption of AI in companies*

Organizational Factors

In addition to technological readiness, factors must also be taken into account that reflect the overall organization’s ability to implement AI. The factors culture and organizational structure were newly discovered by examining the expert interviews, while the factor resources was subdivided into the aspects budget, employees, and data.

Culture. After evaluating the interviews, it became evident that the adoption of AI in a company is strongly influenced by the culture in the company. In addition to top management support the introduction and implementation of an innovative culture in the company are also relevant. In this context, aspects of *change management* to achieve an innovative culture within the company were mentioned frequently by the interviewees. The functionality of an intelligent application is based on the input of already existing, high-quality data as well as the training which has to be carried out by the employees over time (Crowston and Bolici 2019). Only if there is a willingness to use the technology in the long run, the quality of the answers and decisions made by the machine will improve.

“In the beginning the model is bad. You have few answers that reach this threshold. But by constantly saying as an employee that this was right or by correcting, you are building a knowledge base.” – P-08

If the path to an *innovative culture* is not successful, there is a danger of missing out on new, important technologies and trends. The factor of missing an absorptive capacity to adopt new technologies is evidenced by the following statement:

“In such a large corporation you have the tendency to say again and again ‘well, we make money with the model we have! Why should I come up with something new now?’.” – P-05

Resources. The adoption of AI in a company does not only depend on the culture, but also results from slack resources, which should be further subdivided. Comparable to other innovations (e.g., Bremser et al. 2017), the available financial resources through a *budget* are an important aspect that generally determines the implementation of new technologies in projects. A high budget can enable capacities, create financial freedom and help to build know-how. On the other hand, obligations also arise from financial resources. This problem can in turn jeopardize the successful introduction of AI, since the course of projects with AI is unpredictable and strongly dependent on the data used. The restricting influence of budget is demonstrated by the following statement:

“The second point is the budget. The moment your management or the person responsible for the budget asks the question ‘what is the return on investment?’. And ‘what happens if I don’t do it?’ You are no longer on the move agilely, but you are immediately arrested in a major project. The demand or the requirements are already defined, there’s a price tag on it and there’s a timeline on it. No more room for adjustments.” – P-01

In addition to the budget, a second aspect should be considered as one of the most frequently discussed factors within the sample: the *employees* of a company who have the necessary know-how to apply the technology. Basically, it should be noted that the staff should have both, the professional qualifications and programming knowledge in the field of AI (e.g., utilizing libraries such as TensorFlow, PyTorch or Keras) as well as a domain understanding of the respective organization. It should also be considered that many companies have problems recruiting professionals such as data scientists, who demand high salaries and are potentially disloyal to their employers due to a high demand on the labour market. The necessity of these occupational groups for implementation of AI is also addressed by previous studies (e.g., Kruse et al. 2019). Additionally, interviews show that AI projects cannot simply be outsourced as they require the company’s domain knowledge as described. Therefore, an expert suggests to train the employees in the company who already have a domain specific knowledge (e.g., controller, statisticians) in the field of machine learning. The problem set is evidenced by the following statement:

“This is one of the most important things: you need the people! In this day and age you can no longer outsource. Especially not with machine learning and artificial intelligence. That doesn’t work. You need the experts. You need the people – who actually don’t have the time.” – P-01

The third subcategory that can be seen as a resource is the *data* used to train the AI. Data was among the factors most often mentioned by all interviewees across firms and positions and is also frequently considered in current literature (e.g., Crowston and Bolici 2019). Various problems have been extracted while examining the qualitative interviews: Data must first be made accessible. Both *data availability* and data protection play an important role. Often the data must be made usable from different old systems. Furthermore, it is necessary to extract the data in a scalable form, because AI projects require as many data records as possible. According to the experts, these requirements can account for up to two thirds of the workload of an AI project. The following statement illustrates how time-consuming and difficult the provision of data can be:

“We also often [...] first had to think about ‘where does the data actually come from?’ [...] We actually had to deal with three or four different legacy systems from which we had to get the data out.” – P-05

In addition to the technical aspects of data availability, *data protection* also plays an important role. Often, it is mainly larger corporations that experience difficulties implementing an open data policy. In these kind of companies, a deliberate isolation of the individual departments takes place, which makes the successful introduction of AI more difficult:

“We’re going to have to make sure that we stop pursuing a silo mentality.” – P-01

Once the data is available, the quality of the data becomes relevant. This aspect was brought up very often by the interviewees, who point out that *data quality* is regularly a problem, as it is not fully possible to assess the data sets before the project is indeed implemented. Only a few incomplete metrics exist to evaluate the data in advance. This is particularly problematic because historical data often does not have the required quality and degree of detail due to time and cost pressure when data was generated.

“We also have customers who say yes, we have the CRM here, our system here, our old system. Maybe an old application. But we don’t really want to take the data with us, because we know that the service staff often just entered something hurriedly due to a lack of time, and that it’s not right.” – P-10

Organizational Structure. The culture of the company is closely linked to its structure. As in the statement above, large corporations have problems setting up new AI projects because of their “everything is fine” mentality. Many companies therefore go the way of circumventing old, inhibiting structures by establishing a lab or hub within the organization. However, problems can also arise as a result of this procedure, which is made clear by the following statement by an expert:

“Is this somehow a lab in Silicon Valley, where clever people are all sitting around building something without being subject to the restrictions of the traditional company? The advantage of this is that they are very fast. This has the disadvantage that the integration into the slow company will fail later. [...] On the other hand, if you try it out of the existing IT, which is historically very cost-driven and very innovation-free, then it won’t work either.” – P-14

Therefore, it is suggested to use a hybrid model, in which a hub serves as a starting point for new ideas and technologies, but where an intense communication between the lab and the company still exists.

As shown above, organizational readiness factors influence decisions regarding AI adoption of companies strongly. Hence, we posit:

Proposition 2: *A dedicated AI budget, which does not entail any obligations to meet performance targets, will have a positive impact on the adoption of AI in companies*

Proposition 3: *The availability of data scientists and developers with appropriate expertise, domain knowledge as well as the willingness of users to train AI systems over time will have a positive impact on the adoption of AI in companies*

Proposition 4: *The availability of extensive, meaningful and high quality data will have a positive effect on adoption of AI in companies*

Proposition 5: *Departments who keep relevant data to themselves, an overreliance on status quo as well as slow and bureaucratically shaped corporate structures will have a negative effect on the adoption of AI in companies*

Environmental Factors

Looking at environmental readiness, the known factor government regulations is divided into two main aspects (GDPR and employees’ council) and the categories industry requirements as

well as customer readiness are newly filtered out by coding the expert interviews. The extensions of the original framework are explained in more detail in the following section.

Government Regulations. As already indicated, the introduction of AI must also consider several legal aspects. A relevant regulation that was enforced in May 2018 is the *General Data Protection Regulation (GDPR)*, which regulates activities like the processing of personal data. The handling of the new legal situation is addressed by many experts in the interviews as companies struggle to provide personal data for the training of their intelligent machines. In this context, many data sets need to be anonymized, which makes the use of intelligent, self-learning algorithms more difficult or even impossible. The following statement expresses the impact that such a regulation can have on the European economy:

“This shock with the General Data Protection Regulation [...] to make everything bad per se and excessively laborious, that also contradicts any reality. Also, we have to be careful that we don’t lose track of others with all these AI topics, because they will do it. We would like to, but we’re getting a bit in ourselves’ way.” – P-11

In addition to legislation concerning the handling of personal customer data, the protection of employees must also be taken into account by firms. Many applications in the field of AI are based on learning from data. If intelligent software is used in the company to support employees, it can access a lot of information from their daily work routine. Thus, there is a danger that the personnel could be monitored. In addition, as a result of the progressive automation by AI, a large scope of duties is taken over gradually by machines. Although it was one of the less prominent constraints mentioned by all interviewees, these effects of intelligent algorithms ultimately lead to the fact that the introduction of AI is inhibited by *employees’ council and employee representatives* in companies to protect employees’ workplaces.

“Because, of course, a system of this kind, which logs data without limits, could of course also store the information. That X makes three mails in one day and Y makes 30. And her completion rates are much higher. Okay? So the employees’ council is definitely a key stakeholder.” – P-01

Industry Requirements. In addition, each industry has its own specific requirements, which also affect the adoption of AI. These are specific laws, external circumstances affecting the company, and the organization’s interaction with the environment. For example, Kruse et al. (2019) examine the adoption of AI in financial sector taking into account its specific regulations, IT systems and customer group. These influences can encourage or inhibit the use

of AI, depending on their nature. The necessary inclusion of the factor *industry* was evidenced, besides the related literature, by the following statement:

“I also believe that our industry [electricity provision] is simply making a bit of an impact. The challenges facing our industry are simply more complex than what a small retailer might have to solve [...]” – P-02

Customer Readiness. When a company is faced with the decision to introduce AI, the knowledge and acceptance of its customer base must also be taken into account. These requirements apply to B2B as well as to B2C companies, which should both focus on their customer’s benefit. The interviewed experts currently see a development of their customer’s ability and willingness to deal with new technologies. Consumers in particular are increasingly demanding digital and intelligent offers and are acting as disruptors. This is consistent with other adoption literature, which points to changing customer expectations for individualized services and products (e.g., Bremser 2018). But also corporate customers are beginning to innovate. The requirements they will have in the future can be seen from the following statement:

“In 3 to 4 years, when the algorithms are mature, this will become the standard. Then the customers simply expect that such a function [intelligent service] is in the solution.”
– P-10

We thus posit that environmental factors, like the legislation or the readiness of industry and customers, affects AI adoption as follows:

Proposition 6: *Strict laws regarding the processing of personal data will hamper the training of intelligent machines and the review by a strong employee representative body will slow down and inhibit the introduction of new technologies. Thereby both will have a negative effect on adoption of AI in companies*

Proposition 7: *Industry specific properties (e.g., specific regulations, customer group) will, depending on their nature, have both positive and negative effects on the adoption of AI in companies*

Proposition 8: *Demanding customers will nudge the companies to design individualized, intelligent products and thus will have a positive effect on the adoption of AI in companies*

The previous findings will be used in the following to supplement the basic framework (see Figure 4) and to generate an overview of the experts’ statements and thus the special features of AI (see Figure 6).

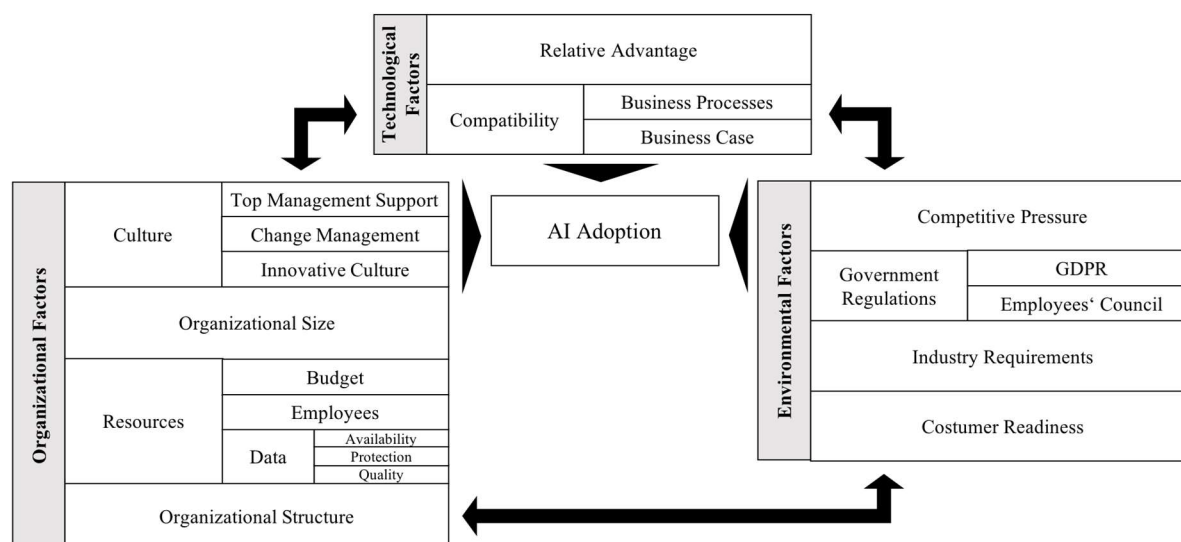


Figure 6: Extended and Deepened Framework for AI Adoption

After the proposed framework has been extensively investigated and extended, the next step is to showcase special features that occur during the introduction of AI in comparison to other technologies and which go beyond the theory of TOE. For this purpose, the statements of experts are investigated via crosstab queries (i.e., filter coded interviews simultaneously by a factor and company type) in order to get an idea about perceptual differences between provider and user firms, which eventually create a gap between supply and demand. The comparison inductively leads to different problem areas where the preconditions, views and attitudes of the provider and user firms differ.

An example of this misconception between those two groups is the differing assessment of consumers. While user firms tend to view their costumers as sceptical about the acceptance of intelligent applications, providers see consumers as disruptors who explicitly demand innovations.

“I believe we must not forget that our clientele is, to a large extent, rather conservative. And such a chatbot would not be suitable for everyone, not even for half of our target group.” – P-02 (UF)

“Very important, I have also become aware of this very often and very clearly, the customers are, as they say, the disruptors. They say exactly how they would like best to work with the brand.” – P-10 (PF)

But it is not only the customers that are assessed differently by the respective category of the firm. There is also a divergence of ideas about the prerequisites within the companies. For example, user firms see the size and bureaucracy of their group as an obstacle to the acceptance

of AI, while the provider designs their products primarily for large firms in mass markets which can generate sufficient amounts of data.

“Because it has been said that we do not see it within our existing group structures, we cannot give the issue the attention it needs.” – P-05 (UF)

“That especially companies that have many service requests benefit from this. [...] I also believe that, for medium-sized companies or something, I do not know. Especially larger companies.” – P-08 (PF)

The evaluation of the interviewees’ statements also shows that the ideas regarding the availability of budget for AI projects diverge. While large user firms state that they have problems providing the required financial resources, the provider firms overestimate the possibilities of their customers.

“I think obstacles, why we have not done it [AI adoption] yet, are certainly the initial expenditures. At the beginning where you would need a small one-off budget, a bit of know-how as a starting point, which might not be there yet.” – P-02 (UF)

“It’s also often the case that large corporations in particular have strategic investment pools, where even a CEO says ‘yes, I have understood that in order to do something there, we now have to take three, four million in to our hands and we’ll take that as play money and start making this initial investment’.” – P-14 (PF)

Another point mentioned by the provider firms is the preference of user firms regarding on premise versus cloud-based solutions. As a result, providers are often unable to train and adapt the intelligent algorithms adequately since access to data and sufficient computing power is constrained.

Considering the differences between user and provider firms, we posit the following propositions:

Proposition 9: *The diverging assessment of consumer’s AI readiness by provider and user firms leads to a different estimation of demand and thus will have a negative effect on adoption of AI in companies*

Proposition 10: *The fact that the companies that have sufficient data volumes and are addressed by provider firms are also trapped in slow structures of their corporations will have a negative effect on adoption of AI in companies*

Proposition 11: *Misconceptions about budget availability and willingness to pay between user and provider firms will have a negative effect on adoption of AI in companies*

Proposition 12: *Differing preferences of cloud-based and on premise applications between provider and user firms result in a negative effect on adoption of AI in companies*

3.5 Conclusion, Limitations, and Future Research

The explorative study showed that the TOE framework is applicable to the adoption of AI. However, some categories show results that are partially contradictory and require further research (e.g., organizational size). Furthermore, we were able to identify new, AI-specific factors (e.g., data) and subcategories for existing ones (e.g., GDPR and employees' council as part of government regulations). Moreover, evaluating the interviews allowed us to provide initial solution approaches to address the problems that could possibly arise while implementing AI. Altogether, a framework for the adoption of AI is proposed, which provides executives with a broad overview of AI related conditions in organizations. This enables companies to carry out a structured analysis of their status quo and identifying areas of improvements to adopt AI successfully in their processes and services. In addition, it is shown how a gap between supply and demand for AI technology can arise due to diverging assumptions of user and provider firms. In order to enable the top management to address this disagreement, it is necessary to expose them and to create the prerequisites needed for a successful implementation of AI in their company. Besides the practical implications, by conducting the first cross-industry exploratory study focusing on factors which enable and impede AI adoption in general, a basis for further research is introduced. This study can be seen as a starting point to conduct additional studies – for example focusing on or comparing special industries (e.g., healthcare, banking and finance) and associated requirements or looking at specific departments and use cases in depth (e.g., HR, Service).

Future research should consider a constitutive quantitative study, to review the given proposals and further examine existing inconsistencies within the factors. This will help to understand the factors' actual impact, making it possible to develop sound strategies and action plans for an integrated AI adoption. Moreover, a framework other than TOE might then be applied to better reflect the specific requirements of AI (e.g., conceptual framework of organizational innovation adoption by Frambach and Schillewaert 2002). In addition, companies across the globe and of various cultures, should be included in the research, although a semi-multinational context already exists due to the fact that the interviewed firms are operating in several countries.

Additionally, we have mainly considered large companies so far, as they currently already have dedicated positions for AI projects and could therefore be easily identified and contacted. However, future research should survey medium-sized and smaller companies, especially as contradictory results on the impact of company size were obtained in the study. Nevertheless, this study ultimately was able to conceptualize an ‘organizational chassis’ for the introduction of AI adoption that enables organizations to move forward in the field of AI.

4 Paper 1.B: Adoption of Machine Learning Systems for Medical Diagnostics in Clinics: Qualitative Interview Study

Title

Adoption of Machine Learning Systems for Medical Diagnostics in Clinics: Qualitative Interview Study

Authors

Pumplun, Luisa; Fecho, Mariska; Wahl, Nihal; Peters, Felix; Buxmann, Peter

Publication Outlet

Journal of Medical Internet Research

Abstract

Background: Recently, machine learning (ML) has been transforming our daily lives by enabling intelligent voice assistants, personalized support for purchase decisions, and efficient credit card fraud detection. Besides its everyday applications, ML holds the potential to improve medicine as well, especially with regard to diagnostics in clinics. In a world characterized by population growth, demographic change, and the global COVID-19 pandemic, ML systems offer the opportunity to make diagnostics more effective and efficient, leading to a high interest of clinics in such systems. However, despite the high potential of ML, only a few ML systems have been deployed in clinics yet, as their adoption process differs significantly from the integration of prior health information technologies, given the specific characteristics of ML.

Objective: This study aims to explore the factors that influence the adoption process of ML systems for medical diagnostics in clinics to foster the adoption of these systems in clinics. Furthermore, this study provides insight into how these factors can be used to determine the ML maturity score of clinics, which can be applied by practitioners to measure the clinic status quo in the adoption process of ML systems.

Methods: To gain more insight into the adoption process of ML systems for medical diagnostics in clinics, we conducted a qualitative study by interviewing 22 selected medical experts from clinics and their suppliers with profound knowledge in the field of ML. We used a semistructured interview guideline, asked open-ended questions, and transcribed the interviews verbatim. To analyze the transcripts, we used a content analysis approach based on the health care-specific framework of nonadoption, abandonment, scale-up, spread, and sustainability in

the first step. In the second step, we drew on the results of the content analysis to create a maturity model for ML adoption in clinics according to an established development process.

Results: With the help of the interviews, we were able to identify 13 ML-specific factors that influence the adoption process of ML systems in clinics. We categorized these factors according to seven domains that form a holistic ML adoption framework for clinics. In addition, we created an applicable maturity model that could help practitioners assess their current state in the ML adoption process.

Conclusions: Many clinics still face major problems in adopting ML systems for medical diagnostics, and thus do not benefit from the potential of these systems. Therefore, both the ML adoption framework and the maturity model for ML systems in clinics can not only guide future research that seeks to explore the promises and challenges associated with ML systems in a medical setting but also be a practical reference point for clinicians.

Keywords

Machine Learning, Clinics, Diagnostics, Adoption, Maturity Model

License

Published under the terms of Creative Commons Attribution 4.0 licence

URL

<https://www.jmir.org/2021/10/e29301/>

4.1 Introduction

4.1.1 *Machine Learning Systems for Medical Diagnostics*

The ongoing digitalization is influencing the everyday activities of almost every individual, both in their private and professional lives. This transformation is particularly evident in health care, where the integration of health information technologies (HITs), such as electronic health records or clinical decision support systems, enables significant improvements in processes such as emergency medical care, diagnostics, and therapy (e.g., Hufnagl et al. 2019; Sun and Qu 2015; Wang et al. 2021). However, the integration of HITs is not a panacea but leads to major challenges in clinics as, fueled by these technologies, physicians have to handle an ever-growing volume of patient data and complexity of interacting systems (Bardhan et al. 2020). Moreover, societal problems further complicate the provision of health services to the population, as age-related diseases are on the rise because of demographic shifts and global pandemics such as the COVID-19 crisis are overburdening clinics, pushing medical personnel to the limits of their capacity (e.g., Kankanhalli et al. 2016; Li et al. 2020).

Artificial intelligence (AI) as the “science and engineering of making intelligent machines, especially intelligent computer programs” (McCarthy 2007, p. 2) could help relieve this burden on physicians as AI is capable of solving tasks previously reserved for human intelligence (Rai et al. 2019). In particular, machine learning (ML), as a subfield of AI, is currently one of the fastest growing technological approaches, opening up a wide range of possibilities for medicine (e.g., Jordan and Mitchell 2015; Shaw et al. 2019). Therefore, in the remainder of this research work, we focus on ML systems, that is, information systems (IS) that learn to perform certain tasks autonomously through experience without receiving explicit human instructions. Instead, ML systems use algorithms to search large amounts of data for patterns to create their own rules and strategies on how to deal with a particular problem. The identified rules can then be applied to solve a task (Brynjolfsson and Mitchell 2017; Jordan and Mitchell 2015; Meskó and Görög 2020; Russell and Norvig 2016). ML systems can be particularly useful in solving problems for which the rules are difficult to derive and express. This is the case, for example, in image recognition; for instance, how can the image of a cat be explained in terms of pixels, what shapes of ears are allowed, and how can they be recognized in a picture (Meskó and Görög 2020). From the prediction of patient admissions in clinics to therapy support, ML systems can help solve various problems in medicine (Shaw et al. 2019; Thrall et al. 2018). However, one application area of particular value to researchers and practitioners in which ML systems could have a major impact on the overall well-being of the population is medical diagnostics (e.g.,

Paton and Kobayashi 2019; Shahid et al. 2019). In this context, ML systems can help identify patterns in medical data (e.g., in medical scans, pathology slides, electrocardiograms, and written diagnoses) and sort possible conditions according to their likelihood (He et al. 2019; Topol 2019). A distinction can be made between ML serving to take over entire areas of responsibility from physicians and supporting them in their decision-making process. In the near future, ML systems will mainly be used as intelligent decision support rather than to automate medical diagnostics fully (e.g., He et al. 2019; Lebovitz 2019; Roski et al. 2019; Shaw et al. 2019). In this sense, current cases in research and practice show that an increasing number of such assistive ML systems are presently finding their way into medical workflows. For example, ML systems are being developed, refined, and deployed to help in the early diagnosis of COVID-19 based on entered symptoms or medical images such as computed tomography scans and algorithms such as deep convolutional neural networks (Jin et al. 2020). These systems raise the hope of making medical diagnostics of COVID-19 and also other diseases faster, more efficient, and consistent, and thus more valuable as they are able to compare patient data with a database that is larger than any physician's experience. Consequently, applying ML systems in patient care could make the difference between life and death by enabling more effective and efficient diagnostics (He et al. 2019; Shaw et al. 2019).

4.1.2 Challenges of Adopting Machine Learning Systems in Clinics

However, despite this enormous promise, the integration of ML systems also poses challenges that have prevented the widespread adoption of these systems in clinics to date (Kuan 2019). More specifically, clinics cannot draw on their experience from adopting other HITs, as ML differs substantially from prior technologies. Specifically, ML systems learn from high volumes of data instead of being explicitly programmed (Russell and Norvig 2016). Although traditional clinical decision support systems rely on rule-based systems that produce deterministic outputs, ML systems derive their solutions based on complex statistical methods, leading to several consequences. First, ML systems are becoming increasingly complex and commonly resemble black boxes; that is, their mechanisms for coming up with predictions are opaque to humans. For example, ML systems based on deep neural networks make predictions using millions of parameters, and humans cannot comprehend each and every calculation. Second, ML systems that learn from data will almost never be able to perform tasks perfectly, for example, make classifications with 100% accuracy (Brynjolfsson and Mitchell 2017; Lebovitz 2019). This is mainly because of the ML system reliance on statistical patterns, which will never be able to cover all edge cases. Third, the operationalization of ML systems in practice is challenging,

largely because complex relationships between different types of artifacts (e.g., data sets, models, and source codes) have to be managed (Amershi et al. 2019). Whereas traditional clinical decision support systems rely on human-defined rules that are instantiated in software code, ML systems are a result of applying algorithms to data, thus creating an additional dependency. All artifacts have to be versioned, and their dependencies must be tracked to comply with regulations and ensure reproducibility. Owing to these complicating factors, organizations in various industries struggle to integrate ML systems into their processes. Therefore, initial research is looking at the challenges that ML systems pose in terms of organizational adoption (e.g., Alsheibani et al. 2019a; Jöhnk et al. 2020; Kruse et al. 2019; Pumplun et al. 2019). However, clinics differ considerably from other organizations, as they not only possess unique structures, management processes, and requirements for HIT adoption but are also responsible for their patients' lives (e.g., Poba-Nzaou et al. 2014). In these medical settings, the characteristics of ML systems are particularly problematic as physicians and patients rely on profound diagnoses and the correct functionality of ML systems at any time (Lebovitz 2019). Consistent with the call of Davison and Martinson (2015) for more context-specific research, studies regarding the adoption of ML systems in clinics must, therefore, reflect on both, the specific characteristics of ML systems and clinics. Such context-specific research on the organizational adoption of ML systems in clinics is becoming more prevalent in recent times (e.g., Alami et al. 2021; Sandhu et al. 2020; Shaw et al. 2019). Thematically, researchers mainly investigate the individual acceptance of physicians (Lebovitz 2019; Sandhu et al. 2020) and the technical specifics of ML systems, such as their lack of transparency (e.g., Arora 2020; Kelly et al. 2019). However, the problem with existing research is that most of these publications are merely reviews and rely on the personal understanding and experience of the authors. Rare exceptions are, for example, Hofmann et al. (2019), Sandhu et al. (2020), and Sun and Medaglia (2019), who made use of qualitative research methods. Hofmann et al. (2019) examined the opportunities and challenges of ML systems in radiology, whereas Sandhu et al. (2020) and Sun and Medaglia (2019) studied the introduction of two specific ML-based diagnostic decision support systems in clinics. Although these publications already offer a first insight into the possible factors along the adoption process of ML systems, they are not sufficient to understand the process in its entirety.

4.1.3 Objectives and Research Approach

In particular, to our knowledge, no work exists that theoretically embeds the organizational adoption process of ML systems in clinics and presents it based on empirical evidence. Rather,

current research focuses on individual acceptance criteria instead of taking a holistic, organizational perspective (e.g., Lebovitz 2019; Sandhu et al. 2020). Therefore, clinics lack an integral overview of the requirements that ML systems imply and that they need to address to harness the potential of these systems for their diagnostic processes. Guided by the call of Shaw et al. (2019) for more research on the adoption of ML systems in clinics and the lack of prior integral research, our study thus aims to answer the following first research question:

RQ1: which specific factors influence the adoption process of ML systems in medical diagnostics?

Moreover, previous research does not elaborate on how these factors may manifest in a range of different stages and how these stages determine an overarching maturity score. However, such a maturity model could shed further light on the adoption process of ML systems in clinics by providing an empirically grounded and operationalized construct to measure adoption progress (Becker et al. 2009; Poepelbuss et al. 2011). Therefore, the maturity model could not only be applied in future empirical research but also allow clinics to assess their as-is situation and evaluate potential courses of action for ML adoption. Therefore, our research sets out to investigate the following second research question:

RQ2: how can the identified factors be used to establish a maturity model for the adoption process of ML systems in clinics?

To answer these research questions, we conducted a qualitative study based on explorative interviews (N=22) with experts working for clinics or suppliers of clinics. To structure the key findings of our empirical investigation, we referred to the health care-specific framework of nonadoption, abandonment, scale-up, spread, and sustainability (NASSS) for a conceptual basis (Greenhalgh et al. 2017). Although this adoption framework provides a foundation, it is not sufficient to represent the full adoption process of ML systems in clinics, given the particular characteristics of ML systems. To provide a more context-specific framework (Davison and Martinsons 2015), we drew on qualitative data to gradually adapt and expand the existing framework by several factors specific to the adoption process of ML systems for clinical diagnostics. Moreover, we used qualitative data to develop a maturity model that can help researchers and clinicians understand the possible range of ML adoption stages in clinics and determine an overarching maturity score. Overall, we aim to provide a practical reference point for clinicians to integrate ML systems more effectively into their diagnostic processes.

In the next section, we describe our qualitative research design, introduce directed content analysis as our basic data analysis methodology, and explain the development process of the

ML maturity model in detail. We then present the empirical results of our study to provide a valuable basis for further research and guidance to clinics aiming to integrate ML systems within their diagnostic processes. Finally, we conclude by discussing the theoretical and practical implications of our study and showing perspectives for future research.

4.2 Methods

4.2.1 Overview

Qualitative data provide a rich source of information that can help to better understand emerging, highly complex research subjects (Greener 2008). Therefore, to understand the complex adoption process of ML systems and derive a maturity model, we used a qualitative approach to “see the world through the eyes of the people being studied” (Greener 2008, p. 17). In this regard, we applied the key informant method and conducted in-depth interviews with experts (N=22) who have particular qualifications and specialized knowledge on the topic investigated (Bagozzi et al. 1991). We led these interviews according to a semistructured interview guideline to ensure that all relevant questions were posed. The questionnaire included general questions about the person, questions about previous knowledge in the field of ML systems, the assessment of potentials and challenges of ML systems for medicine, and further, more detailed questions about the prerequisites in clinics to adopt ML systems for diagnostics. Owing to the qualitative approach, we kept the guideline open and flexible to allow adaptations to the respective interviewed expert, her or his position, and knowledge base (Myers and Newman 2007). We analyzed the qualitative data with the help of *directed content analysis* (Hsieh and Shannon 2005) and the methodological approach for *maturity model development* (Becker et al. 2009). For an overview of the research procedure, please refer to Figure 7.

During the research process, we used several practices to obtain rigor and trustworthiness. To begin with, we defined two clear research questions and a conceptual framework that we used as input for our research design. Furthermore, we followed a theoretical sampling approach by iterating between data collection and analysis until we reached theoretical saturation (Flick 2004). In this way, we drew on the results from preceding interviews to select further experts and, for example, interviewed not only physicians and managers from clinics but also managers from HIT suppliers to obtain a more holistic perspective. In this regard, considering suppliers allowed us to gain an external, less biased perspective on the adoption of ML systems in clinics. Therefore, we found the additional supplier perspective to be particularly useful in triangulating the data and increasing the validity of our findings (Carter et al. 2014). Moreover, different

medical disciplines were considered in the interviews (e.g., radiology, pathology, and internal medicine) to allow for different perspectives on medical diagnostic processes (e.g., interpretation of medical scans, pathology slides, and electrocardiograms) and obtain more generalizable results (Benbasat et al. 1987). The resulting number of interviews is comparable with those of other qualitative studies in IS health care research (e.g., Hofmann et al. 2019; Horan and Schooley 2005; Sandhu et al. 2020; Whittaker 2012). With regard to data analysis, we followed a structured and reproducible approach to evaluate the qualitative data (Becker et al. 2009; Hsieh and Shannon 2005). During this whole process, a multiresearcher triangulation took place to include different perspectives on the research topic (Carter et al. 2014). In that sense, we discussed all data analysis steps and results intensively with the authors and with further qualified researchers from the fields of IS, computer science, and medicine. We recorded the results of these discussions in the form of memos to make them available in the following analysis stages (Urquhart et al. 2010). For later documentation of the results, we decided to include “the voice of participants” (Creswell 2007, p. 182) and thus quote directly from the interviews while presenting our findings. Where possible, we have additionally incorporated existing—so far scattered—literature that backs up and contextualizes particular statements made by interviewed experts, thus demonstrating the relevance of the findings from the interviews (Jöhnk et al. 2020).

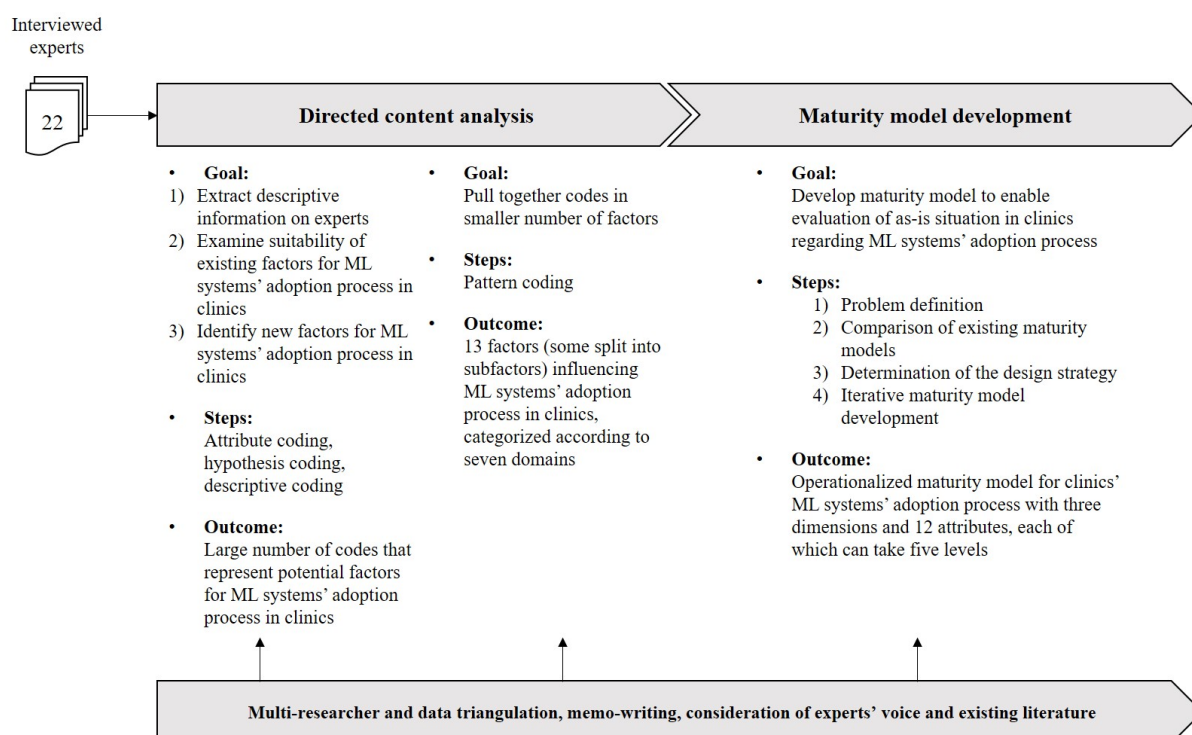


Figure 7: Overview of Research Procedure (Illustration based on Jöhnk et al. 2020)

4.2.2 *Data Collection and Sample*

Qualitative data were collected in two rounds. We conducted a first round of in-depth interviews from the second to the last quarter of 2019. This round of interviews included most participants (15/22, 68% of experts) and formed the basis for content analysis and maturity model development. However, the adoption of ML systems in clinics has progressed significantly in recent times. Therefore, we conducted a further round of interviews (7/22, 32% of experts) in the first quarter of 2021 to capture potential new insights from clinics on the research subject. Moreover, we shared the identified factors and the complete operationalized maturity model with the second-round interview participants to verify and refine the findings from the first panel. All the interviews were conducted in two European countries (Germany and Switzerland).

To identify suitable participants for both rounds of interviews, we searched for experts in professional networks, clinic websites, and at relevant conferences on ML in medicine. We interviewed qualified experts, who had detailed knowledge of clinical processes, had profound experience with ML systems, and were involved in the respective decision-making processes (Bogner et al. 2009). Of the 22 interviewed experts, five (23%) were physicians, eight (36%) held a hybrid position (i.e., physicians with additional leadership responsibilities), and nine (41) worked as full-time managers or information technology staff in the medical field. The participants worked for 11 different clinics and five HIT suppliers. Four clinics are privately financed, and the others are public, providing a view of both privately and publicly funded clinics. All clinics and suppliers are currently running projects related to ML. On average, each expert interview lasted 48 minutes and took place in a private space. The interviews were audio recorded and transcribed after mutual agreement. In three interviews, we only took notes as the participants did not consent to recording. For an overview of the experts, see Table 4.

Table 4: Overview of Interviewed Experts

ID	Position	Specialty	Expertise (years)
Clinics: key informants of clinics			
C-01	Physician	Radiology	3
C-02	Physician	Radiology	15
C-03	Physician	Radiology	8
C-04	Physician	Cardiology	3
C-05	Physician	Neuroradiology	3
C-06	Physician ^a	Neuroradiology	9
C-07	Physician ^a	Internal medicine	19
C-08	Physician ^a	Internal medicine	35
C-09	Physician ^a	Pathology	18
C-10	Physician ^a	Radiology	37
C-11	Physician ^a	Gynecology	40
C-12	Physician ^a	Otolaryngology	25
C-13	Physician ^a	Cardiology	12
C-14	Chief technology officer	Cardiology	8
C-15	Chief technology officer	Biomedicine	20
C-16	Director	Internal medicine	12
Health information technology (HIT) suppliers: key informants of clinics' HIT suppliers			
S-01	Director	Nephrology	20
S-02	Director	Biomedicine	22
S-03	Director	Genetics	10
S-04	Head of research and development	Radiology	2
S-05	System-engineer	Pathology	3
S-06	Innovation project lead	Surgery	3
^a Physician with leadership responsibilities.			

4.2.3 Directed Content Analysis

Our first goal was to identify the factors that are specific to the adoption process of ML systems in clinics and are not yet sufficiently covered by existing theories. As ML systems have an innovative character because of their novel, complex technical characteristics, we followed the steps of directed content analysis to extend existing theory on the adoption of innovations (Hsieh and Shannon 2005).

The process of adopting innovations in organizations is an overarching process that evolves from initial awareness of technology to a solidified interest and a subsequent adoption decision, to its implementation in the organization, and finally to continued adoption (Rogers 1995). Presently, adoption research regarding HITs has started to look beyond the mere awareness of a technology to include the later stages of the adoption process (Greenhalgh et al. 2017). In this context, ML systems own highly specific characteristics that will necessitate a significant change in the organization structure and working routines in the long run (Brynjolfsson and Mitchell 2017; Lebovitz 2019). Therefore, the whole adoption process of ML systems should be considered thoroughly. To capture this, we used the NASSS framework as a conceptual basis. NASSS has primarily been developed for the health care context by combining established health and social care frameworks and can be used to analyze the full adoption process of an HIT, including the implementation phase and continued adoption of the technology. It includes several *domains*, namely *technology* and its features, the *organization* that aims to adopt the technology, the *wider system* of an organization, the *condition* to be diagnosed and treated, the demand and supply side *value proposition* associated with HIT, and the *adopter system* consisting of patients, their relatives, and medical staff. Furthermore, it explicitly conceptualizes the *embedding and adaptation* of the HIT within a clinic over time (Greenhalgh et al. 2017). Each domain, in turn, comprises several *factors* that specify the domain considered. These are, for example, the regulatory issues related to a technology (wider system) or the value a technology can have for a patient (value proposition). The suitability of the NASSS framework for the topic under study is evidenced by recent research calling for the use of the framework for empirical work on the adoption process of ML systems in clinics (Alami et al. 2021). The NASSS framework forms the basis for our research but is insufficient to explain the specific adoption process of *ML systems* in clinics and, therefore, needs to be reconsidered. In this regard, we used the framework as a starting point, and it was adapted and expanded, taking into account the qualitative data (Hsieh and Shannon 2005).

Specifically, we applied an iterative multicycle coding process that is in line with directed content analysis, which consists of two coding cycles, between which we moved back and forth (Saldaña 2009). The first cycle comprised three different types of coding. Using *attribute coding* enabled us to receive descriptive information concerning the participant. *Hypothesis coding* was used to consider the prespecified conceptual framework (i.e., NASSS) and to examine the suitability of existing domains and factors regarding the adoption process (e.g., domain: value proposition; factor: patients' value through ML). In contrast, the *descriptive coding* approach allowed us to identify new aspects that go beyond the conceptual framework

by disregarding formerly identified domains and factors. As the coding procedure during the first cycle has led to a large number of constructs, we used *pattern coding* within the second coding cycle to pull together the codes into a smaller number of factors (Saldaña 2009). We performed the analysis using the NVivo 12 (QSR International) software. The result of the analysis is a holistic overview of domains, factors, and subfactors that influence the adoption process of ML systems for diagnostics (see section *Factors Influencing the Adoption Process of ML Systems in Clinics*).

4.2.4 Maturity Model Development

In a further step of our data analysis, we aimed to use (a subset of) the factors identified during content analysis to create a maturity model that can help clinics to assess their current state in the ML system adoption process. Organizations can have different maturities with regard to the management of technologies. To determine the maturity score of an organization regarding a certain type of technology, specified maturity (assessment) models can be used (Becker et al. 2009). These models constitute an instrument for organizations to “measure and assess domain capabilities at a given point in time” (de Bruin et al. 2005, p. 18). In this context, maturity models are valuable tools for organizations to assess and document their as-is state and, based on this, achieve directions for transformation and prioritization of potential investments (Becker et al. 2009; Paulk et al. 1993). Therefore, a maturity model comprises different *dimensions* that are subdivided according to specific *attributes*, each of which can take different *maturity levels*. Dimensions represent capability areas, for example, in the field of technology management, that should be exhaustive and distinct from each other. Attributes further specify these dimensions and represent practices, activities, or measures that can be taken by the organization and contribute to an organization’s maturity. Levels, on the other hand, are archetypal degrees of maturity which are often represented as a 5-step sequence of stages expressed by different labels (Becker et al. 2009; Fraser et al. 2002; Mettler and Rohner 2009; Raber et al. 2012). Becker et al. (2009) differentiated 5 levels, namely, (1) *initial*, (2) *assessing*, (3) *determined*, (4) *managed*, and (5) *optimizing*. The descriptions characterizing these levels may vary depending on the level definitions and the subject of investigation. However, in general, an attribute is considered to be at an *initial* (1) level if the processes investigated are still in their infancy, chaotic, and not consciously controlled by the organization, whereas the most advanced level *optimized* (5) stands for those attributes whose processes are already actively and continuously improved with the help of standardized feedback mechanisms (Fraser et al.

2002; Humphrey 1988). The overall maturity score of the organization, which can take one of the 5 levels described, results from the compilation of the individual attribute levels.

In recent years, maturity models have made their way into the health care sector. A literature review conducted by Carvalho et al. (2016) showed that clinical researchers and practitioners have established and applied various specified maturity models to understand and evaluate the integration of different HITs. However, there are no studies in the existing literature or insights from practice on a specific maturity model related to ML systems in clinics. To create a new maturity model for the ML adoption process in clinics, we followed the systematic development process outlined by Becker et al. (2009), which is loosely based on the design science methodology of Hevner et al. (2004). This methodological approach includes four steps that structure the development of maturity models and four more that accompany the application of maturity models in practice. As our primary goal was to create a maturity model for the adoption process of ML systems in clinics rather than the subsequent application of the model in clinical practice, we focused primarily on the first four steps.

The first step of the maturity model development process by Becker et al. (2009) is to define the problem underlying maturity development. The aim of this study was to provide researchers and clinics with the opportunity to evaluate the clinic status quo in the adoption process of ML systems. As clinics still struggle to integrate ML systems into their processes, we consider this problem particularly relevant and topical (Kuan 2019). After defining the problem domain and the target group, we searched for existing maturity models from adjacent research fields. In particular, we identified three maturity models that, although not specific to clinics, are drawn from the field of AI: the *artificial intelligence maturity model* by Alsheibani et al. (2019b), the *five maturity levels of managing AI* by Lichtenthaler (2020), and the *machine learning maturity framework* established by Akkiraju et al. (2020). All of them use a five-level maturity scale ranging from an *initial* (1) level to *optimized* or *integrated* (5). Although the framework by Akkiraju et al. (2020) was strongly technically oriented, Alsheibani et al. (2019b) and Lichtenthaler (2020) incorporated a management perspective as well. Although the identified maturity models helped provide a structure for the model to be built (e.g., levels and potential attributes) and specific wordings that could be used (e.g., “no data exists to train AI”; Alsheibani et al. 2019b, p. 7), no model is complete in itself or tailored to clinics. As clinics are highly specific in their structures and processes (e.g., Poba-Nzaou et al. 2014), we took initial ideas from the existing models but widely supplemented and concretized these ideas with the help of the content analysis results. In particular, we designed a new maturity model that is

specific to ML adoption in clinics, but which incorporates some basic structures and descriptions from existing models. In the following core step, the actual development of the maturity model takes place. We adopted an iterative approach that included four substeps: design-level selection, approach selection, model design, and testing. In total, three iterations were performed to develop the maturity model. In the first iteration, the existing maturity models and the results of the directed content analysis were considered to build a basic concept. In the second iteration, additional researchers from the field of IS and computer science were brought in to discuss and optimize the maturity model. In the third round, the maturity model was shared, discussed, and tested with eight of the medical experts (Becker et al. 2009). Within these iterations, we decided to adopt a multidimensional maturity assessment based on the results of the previously conducted content analysis. In particular, a subset of three domains was used for the dimensions of the maturity model; the corresponding factors or subfactors form 12 attributes that further specify these dimensions. Thereby, only those domains and factors were selected that clinics can modify themselves and are not set by external forces that are beyond the clinics' reach (e.g., from the wider system). The resulting attributes were then populated with individual-level descriptions using the qualitative interview data. Therefore, we started with the two extreme levels *initial* (1) and *optimized* (5) for each attribute, and the formulations for the levels in between were derived from the interview data, the existing maturity models and literature, or logical inference. The complete maturity model, including dimensions, attributes, and levels, was then discussed with eight of the medical experts, who confirmed its comprehensiveness, consistency, and adequacy. Following Joachim et al. (2011), the maturity model was mathematically operationalized to enable clinics to calculate an overall maturity score. In addition, we have developed a web application for using the maturity model that clinicians can apply to calculate their maturity level in the process of ML system adoption. The result of these iterative development steps is an evaluated applicable maturity model that can help researchers and clinics assess the current state of clinics in adopting ML systems (see section *A Maturity Model for ML Systems in Clinics*).

4.3 Results

4.3.1 Factors Influencing the Adoption Process of ML Systems in Clinics

4.3.1.1 Overview

As diagnostic procedures can differ within different medical specialties, the data analysis focuses on common factors that affect the adoption process of ML systems for diagnostics in

clinics and can be derived across all disciplines. An integrative overview of these factors is shown in Figure 8.

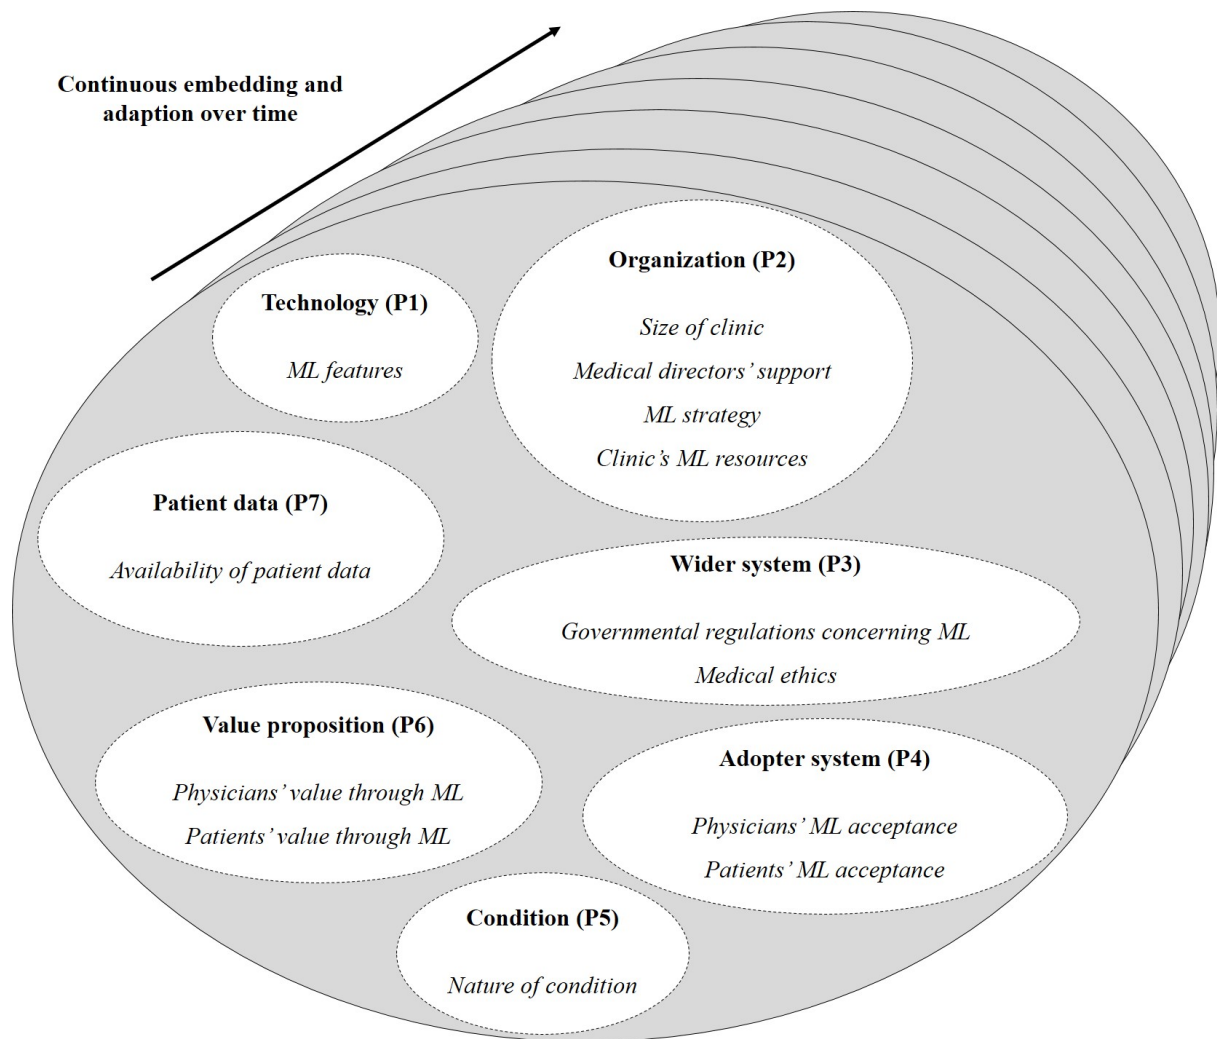


Figure 8: Integrative Framework for the Adoption Process of Machine Learning Systems in Clinics

In the following section, we present and discuss the results of our directed content analysis. For this purpose, we structured our findings according to the domains: technology, organization, wider system, adopter system, condition, value proposition, and the new domain patient data. The aforementioned domains interact with each other to enable the continuous embedding and adaptation of ML systems in clinics over time (e.g., Greenhalgh et al. 2017, 2018; James et al. 2021). In line with the existing literature, we thus did not formulate a separate domain to address the deep integration of ML systems across time. Rather, we assumed the embedding and adaptation over time to be a dynamic process in which, depending on the phase in the adoption process, specific domains and associated challenges are particularly relevant.

4.3.1.2 Technology

The features of technology are factors that are already considered within the original NASSS framework (Greenhalgh et al. 2017). Nevertheless, as outlined earlier, ML systems encompass several highly specific characteristics that cannot be compared with those of other HITs. Therefore, the existing general technical features factor is not sufficient to capture the properties of ML and has to be specified further.

As one subfactor of *ML features*, the interviewees pointed out the *lack of transparency* of ML systems as a major obstacle for the clinic's adoption of ML systems. ML systems based on neural networks can consist of multiple processing layers and up to billions of numerical weights, hampering the comprehensibility of ML systems to humans (e.g., Arora 2020; Brynjolfsson and Mitchell 2017; Kelly et al. 2019). Especially in high-stakes decision-making processes such as medical diagnostics, this can lead to major issues, as ML systems do not always provide correct suggestions (S-05). As a result, the experts state that physicians need to know exactly what the critical features considered by ML systems are and how identified patterns lead to conclusions. This is required so that physicians can assess the ML system's recommendations and suggest an appropriate diagnosis and therapy. One of the experts underlines this aspect:

“You will never make these existential decisions dependent on a black box, where it is not possible to understand what led to the recommendation.” [C-08]

Another subfactor of ML features is the *ability to adapt* their functioning if being retrained on novel data. This can become relevant either when the ML system is transferred to another context (e.g., another clinic) or needs to be retrained after some time; for example, new medical research results are gained or the patient demographic structure shifts. Clinics thus have to deal with an opaque system that is able to change its reasoning, making the outcome of an ML system unpredictable. Accordingly, experts see the adaptability of ML systems as another factor that has to be addressed by clinics (C-08, S-01, S-03, and S-05). To adopt ML systems, clinics need to have a clear strategy in place on how to cope with the opacity and adaptability of self-learning ML systems. Thus, we state our first proposition:

P1: *The features of ML systems (i.e., lack of transparency and adaptability) will impede their adoption in clinics.*

4.3.1.3 Organization

Looking at the organization domain, 4 factors emerged during the interviews. These are the *size of a clinic*, *medical directors' ML support*, *ML strategy*, and *clinic's resources for ML*.

The size of a clinic is a newly identified factor that was not specifically considered in the original NASSS framework. However, the interviewed experts emphasize that small clinics usually have fewer resources than large clinics, which could hamper the adoption of ML systems (C-15). In the specific context of ML systems, larger clinics further care for a higher number of different patients and thus have access to more patient data, which are needed to train ML systems appropriately (S-01).

Furthermore, experts state that clinic medical directors need to support the adoption of ML systems for diagnostic processes to guarantee financial and nonfinancial support for the new technology (C-03). In this regard, ML systems for medical diagnostics affect the core business of clinics and thus have strategic relevance (Zhu and Kraemer 2005). As medical directors develop the clinic's strategy, they are responsible for paving the way for the readiness of clinics to adopt ML systems. This is in line with prior research that states the significance of medical directors' support regarding the adoption of strategically relevant HITs in clinics (Lian et al. 2014; Yang et al. 2015).

As ML systems are a strategically relevant innovation, not only is the support of the directors necessary but also the establishment of an overarching, long-term ML strategy. The importance of an innovation strategy is also confirmed by an expert who emphasizes its relevance, especially against the background of the adoption of ML systems in a hospital network:

“When I want to launch it to the 1900 other hospitals, I have to think about a classic transformation strategy.” [C-16]

Such a strategy should include a plan of structured activities that contribute to the successful adoption of ML systems over time and should be supported by the clinic's medical directors (C-03).

One of the most frequently stated factors within the domain organization is the clinic's resource. This factor is similar to the factor capacity to innovate already included in the original NASSS framework but is subdivided into novel subfactors (i.e., *clinic's technical infrastructure*, *clinic's financing structure*, and *clinic's medical and ML methods expertise*). In line with existing literature (e.g., Panch et al. 2019; Shaw et al. 2019), some of the experts report that

clinics frequently rely on a wide range of clinical legacy systems, which are often proprietary to the suppliers, not connected, and based on outdated software and hardware:

“The primary challenge [...] is that the clinic usually consists of [...] million proprietary systems that are not connected.” [C-01]

This difficulty is not only present within the clinic itself but also translates to the interorganizational level. Although some experts state that their clinics already have some special data networks in place, almost half of the experts stress that health care organizations have not yet connected their data to systems in and outside the clinic (C-01, C-03, C-04, C-05, C-06, C-08, C-09, C-13, C-15, and S-04). However, experts emphasize the importance of having a high-performance technical infrastructure that can efficiently access data from multiple sources, for example, via secure internal (within clinic) and external data networks (e.g., clinic-to-primary care), which has the computing capacity needed to train ML systems (C-01, C-03, C-04, C-05, C-09, C-13, and S-04). Therefore, a clinic’s existing technical infrastructure could pose a major challenge to the adoption of ML systems.

Furthermore, the interviewed experts pointed out the problem of the current financing structure of clinics, which leads to strict budgetary constraints, especially in publicly funded institutions (C-04, C-05, C-11, C-12, and C-13). In this regard, an interviewee states that one part of their budget is assigned to daily costs, such as medication. The other part of the budget can be used to purchase large-scale medical equipment, such as x-ray systems. Thus, the development and setup of ML systems are not covered by either of the 2 parts, and often, no specific ML budget can be claimed (C-08).

Beyond that, there is a lack of personnel in clinics having expertise in both medicine and ML methods such as data science or data engineering:

“The shortage of medical specialists hits us twice as hard. We feel this at the medical professional side [...], but it is also very apparent at the technical side.” [C-14]

Both fields of knowledge are regarded as highly important for the adoption of ML systems by many experts (C-01, C-04, C-05, C-14, and S-02). Although a medical background can help identify relevant training data or assess the functionality of the ML system, ML method expertise is needed to train, integrate, and operate ML systems as presently, only scattered out-of-the-box ML systems exist for application in medicine, requiring clinics to develop and maintain ML systems by themselves (C-01, C-14, and S-02). Therefore, clinics need specific

expertise in the field of ML methods in addition to their medical understanding to develop, set up, and run ML systems in clinics. In sum, we propose the following:

P2: *A larger clinic size, medical directors' ML support, formulation of an ML strategy, and availability of resources for ML (i.e., technical infrastructure, ML budget, expertise in the field of medicine, and ML methods) will facilitate the adoption of ML systems in clinics.*

4.3.1.4 Wider System

With regard to the wider system, there are two relevant factors influencing the adoption of ML systems: *governmental regulations concerning ML* and *medical ethics*. Governmental regulations are a factor already known from the original NASSS framework. Nevertheless, the interviews revealed some particularities that were not covered by the general concept and are described below. Medical ethics is a factor that has not been captured by the NASSS so far but has been identified through our study.

In the field of medicine, there are several governmental regulations that must be taken into consideration when adopting ML systems. The following subfactors could be identified: *medical approval of ML systems, accountability, and the protection of sensitive personal data.*

The experts drew attention to the fact that HIT offered in the market and used in clinics is subject to several laws. This includes the need for medical approval conducted by legal authorities or HIT suppliers themselves (C-03, C-05, and C-12). In the United States, the Food and Drug Administration is responsible for the admission of medical products. In Europe, the HIT suppliers themselves need to perform a conformity assessment procedure, for example, based on the Medical Device Regulation (Food and Drug Administration 2021; Migliore 2017). As mentioned before, most ML systems are currently being developed by the clinics themselves and have not undergone any approval process (C-03). However, legal approval of ML systems is not trivial, as the systems can learn from new experiences and adapt themselves as described above:

“It is not obvious how evidence can be obtained for an [ML] model that differs significantly at the beginning, middle, and end of the study. If you want to approve a medical device today, you have to describe the intended use in detail.” [S-01]

The Food and Drug Administration addresses this legal uncertainty in an official statement that proposes an action plan for innovative approaches to more effectively approve adaptive ML systems (Food and Drug Administration 2021). The European Medicines Agency is also still in the early stages of defining and establishing an approval process for ML systems (European

Medicines Agency 2020). Therefore, legal ambiguities could represent a hurdle for clinics to adopt ML systems for diagnostics.

In addition to the medical approval of an ML system, there is the question of accountability for diagnoses. The experts interviewed indicated that it is questionable who takes over responsibility if the diagnosis prepared by an ML system is inaccurate (C-06, C-14, and S-05). It is also unclear who can be held liable—the HIT provider, the clinic, or the physician who is providing the medical diagnosis. An expert underlines this aspect with the following words:

“Then there are certainly [...] legal problems, for example: who is responsible for the interpretation and possibly wrong results of the ML model?” [C-14]

According to the current state of the art, ML systems cannot be held responsible for their output, as a registered physician is always obliged to validate and interpret the system’s results and perform the final diagnosis (C-16). However, it would ease the decision of clinics to opt for ML systems if there were a legal specification, especially if ML systems are increasingly able to automate steps of sensitive processes such as diagnostics (C-14 and C-15).

Another subfactor of governmental regulations, which could be identified as relevant for the adoption process of ML systems for diagnostics, is the protection of sensitive personal patient data. Patient data are widely considered as highly sensitive (Fox and James 2020) and are under special protection by national and international laws (C-02, C-04, C-13, S-02, and S-05). For example, the General Data Protection Regulation in Europe only permits the processing of health data if the patient explicitly accepts or if the clinic can provide particular reasons for the use of the data (European Parliament 2016). Thus, the respondents emphasized the clinics’ concerns in obtaining the necessary patient data to train the ML system (C-02, C-10, and S-06).

Using ML systems for diagnostic processes fueled medical ethics concerns among interviewees. On the one hand, ML systems are able to improve the efficiency and effectiveness of diagnostics (C-15, C-16, and S-02) and, on the other hand, the suggestions provided by ML systems are deduced based on statistical methods that recognize patterns in patient data that can be biased (C-15). Furthermore, the experts claimed that ML systems that are fed with patient data could determine whether a patient tends to develop a disease. This type of medical application would contradict the “patient’s right not to know” (C-15). Summarizing these remarks, we set up the proposition:

P3: *Uncertainties in governmental regulations, strict requirements for the protection of sensitive patient data, and existing medical ethics will impede the adoption of ML systems in clinics.*

4.3.1.5 Adopter System

The NASSS framework suggests that the successful adoption of ML systems is strongly influenced by individuals who are supposed to use the system or are affected by their suggestions. In this context, two ML-specific factors turned out to be relevant according to the interviews, which further specify the domain: *physician* and *patient ML acceptance*.

More than half of the interviewed experts stated that physicians' acceptance is essential for the adoption of ML systems in clinics (C-01, C-02, C-03, C-05, C-06, C-08, C-09, C-12, C-14, C-15, S-03, and S-06). As ML systems have the ability to solve tasks that were previously performed by humans, physicians might feel interchangeable in their job (C-03, C-07, S-03, and S-05). ML systems are trained on large sets of data that exceed the experience of any single physician, setting new standards for medical diagnostics. In this regard, most experts are concerned that physicians could reject ML systems for their daily work:

“As a doctor who may have ten or 20 years of experience [...], would I like to be taught by a machine [...]?” [S-03]

These concerns have recently found their way into pertinent research, demonstrating the relevance of the topic (e.g., Alami et al. 2021; Hofmann et al. 2019; Lebovitz 2019; Sandhu et al. 2020). However, it is also evident that the acceptance of ML systems differs among different age groups. In particular, physicians who belong to the group of digital natives are more willing to understand and ultimately use ML systems (S-04 and S-06).

Most interviewees stated the importance of patients' views on the use of ML systems for medical diagnostics. Although a physician is still involved in the decision-making process, patients might refuse the use of an ML system as the physician may be influenced by suggestions for possible conditions that are derived statistically and could be affected by biases. Furthermore, personal, sensitive patient data have to be processed to gain results. Therefore, experts state that patient acceptance of ML systems is highly relevant for the adoption of ML systems for diagnostics (C-02, C-06, and C-14). We thus conclude the following:

P4: *Physicians' and patients' acceptance of ML systems will facilitate the adoption of ML systems in clinics.*

4.3.1.6 Condition

As specified within the NASSS framework, *patient condition* affects the applicability of a technology. This is not only the case for conventional HITs but also holds true for ML systems, as stated by the interviewed experts (C-02 and C-09). ML systems have a narrow focus and can only deal with specific delimited problems (Brynjolfsson and Mitchell 2017; Russell and Norvig 2016). However, the human body is a highly complex and not fully understood system that can hardly be delineated. Medical conditions can be complex, poorly understood, or even unpredictable, for example, when multiple comorbidities are involved, making it difficult for ML systems to provide a clear diagnostic recommendation (C-02 and S-02). Therefore, the nature of the condition affects the applicability of ML systems, which can only handle delimited problems in the diagnostic process. Thus, the use of ML systems will be limited to the diagnosis of certain conditions:

P5: The limited applicability of ML systems for the diagnosis of specific conditions will impede the adoption of ML systems in clinics.

4.3.1.7 Value Proposition

The value proposition is another domain of the NASSS framework that we were able to concretize by analyzing the interviews. According to the experts, the adoption of ML systems could result in the creation of *value* for both *physicians* and *patients* (C-03, C-10, and C-14).

Integrating ML systems in their daily work enables physicians to improve the effectiveness and efficiency of their diagnostics as they can base their decisions on a broad database that is evaluated within a few seconds (C-16):

“If you have the choice among a pathologist who has already looked at 10,000 cuts [...] compared to one who has created only 500 findings, whom would you chose? But [...] AI has not only 10,000 but 500,000 findings in its memory.” [C-08]

In this regard, ML systems that are, for example, based on image recognition algorithms can surpass the ability of the human eye to capture details and patterns in x-rays (e.g., Akcay et al. 2018). If used for a second opinion, ML systems thus increase the quality of physicians’ work (C-02 and C-11).

In addition, patients could directly benefit from a decision that is faster and more informed if physicians use ML systems for diagnostics as a supportive tool (C-10 and C-16). We thus propose the following:

P6: *The additional value for physicians and patients created through ML systems will facilitate the adoption of ML systems in clinics.*

4.3.1.8 Patient Data

During the interviews, nearly all experts stated the *availability of patient data* as crucial for the adoption of ML systems for diagnostics. In this regard, patient data have to be available to develop and train the ML system in the first place and subsequently retrain it during use. This factor comprises various subfactors (i.e., *digitization of patient data*, *unified data formats*, *data quality standards*, *data anonymization* and *representativeness of training data*) which are described in the following section.

According to the experts, most clinics generate high volumes of patient data through their daily diagnostic processes (C-03, C-05, S-01, S-04, S-05, and S-06), which is basically a positive feature as an appropriate amount of data is needed to train ML systems (e.g., Brynjolfsson and Mitchell 2017; He et al. 2020; Roski et al. 2019; Sun and Medaglia 2019). However, although high volumes of data are generated, many processes in clinics are still paper-based, which lowers the proportion of patient data available in digitized form:

“Data are often not digitized, much is still in paper files, not structured, which means that the data availability is really extremely [...] poor.” [C-03]

This observation is in line with prior research concerning clinics that are lagging behind at using digitized technologies and digitizing patient data (e.g., Hufnagl et al. 2019). As a consequence, the interviewed experts see the integration of an electronic medical record system as a prerequisite for the application of ML systems (C-16, C-03, C-04, and C-13).

Furthermore, interviewing the experts revealed that medical patient data, if available in digitized form, are usually provided in a variety of proprietary data formats as many disparate clinical legacy systems from different suppliers have to interact to enable physicians to provide laboratory tests, diagnostic images, or clinical notes. These proprietary data formats are often difficult or impossible to convert, making the generation of consistent formats highly problematic (C-03, C-04, and S-04). The problem of differing data formats in clinics has already been recognized outside the ML context, for example, in research on the adoption of cloud solutions in health care environments (e.g., Zhang et al. 2017). Nevertheless, it is particularly critical for the introduction and use of ML systems that the patient data be processed for training and retraining the system. Although the first research has been conducted to allow for the transformation of different medical data types in one format (Lee et al. 2019), most clinics have

not yet been able to implement unified standards for patient data to enable processing and analysis by ML systems.

Furthermore, digitized patient data are often stored in unstructured file types, such as images, texts, or videos (C-01, C-03, C-07, C-13, C-15, and S-04). The experts cautioned that the quality of unstructured data is highly dependent on the particular clinic where the data are generated and their clinical staff (C-06, C-07, and S-04). For instance, physician letters are frequently written in free text formats, which are filled with synonyms and can be interpreted individually. More specifically, personal formulations are used, such as the description of a tumor size as compared with that of a walnut (C-07). Thus, patient data are not only hard to harness and have to be transferred to a machine-readable format first (C-03 and C-04) but also lack common quality standards (S-04), impeding the extraction of generalizable patterns through ML. Clinics aiming to adopt ML systems to support their diagnostics should therefore set standards for data creation, for example, by establishing a common language that physicians use when creating free texts. Such efforts are already being driven by some in-clinic as well as national initiatives (C-12 and C-16). In addition, other primary structured data sources could be connected, such as data from laboratory findings, to complement the unstructured data (e.g., Varghese 2020).

Moreover, the experts strongly emphasize that clinics that want to use patient data to train ML systems need to anonymize the sensitive data before processing them through an ML system (C-15 and S-06). However, anonymizing data might remove valuable information, which could be important for obtaining a diagnosis. For instance, information about a person's residence could facilitate a diagnosis if a disease is more prevalent regionally (C-15). Therefore, it is necessary for clinics to find the right balance of anonymization and information value to be able to use the data despite data protection regulations and still preserve all the information necessary to find meaningful correlations through ML systems. The first steps are already being taken in technical research to balance protection and the quality of sensitive data effectively (e.g., Prasser et al. 2014; Schneider and Tkachenko 2018).

According to the experts, the selection of the right training data is especially important in a health care context, as wrong diagnoses may have an impact on patients' lives. This leads to another aspect of patient data to be considered: the representativeness of training data. Patients in clinics vary in many aspects, from an outer perspective (e.g., age, gender, and hair color) as well as from inner functioning (e.g., size of organs and blood values; C-01 and S-05). If ML systems are trained based on an external database (e.g., collected via data exchange) that is demographically or regionally skewed compared with the clinic's conditions, false conclusions

could be drawn by the system. In this context, an expert raised the example of an ML system supporting the detection of skin melanomas, which is mainly trained on a sample of patients with a similar phenotype. Therefore, this pretrained ML system cannot be easily transferred to patients of other ages or with other skin pigmentations (C-01). In addition, the representativeness of the data is affected when different clinical systems, such as different radiographic systems, collect data as the resolution of the medical equipment may vary from provider to provider (S-04). As training data for supervised learning need to be labeled by humans, the same could be said regarding the expertise and working philosophy of physicians, which could be highly heterogeneous depending on the physician's knowledge state and working environment (C-09, C-14, and S-05).

The availability of patient data is a factor that is decisive for the adoption process of ML systems that need to be fed and retrained:

***P7:** The availability and exchange of a large amount of digitized patient data for training (that are uniformly formatted, of high quality, anonymized but informative, and representative of the clinic) will facilitate the adoption of ML systems in clinics.*

4.3.2 A Maturity Model for ML Systems in Clinics

4.3.2.1 Overview

Against the background that no maturity model for the adoption process of ML systems in clinics could be found in research and practice, we created a concept for a maturity model and present the model below. On the basis of our empirical results, the model is intended to enable researchers and clinics to quantify the overall maturity of clinics within the adoption process of ML systems. We followed the design process of Becker et al. (2009) to conceptualize a maturity model that comprises three dimensions and 12 attributes, each of which is operationalized by five corresponding levels (Table 5). The dimensions and attributes are derived from a subset of the results presented in the previous section, whereby the dimensions were inferred from the domains and the attributes from the factors or subfactors that can be modified by the respective clinic itself. Specifically, the dimensions organization (P2), adopter system (P4), and patient data (P7) and their respective subfactors were taken into account, as these can be controlled by the clinic itself, whereas the technology (P1), the wider system (P3), the condition (P5), and the value proposition (P6) are influenced by factors that are not in the hands of a single organization.

Table 5: Maturity Model for Machine Learning Systems in Clinics

Dimension	Attributes	Level: initial (1)	Level: assessing (2)	Level: determined (3)	Level: managed (4)	Level: optimized (5)
Organization						
	Medical directors' ML ^a support	No ML support of medical director	Low ML support of medical director	Medium ML support of medical director	High ML support of medical director	Extensive ML support of medical director and dedicated ML representatives (eg, Chief AI ^b Officer)
	ML strategy	No ML strategy in planning	ML strategy planned but not yet developed	ML strategy is under development	ML strategy is under implementation	ML strategy is largely implemented and supported by all stakeholders (eg, medical directors and physicians)
	Clinic's resources for ML					
	Clinic's technical infrastructure	Unconnected and outdated software systems	Solutions for connecting outdated software systems for data exchange within the clinic	Up-to-date software systems connected for data exchange within the clinic	Up-to-date software systems connected for data exchange within the clinic and with selected partners (eg, other clinics and academic institutions)	Up-to-date software systems deeply integrated in data networks for regular and extensive data exchange within the clinic and external institutions (eg, other clinics and academic institutions)
	Clinic financing structure	Strict budgetary constraints exclusively for medication and large medical equipment	Exceptions to budget constraints	Partially free allocation of budget	Free allocation of budget	Dedicated ML budget
	Clinic medical and ML methods expertise	Medical expertise only and no ML methodological expertise	Medical expertise and first medical experts with self-developed ML methodological expertise in data science	Medical expertise and first dedicated ML method experts in data science	Medical expertise and dedicated ML method experts in data science and engineering	Medical expertise, dedicated ML methods experts in data science and engineering, and all have insights into each other's area of expertise (eg, through training and academic courses)
Adopter system						
	Physicians' ML acceptance	Skepticism of physicians toward ML	Conservative attitude of physicians toward ML	Pragmatic attitude of physicians toward ML	Openness of physicians toward ML	Enthusiasm of physicians toward ML
	Patients' ML acceptance	Skepticism of patients toward ML	Conservative attitude of patients toward ML	Pragmatic attitude of patients toward ML	Openness of patients toward ML	Enthusiasm of patients toward ML
Patient data						
	Availability of patient data					
	Digitization of patient data	Nondigitized data (eg, paper-based data)	Low level of data digitization (eg, only central organizational processes)	Medium level of data digitization (eg, central diagnostic values)	High level of data digitization (eg, time series data on the course of diseases)	Fully digitized data
	Unified data formats	Many different proprietary data formats that are not convertible	Many different proprietary data formats, some of which are convertible	Many different but common data formats that are convertible	Many common data formats	Unified standards for medical data formats (eg, DICOM ^c and FHIR ^d)
	Data quality standards	No uniform quality of data (eg, handwritten, synonym-biased diagnoses)	First requests to unify quality of data within clinic	Clinic-internal quality standards (eg, for key terminologies)	Inter-clinical quality standards (eg, for key terminologies)	International quality standards (eg, UMLS ^e)
	Degree of information value despite anonymization	Nonanonymized personal data	Anonymized data, with almost complete loss of diagnostically relevant information value	Anonymized data, with low level of diagnostically relevant information value	Anonymized data, with medium level of diagnostically relevant information value	Anonymized data, with preserved diagnostically relevant information value
	Representativeness of training data	Nearly no match of demographics, region, and work culture in training data related to patient base and clinic	Low match of demographics, region, and work culture in training data related to patient base and clinic	Medium match of demographics, region, and work culture in training data related to patient base and clinic	High match of demographics, region, and work culture in training data related to patient base and clinic	Full match of demographics, region, and work culture in training data related to patient base and clinic
^a ML: machine learning. ^b AI: artificial intelligence. ^c DICOM: Digital Imaging and Communications in Medicine. ^d FHIR: Fast Healthcare Interoperability Resources. ^e UMLS: Unified Medical Language System.						

It is necessary to operationalize the model mathematically to render the maturity model applicable for research and practice. To this end, we followed the approach of Joachim et al. (2011), which has already been used for the operationalization of other maturity models (e.g., in the area of business intelligence; Raber et al. 2013). We assume that maturity evolves linearly in five levels $l \in L$ with $L = \{1, 2, 3, 4, 5\}$, starting with *initial* (1) and ending with *optimized* (5; Raber et al. 2013). The maturity model for the adoption of ML systems in clinics consists of three dimensions, d , each of which consists of a set of attributes I_d in turn. Therefore, the overall maturity score of a clinic is composed of the maturity score of all dimensions, whereby the maturity of each dimension d depends on the maturity within the corresponding attributes $a \in I_d$. As a clinic can have different maturities in the different dimensions and attributes of a

dimension, a stepwise estimation of the overall maturity score must be made. Therefore, a two-step process is followed in which (1) the *maturity score of the dimensions* (i.e., Mat_d) is determined first based on the respective attributes, followed by (2) the calculation of the *overall maturity score of a clinic* (i.e., Mat).

4.3.2.1.1 Maturity Score of the Dimensions

At the lowest layer, each attribute a can take a value $x_a \in A$ with $A=\{1,2,3,4,5\}$ depending on the actual maturity of the clinic regarding the attribute, ranging from initial (1) to optimized (5). To determine the actual maturity value of each attribute in a dimension, a clinic must assess its own as-is situation by comparing the level descriptions (within each attribute) with their current adoption state in the clinic (Table 5). For example, a clinic has a maturity value of $x_a=1$ for the attribute *digitization of patient data* if it has nearly no digitized data available for training ML systems and is thus at an initial level of maturity. In the next step, all maturity values x_a of the attributes within a dimension d are compared with all possible maturity levels l to determine the level with the smallest distance to the set of attributes of a dimension. To operationalize the comparison, a weighted Euclidean distance metric $Dist_d(l)$ is used in line with prior research (Joachim et al. 2011; Raber et al. 2013):

$$Dist_d(l) = \sqrt{\sum_{a \in I_d} (x_a - l)^2} \text{ for } 1 \leq d \leq n_d = 3 \text{ and } 1 \leq l \leq n_l = 5 \quad (1)$$

where n_d represents the total number of dimensions and n_l is the total number of levels. As a result, each clinic receives five distance values (for five levels, l) per dimension. To obtain the maturity score of a dimension Mat_d , the level m associated with the minimum of these distance values needs to be selected per dimension:

$$Mat_d = m, \text{ such that } Dist_d(m) = \min_{1 \leq l \leq n_l} (Dist_d(l)) \text{ for } 1 \leq m \leq n_l = 5 \quad (2)$$

4.3.2.1.2 Overall Maturity Score of the Clinic

On the basis of the distinct maturity scores Mat_d of the three dimensions, the overall maturity score Mat can be calculated in the second step. Again, we use a Euclidean distance metric $Dist(l)$ to compare the maturity scores of the dimensions with levels l (Equation 3). The final overall maturity score of a clinic striving to adopt ML systems is determined by the minimum distance (Equation 4):

$$Dist(l) = \sqrt{\sum_{d=1}^{n_d=3} (Mat_d - l)^2} \text{ for } 1 \leq l \leq n_l = 5 \quad (3)$$

$$Mat = m, \text{ such that } Dist(m) = \min_{1 \leq l \leq n_l} (Dist(l)) \text{ for } 1 \leq m \leq n_l = 5 \quad (4)$$

To make the maturity model easily applicable for practitioners from clinics and researchers in the field of adoption science, we have developed a free-access web application based on the described mathematical operationalization, which calculates the maturity level of a clinic based on a questionnaire (Figure 9 and Figure 10). This questionnaire includes the attributes as well as their level descriptions and is provided on the web (Pumplun et al. 2021).

Determine your clinic's readiness for ML-supported diagnostics

Machine learning holds the potential to improve diagnostics in clinics. However, few clinical ML systems have been deployed yet, since their adoption process differs significantly from prior health IT. Our model helps you to calculate an ML maturity score for your clinic.

[Get started](#)

Organization
For each question, please select the attribute most accurately describing the situation in your clinic.

Medical directors' ML support

Level of medical directors' ML support

☐ No ML support of medical director

☐ Low ML support of medical director

☐ Medium ML support of medical director

☐ High ML support of medical director

Figure 9: Determine Your Clinic's Readiness for Machine Learning-Supported Diagnostics (Screenshot 1 of the Web Application)

Thank you for using the maturity model.

Please have a look at the results summary below. Check back in the future in order to receive recommended actions on your journey towards implementing ML-supported diagnostics.

Maturity score
Level 1: Initial

Your dimension scores

Organization
Level 1: Initial

Adopter system
Level 1: Initial

Patient data
Level 1: Initial

Figure 10: Thank You for Using the Maturity Model (Screenshot 2 of the Web Application)

4.4 Discussion

4.4.1 Principal Findings

ML has an impact on all areas of human life, including the health care system. In this regard, ML systems offer the opportunity to make diagnostics more efficient and informed. However, to harness ML for such an application, clinics need to deeply integrate ML systems into their clinical practice, a challenge that most clinics have not yet been able to overcome (Roski et al. 2019). As clinics own highly individual, patient-oriented processes, it is crucial for researchers to reflect on this specific context (Davison and Martinsons 2015; Poba-Nzaou et al. 2014). However, prior research is lagging behind to provide empirically proven factors that influence the adoption process of ML systems in clinics for diagnostic processes. To address this shortcoming, we set up a qualitative study to (1) establish an integrated overview of factors specific to an ML system adoption process in clinics based on the NASSS framework and (2) create an operationalized maturity model that clinics can apply to assess their as-is state of ML adoption progress to decide on further actions and prioritize investments.

4.4.2 Limitations and Future Research Opportunities

Before we discuss our contributions to theory and practice in detail, it is necessary to clarify the limitations of this study and show room for further research. As we pursued a qualitative approach, our results are based on the expertise of the 22 interviewees. To counteract potential problems of generalizability, we have not only applied various criteria to ensure rigor and trustworthiness of our study (e.g., theoretical saturation, multiresearcher and data triangulation, and inclusion of multiple medical disciplines) but also carefully selected only highly involved experts. Nevertheless, it might be interesting for further research to perform a follow-up study to validate the proposed framework and maturity model quantitatively. In this regard, it might be informative to evaluate the derived maturity model by applying it in clinics. In doing so, it could also be investigated whether practitioners attach different importance to attributes and dimensions. On the basis of these findings, the maturity calculation could be adjusted by introducing weights for attributes and dimensions.

Moreover, we conducted the interviews in only two European countries. As health care systems vary across nations, interviewing experts from other regions with different economic and cultural prerequisites could lead to differing results. Nevertheless, the relevance of the findings for the international context was substantiated with the help of existing literature and practice contributions from international authorities, which are cited in the *Results* section. For example,

the report of the Food and Drug Administration shows that the issue of medical approval of ML systems is also being discussed in the United States (Food and Drug Administration 2021). However, replication of this study in other countries would be useful to highlight possible differences within the adoption process of ML systems in clinics.

In addition, the rapid development of increasingly advanced ML algorithms could lead to systems that can not only augment but also automate diagnostic processes. Investigating automated diagnostics, which has not yet been applied in clinics, could produce different findings, although the results obtained in this study could provide first indications.

4.4.3 *Theoretical Contributions*

Despite the limitations discussed, our study makes several important contributions to research. To begin with, we demonstrated that the NASSS framework can be applied but has to be adapted and expanded to explain the full adoption process of ML systems for diagnostics in clinics. To the best of our knowledge, this is the first study to provide an empirically proven and integrative overview of the factors determining the adoption of ML systems for clinical diagnostics and thus show what clinics need to consider to effectively integrate ML systems into their processes. Therefore, we contribute to and extend prior adoption research in health informatics, which has recently called for looking at the entire adoption process of HITs rather than just the initial awareness of the technology (Greenhalgh et al. 2017). Although the identified factors are specific to diagnostic processes, it is conceivable that they may be applicable to other scenarios in which the cost of errors is high, such as ML-based treatment recommendations or medical prognoses in clinics.

Moreover, we have developed the first maturity model for ML system adoption in clinics, which contributes to the IS and medical body of knowledge by providing an empirically grounded and strategically derived artifact that depicts medical and ML-specific attributes and their level descriptions in detail. More specifically, the maturity model shows which attributes determine the status quo of clinics in adopting ML systems, how these attributes may manifest in descriptors according to five different maturity levels, and how clinics can evaluate their as-is state in the adoption process of ML systems. Researchers can apply the developed maturity model, for example, as an instrument in statistical studies investigating the adoption of ML systems in clinics. More specifically, the model can be used to operationalize the dependent variable in structural equation models or as a variable for multigroup comparisons (e.g., Ply et al. 2012), for example, to study the antecedents of clinical adoption of ML systems. Therefore,

both the adoption framework and the maturity model for ML systems in clinics can guide future health care–centric research that seeks to explore the promises and challenges associated with ML systems in a medical setting.

4.4.4 *Practical Contributions*

In addition, the empirically based results hold relevant findings for practitioners, who are increasingly facing rising health care costs, demographic changes, and overcrowding of the clinics, and thus need to improve the efficiency and effectiveness of their clinical processes. ML systems could be a solution to these problems but have so far only been sporadically integrated into clinics (Kuan 2019). In fact, our qualitative study shows that most clinics still have major problems integrating ML systems into their diagnostics. In this regard, the derived framework provides medical directors with a holistic overview of potential enablers and inhibitors during the adoption process of ML systems in clinics and could provide a roadmap for practitioners.

Moreover, the developed maturity model can be used by clinics to obtain the first impression of their as-is situation in the adoption process of ML systems and to quantify it in an overall maturity score (see the website to easily apply the model; Pumplun et al. 2021). Assessing the maturity score with the help of the model not only helps to make external comparisons between clinics but also to identify internal deviations of certain attributes from the overall status. This allows clinics to invest especially in these attributes that are far from the present overall performance and lower the clinic’s maturity score significantly to date. Thereby, the maturity model allows practitioners working for clinics to analyze their clinic’s current status quo, identify shortcomings, prioritize possible courses of action, and efficiently allocate scarce resources depending on the respective degree of maturity. In this way, our research can help practitioners identify tailored requirements for the successful adoption of ML systems in clinics and build relevant capabilities and resources needed in the age of AI.

5 Paper 2.A: Acting Egoistically in a Crisis: How Emotions Shape Data Donations

Title

Acting Egoistically in a Crisis: How Emotions Shape Data Donations

Authors

Pumplun, Luisa; Wagner, Amina; Olt, Christian; Zöll, Anne; Buxmann, Peter

Publication Outlet

Hawaii International Conference on System Sciences (HICSS)

Abstract

The spread of COVID-19 has affected all of us, be it financially, socially, or even physically. It has caused uncertainty and anxiety, which has put people into a “hot” mental state. Referred to as an empathy gap, people are assumed to make emotion-driven decisions in “hot” states compared to “cold” states, which contrasts with the normative assumption of rational decision-making in privacy research. Based on an experimental survey study among 445 participants, we investigate whether people’s mental state interacts with individuals’ information disclosure decision-making. We measure our research model in the context of actual health data donation, which constitutes a critical surveillance factor in the COVID-19 crisis. Thereby, we contribute to research by (1) investigating data donation behavior amid a crisis and (2) helping to explain further nuances of privacy decision-making and the importance of trust as a context-dependent driver of data donation.

Keywords

COVID-19, Epidemiologic Surveillance, Data Donation, Privacy Calculus, Trust

5.1 Introduction

In the years 2020 and 2021, the coronavirus infected millions of people worldwide. At the time of writing this paper, more than 3.7 million died because of it. To combat the rapid spread of COVID-19, governments worldwide implemented countermeasures like closing educational institutions and restricting public life. Some even imposed a lockdown and forced their citizens to stay at home for several weeks (Ågerfalk et al. 2020). Information systems have played a significant role in managing various crises (e.g., Fagherazzi et al. 2020). Research institutes and governments could particularly benefit from information systems that gather and store people's health data to predict and control the global spread of COVID-19 today—but also in the following years (Raghupathi and Raghupathi 2014). Accessing personal health data in real-time and processing it with the help of data analytics techniques (e.g., machine learning) promise new opportunities to measure and predict the spread of COVID-19. However, for such data-analytical endeavors to be successful, research institutes rely on the population's health data (Ågerfalk et al. 2020). Thus, to contribute to the current COVID-19 pandemic, this manuscript revolves around individuals' data donation decisions. Data donation describes the act of people willingly disclosing their personal health information for research institutes with no apparent flow of value in return (Skatova et al. 2014; Wessels et al. 2019). During the COVID-19 pandemic, many institutions attempt to collect and store personal health information. For instance, research institutes offer smartphone applications, which enable users to voluntarily share their current health status by answering questions, for example regarding their body temperature (e.g., Drew et al. 2020). Other applications rely on fitness data collected by fitness-trackers to detect infectious diseases even before users show first symptoms. Even though these emerging technologies seem to be promising, they spread slowly among the population, which counteracts their effectiveness (Kaspar 2020). Therefore, it is of utmost importance to understand individuals' decision-making to donate health data to better support the public health system. However, research currently does not provide guidance on how people form data donation decisions, especially during a crisis such as the COVID-19 pandemic. Previous research mainly relies on the privacy calculus (Dinev and Hart 2006; Laufer and Wolfe 1977) to explain the antecedents of information disclosure (e.g., Anderson and Agarwal 2011; Krasnova et al. 2010). In this vein, established privacy research characterizes internet users as rational agents capable of making rational disclosure decisions (Dinev and Hart 2006; J. H. Smith et al. 2011). However, a small body of research based on behavioral economics started to question individuals' capability of making rational decisions regarding their privacy (e.g., Acquisti et al. 2013; Alashoor et al. 2018; Brakemeier et al. 2017; Dinev et al. 2015).

They provide evidence that behavioral biases like immediate gratification (Acquisti 2004), overconfidence (Wagner and Mesbah 2019), and stereotypical thinking (Gerlach et al. 2019) lead to corresponding errors in privacy-related decisions.

Pandemics cause heightened uncertainty levels among the affected population leading to stress, anxiety, and frustration as an emotional response (Bavel et al. 2020; Naidoo 2020). Faced with the COVID-19 pandemic, we expect an *empathy gap* to distort individuals' rational and stable disclosure decisions. Empathy gaps occur if people fail in behaving in accordance with their *cold* mental state (low intensity of emotions) when being in a *hot* state (high intensity of emotions; Van Boven and Loewenstein 2005; Loewenstein 2005; Metcalfe and Mischel 1999). Specifically, if the intensity of emotional influences arises, these emotions drive intuitive, automatic, and spontaneous actions rather than rational ones (Loewenstein 1996, 2005; Slovic et al. 2002). Similar to how the word *hate* triggers negative emotions (Dickert and Slovic 2009; Slovic et al. 2002), asking people to donate data for research on a *current pandemic* can put people in a hot state and, in turn, distort rational decision-making. In contrast, people who make a donation decision for a more distant (less emotionally arousing) research purpose are in a cold state and act more calmly, rationally, and receptive (Loewenstein 2005; Slovic et al. 2007) as formerly proposed by privacy research (Dinev and Hart 2006). Motivated by investigating further nuances of the privacy calculus (Dinev et al. 2015) and guided by behavioral economics research (Loewenstein 1996, 2005; Slovic et al. 2007), our research questions (RQ) are:

RQ1: *Are data donation decisions formed differently when people are in a cold state compared to a hot state?*

RQ2: *And if so, how do these states influence the magnitude of actual data donation behavior?*

To answer both research questions, we conduct an experimental survey study among 445 participants. We manipulate respondents' mental state by asking participants to donate data for either research on the current threat of COVID-19 (causing a hot state) or Ebola as a more distant research purpose (causing a cold state). Specifically, we test the effect of the mental state on the link between institutional trust and data donation behavior and the relationship between privacy risks and data donation, respectively. We show that trust is only significantly related to data donation behavior for research on Ebola (cold state). In the case of COVID-19 (hot state), participants focus primarily on privacy risks that inhibit self-disclosure. Therefore, the magnitude of actual data donation in a hot state is lower than in a cold state.

Understanding how mental states influence data donation decisions is a crucial and timely societal topic that provides several implications. We highlight that—amid a crisis—research

institutes that rely on real-time health data should promote a strong image of privacy-friendliness to lower perceived privacy risks. Surprisingly, trustworthiness does not positively affect people's decision to donate data while in a hot state, lowering the total amount of data donated to crisis-related causes compared to non-crisis-related donation purposes. In the long run, research institutes should establish a continuous, more passive data gathering process that allows collecting data across crises and non-crises periods. Moreover, even though fear appeals effectively motivate people to follow governmental rules (Rowe et al. 2020), it is also necessary to mark *COVID-19* not only with negative emotions causing a hot state but also with hope and social cohesion to strengthen prosocial behavior. Furthermore, we contribute to privacy research by exploring context-specific disclosure decisions. Hence, our findings contribute to the literature on paradoxical information disclosure decisions and respond to the need to investigate data donation decisions amid a crisis. We particularly provide insights on reasons for contradictory study results regarding the linear link between trust and disclosure decisions (Kehr et al. 2015; Norberg et al. 2007).

5.2 Theoretical Background

Research on information privacy typically assumes that Internet users disclose personal information based on a rational trade-off between their benefits, privacy risks, and institutional trust (Culnan and Armstrong 1999; Dinev and Hart 2006; Krasnova et al. 2010). Referred to as the privacy calculus, Internet users release their personal information only if benefits exceed perceived privacy risks (Dinev and Hart 2006; Xu et al. 2009). Simultaneously, users' institutional trust may countervail privacy risks and thus leverage information disclosure (Dinev and Hart 2006). While ones' benefits are immediate and vary across different contexts (J. H. Smith et al. 2011), users' perceptions of privacy risks and trust are the major overarching influential factors for information self-disclosure (Dinev and Hart 2006; Wakefield 2013; Xu et al. 2009).

Data donation constitutes a sub-form of self-disclosure where individuals consciously decide to voluntarily disclose their personal health related information to a collective dataset for free to support research institutes (Bietz et al. 2019; Skatova et al. 2014). Driven by altruistic motives, similarities can be drawn between donating data, blood, or organs for the benefit of others' health (Skatova et al. 2014). Especially in the context of data donation, privacy risks and trust perceptions should dominate data donation decisions since neither an exchange of value occurs nor are benefits offered in return for releasing personal information. Instead, the benefits of donation are more complex, distant, and abstract in nature (e.g., supporting research;

Siminoff et al. 2007). To investigate trust and privacy risks perceptions in more ambiguous situations (Norberg et al. 2007), we analyze the extent to which an empathy gap leads to differing trust and privacy risks assessments.

Cold and hot states delineate the duality between the *cold* rational and the *hot* emotional parts of ourselves. It results from individuals' inability to empathize with their different states, which is termed an empathy gap (Van Boven and Loewenstein 2005; Loewenstein 2005). Making decisions in cold and hot states leads to different behavioral outcomes (Loewenstein 1996, 2005). People deliberately analyze all available information in a cold state to arrive at a rational decision (Loewenstein 1996, 2005; Metcalfe and Mischel 1999). People's decision-making is guided by reasons, logical connections, and past experiences (Metcalfe and Mischel 1999; Slovic et al. 2002). They make conscious decisions based on their long-term preferences (Metcalfe and Mischel 1999). This perspective is also applied in privacy research, where people are assumed to make a rational analysis to arrive at a disclosure decision (Culnan and Armstrong 1999; Dinev and Hart 2006). However, psychologists emphasize that as soon as the intensity of the felt emotion arises, it reduces individuals' motivation to process all available information rationally (Loewenstein 1996, 2005). Instead, individuals act based on their gut feelings and following their acute short-term goals (Metcalfe and Mischel 1999; Slovic et al. 2002). Having intense emotions puts users in hot states where information is processed automatically, intuitively, and driven by emotional factors, which lead to a spontaneous decision (Loewenstein 2005). These intense emotions occur if a stimulus triggers emotions associated with that specific object (Slovic et al. 2002).

Indeed, donation research has shown that emotional stimuli can influence donation behavior. For instance, within the organ donation context, Kopfman and Smith (1996) provide evidence that non-holders of organ donor cards are primarily fearful of negative consequences attached to the sign-up for organ donation. Even though their study participants state a high intention to sign up for a donation card and have a positive attitude towards organ donation, they finally fail to transform their intention into action. Several scholars discuss in this context how external stimuli of "sensationalistic, negative media portrayals" (Morgan et al. 2005, p. 674) amplify fears related to organ donation. For instance, spreading myths about prematurely declared deaths or corruption in the allocation of organs eventually leads to fewer organ donations (Morgan et al. 2005). Due to the current shortage of organ donors, researchers analyze how to raise awareness and reduce negative emotions related to organ donation using social media (Murphy et al. 2020). A similar phenomenon is shown in the context of blood donation, where

potential donors reject blood withdrawal due to their fears related to the process (France and France 2018). Specifically, confronting potential donors with the critical decision to donate blood can evoke anticipatory anxiety associated with the “fear of blood, needles, pain, and fainting” (France and France 2018, p. 114). Therefore, powerful negative emotions can hinder people from making sound decisions based on their long-term preferences (e.g., helping others; Loewenstein 1996).

Even though donating organs or blood differs from the context of data donation, these findings provide first insights into possible phenomena related to emotions and donation behavior. Indeed, “one of the central emotional responses during a pandemic is fear” (Bavel et al. 2020, p. 461), potentially mitigating the donation behavior for research on COVID-19. Building on the above assumption that a negative emotional stimulus interacts with donation decisions, we extend the privacy calculus by investigating how different mental states influence data donation decisions.

5.3 Hypotheses Development

Defined as users’ worry about losing control over their personal information (Malhotra et al. 2004), privacy risks have been studied as the major impediment associated with information disclosure (Anderson and Agarwal 2011; Krasnova et al. 2010; Xu et al. 2009). However, prior privacy studies on the link between privacy risks and disclosure are inconsistent (Brakemeier et al. 2017; Dinev et al. 2015; Kehr et al. 2015). Recently, this influence is shown to be affected by emotions present at the time of decision-making (Kehr et al. 2015; Li et al. 2017). For instance, Alashoor et al. (2018) investigated the effect of an individual’s mood state on the link between perceived privacy risks and the intention to disclose information. They found evidence that negatively inclined users take greater account of perceived privacy risks. Indeed, people are more likely to narrow their attention on risks if an emotional stimulus triggers negative feelings (Loewenstein 1996). In this regard, emotions constitute an essential means of information (Hanoch 2002; Pachur et al. 2012). For instance, being anxious signals that something might be wrong and promotes attention to potential risks (Alashoor et al. 2018; Slovic et al. 2002). With regard to the COVID-19 pandemic, people have intense negative feelings such as financial uncertainty, frustration, depression due to social isolation, or even fear of one’s death (Bavel et al. 2020; Brooks et al. 2020; Naidoo 2020). Therefore—similar to how the word *war* triggers intense negative emotions (Dickert and Slovic 2009; Slovic et al. 2007)—the research purpose *COVID-19* might serve as an emotional stimulus that evokes strong negative emotions and consequently puts people in a hot state. These intense negative

emotions might hinder individuals' willingness-to-donate their data. In this vein, individuals narrow their attention, act more egotistically based on self-related fears, and cannot account for others' needs (Van Boven and Loewenstein 2005; Loewenstein 1996, 2005). As a result, if one is asked to donate health data to institutes researching COVID-19, we presume that people have a higher disposition to privacy risks and are less willing to donate data. In contrast, concerning less emotional research purposes, people are less likely to be biased towards risks as they act in a relatively calm situation (a cold state) where risks seem less salient (Pachur et al. 2012). Therefore, we postulate that being in a hot state leverages the effect of privacy risks on data donation behavior:

H1: The direct link between privacy risks and data donation behavior will be stronger in hot states.

Defined as an individual's belief that the other party is competent, benevolent, and honest, trust has been regarded as a decisive determinant to disclose personal information (Dinev and Hart 2006; McKnight et al. 2002). While a great body of studies could substantiate the trust-disclosure link (Bansal et al. 2010; Metzger 2004; Wakefield 2013), others indicate that trust is not significantly linked to disclosure decisions (Kehr et al. 2013; Norberg et al. 2007). The reasoning behind this is that the role of trust depends largely on the concrete institution in which trust is placed, on the time, place, and circumstances underlying the decision, and on the emotions attributed to a particular situation (Komiak and Benbasat 2004; Möllering 2005; Wakefield 2013). In this context, intense negative emotions such as anxiety can impair the relevance of trust (Bansal et al. 2016; Deutsch 1958). For example, put yourself in the scenario that you would like to do a bungee jump for the first time. You select a renowned provider with whom thousands have already taken the leap. In advance, you would rate the provider to be trustworthy. However, when you find yourself standing on the cliff, this previously granted trust might be irrelevant. Instead, your fear of heights prevents you from considering trust for your decision. Therefore, the effect of trust that has been granted based on rational expectations and available evidence in a calm situation (Komiak and Benbasat 2004) cannot simply be transferred to a more emotionally charged setting (Möllering 2005). Similarly, facing the COVID-19 pandemic, individuals are confronted with inevitable financial uncertainty, infection fears, and social isolation, which represent a completely new, so far unknown extreme situation with intense negative emotions attached (Bavel et al. 2020; Naidoo 2020). In such a scenario, the actors cannot draw on existing knowledge to extrapolate paths for rational decisions as the future is insecure (Loewenstein 2005; Möllering 2005). Since there are no

proven routines to rely on and fear determines actions, individuals no longer rely on trust but emphasize other, more intuitive, and salient antecedents to decide on their behavior (Möllering 2005). Therefore, while in a decision-making situation that is less emotionally charged and characterized by little uncertainty (cold state), one's behavior strongly depends on the established trust, we assume that this determinant has less influence on the decision to donate data in a hot state:

***H2:** The direct link between institutional trust and data donation behavior will be weaker in hot states.*

5.4 Methodology

To test our hypotheses empirically, we designed an online survey, which is a widely used methodology within privacy research (e.g., Dinev et al. 2013). At the beginning of the survey, we explained the study's purpose to every participant and promised full anonymity of all respondents' answers. The survey consists of two parts: First, we asked all participants to donate health data to us as a research institute (i.e., a university) to combat the spreading of a virus. In this vein, we were able to measure actual data donation behavior. Second, we ran a questionnaire consisting of our main measurements, demographics, and controls. We debriefed all participants at the end.

Based on the preceding arguments, we presume that the research purpose triggers emotions putting participants in different mental states depending on the emotions' intensity. Hence, by varying the purpose of the data donation as an emotional stimulus, we deem to manipulate the intensity of emotions and, thus, the mental state of participants that shapes the decision to donate data. For every participant, we randomly assigned either fighting against COVID-19 or Ebola as the data donation purpose and included a short informative description of the respective virus (i.e., the current level of dissemination, the existence of vaccine). We chose Ebola as the second research purpose for two reasons: First, the disease is comparable to COVID-19, both medically and in terms of countermeasures. Notably, countermeasures like data collection applications have been first used during the Ebola epidemic (Rowe et al. 2020). Second, Ebola is expected to evoke less intense emotions in European participants because it has been particularly prevalent in West Africa from 2013 to 2016, has a lower risk of transmission, and an effective vaccination exists already (Naidoo 2020). In contrast, we assume that COVID-19 causes a hot state among participants as it is highly transmissible, infects people worldwide, countermeasures are still active, and no licensed vaccination had been developed

during the time of our study. To test the success of the manipulation, we included a manipulation check on the intensity of emotions attached to the data donation purpose. We leaned on the negative affect scale from the PANAS-X list from Watson and Clark (1999) to distinguish the emotional response triggered by each scenario. After having read the welcome page along with the research purpose, we asked respondents to rate the following four statements along a 7-point Likert scale: (1) *COVID-19/Ebola distresses me*, (2) *I am afraid about the spreading of COVID-19/Ebola*, (3) *Thinking of COVID-19/Ebola makes me nervous*, and (4) *COVID-19/Ebola upsets me*. Being solely asked about their intentions to behave, participants frequently fail to predict their actual behavior. This *intention-behavior gap* is particularly pervasive in privacy research (Norberg and Horne 2007). To overcome this limitation of the privacy domain, we measure actual data donation behavior. Inspired by Acquisti and Grossklags (2005), we asked 14 different health-related questions of varying sensitivity, adapted from existing health data donation apps (e.g., RKI 2020). Our list included questions on basic health data (e.g., body weight), lifestyle data (e.g., drinking behavior), medical history (e.g., medication), as well as on virus-specific data (e.g., frequency of coughing). For every question, the participants could either answer to make a donation or refuse to respond and thus decline to donate. Hence, we measure *how many pieces of health information* every participant chose to donate. For ethical reasons, we did not store the individuals' health data donated to us at any time but only counted how many questions were answered (a number between 0 and 14). After measuring actual behavior as a dependent variable, participants answered a questionnaire with different measurement scales to assess their actual data donation determinants. All measurement scales are based on prior literature and are adapted to fit the context of data donation. We list our main scales in Table 6. The participants' answers are scored along a 7-point Likert scale ranging from *strongly disagree* to *strongly agree*. Furthermore, we included the following control variables into our questionnaire, which are based on established measurement instruments: age, empathy as a trait (Davis 1980), personal feeling of informedness about the virus (S. P. Smith et al. 2011), prior privacy experience (Xu et al. 2009), as well as general perceived health (Anderson and Agarwal 2011). We included several preventives to counteract possible common method bias (CMB) stemming from the survey-based methodology (O'Guinn and Faber 1989; Podsakoff et al. 2003). Moreover, a theoretically unrelated variable, namely fantasizing, was included as a marker to allow for testing for CMB retrospectively (O'Guinn and Faber 1989). To ensure that we only included attentive participants in the analysis, we integrated two attention checks.

Table 6: Measurement Scales

Var.	ID	Item
Privacy Risk (Dinev and Hart 2006)	PR1	I am concerned that my donated health records could be misused.
	PR2	I am afraid that others might gain access to my donated health data.
	PR3	I am concerned about donating my health data, because of what others might do with it.
	PR4	I am concerned about donating my health data, because it could be used in a way I did not foresee.
Trust (Everard and Galletta 2005)	TS1	This university is trustworthy.
	TS2	This university wants to be known as one who keeps promises and commitments.
	TS3	I trust this university keeps my best interests in mind.
	TS4	This university has more to lose than to gain by not delivering on its promises.
	TS5	This university's behavior meets my expectations.

We collected data from May to June 2020 in Germany with the assistance of a market research institute. Involving a market research institute enabled us to recruit participants who were not previously primed with the survey topic or originated from our network, potentially leading to biased results (Lowry et al. 2016). We included quotas (per manipulation group) to receive a data sample representing the average European Internet user (Eurostat 2018). Before analyzing the data, we dropped 17 cases because those respondents failed the attention checks. The resulting sample included 445 participants; 227 cases on COVID-19 and 218 cases on Ebola research. The characteristics of the sample can be found in Table 7.

Table 7: Characteristics of Sample

Characteristic		COVID-19 (n=227)	Ebola (n=218)
Age	18-30	25.6 %	34.9 %
	31-40	29.9 %	22.4 %
	41-50	21.6 %	22.1 %
	51-60	11.9 %	14.6 %
	> 60	11.0 %	6.0 %
Gender	Male	55.5 %	53.7 %
	Female	44.5 %	46.3 %
Monthly income (EUR)	< 800	19.0 %	22.8 %
	801 - 2000	39.7 %	35.8 %
	2001 - 3000	25.1 %	26.1 %
	3001 - 4000	8.8 %	7.3 %
	> 4000	7.4 %	8.0 %
Educational level	Intermediate mat.	24.6 %	26.1 %
	High school	26.0 %	21.1 %
	Vocational train.	22.0 %	24.8 %
	Univ. degree	26.9 %	27.5 %
	Other	.5 %	.5 %

5.5 Results

Before we share our results regarding our hypotheses, we elaborate on the data donation behavior variable, which we transformed to achieve a normal distribution. Furthermore, we checked that our manipulation regarding the mental state was successful. Afterward, we validated our measurement model. Finally, we ran a multi-group analysis (MGA) to test our hypotheses.

As our data for the actual data donation extent (i.e., how many pieces of information people donated) was not distributed normally (skewness=-2.676 and kurtosis=7.089), we calculated the fractional rank of the variable, resulting in uniformly distributed probabilities, and applied an inverse-normal transformation to form a variable of normally distributed z-scores (Templeton 2011). The resulting skewness is .967, and the kurtosis is .152, which are within the range between ± 2.0 for a normal distribution (Gao 2008).

In order to check whether the manipulation of the participants' mental state (cold vs. hot) was successful, we first averaged the items based on the negative affect scale from PANAS-X (Watson and Clark 1999). Subsequently, we compared the negative affect between both groups using a t-test. We justify this approach with a sufficiently large sample size within each group (Lumley et al. 2002). We found that the negative affect is significantly higher in the COVID-19 group (mean=4.257; sd=1.474; median=4.500) compared to the Ebola group (mean=3.314; sd=1.386; median=3.250) with $t(445)=6.950$, $P<.001$. Thus, successful manipulation of the mental state can be assumed.

Finally, we compared the data donation extent between the COVID-19 and the Ebola group based on a t-test of the normalized variable. The mean value of donation behavior for the Ebola group is significantly higher than the COVID-19 sample with $t(445)=-2.926$, $P=.004$. In particular, participants in the Ebola group donated on average 12.9 pieces of information, while participants in the COVID-19 group donated 12.0.

Before testing our hypotheses H1 and H2, we validated our measurement model for reliability, convergent validity, and discriminant validity. The Kaiser-Meyer-Olkin statistics (.890) confirmed that the data set is adequate for exploratory factor analysis. Concerning the item reliability, we checked the items' loadings with their respective constructs, which ranged between .870 – .953 for privacy risks and .662 – .848 for trust and hence exceeded the threshold of .6 (Nunnally 1967). We assessed convergent validity and discriminant validity by checking that Cronbach's alpha is greater than .7 and composite reliability is at least .7 (Bagozzi and Yi 2012). Moreover, the average variance extracted (AVE) should be higher than the threshold of

.5 (Bagozzi and Yi 2012). The following Table 8 shows that our measurement fulfills the requirements.

Table 8: Validity Measurements

Construct	Cr. Alpha	Comp. Rel.	AVE
Privacy Risk	.955	.955	.842
Trust	.876	.899	.642

To analyze for discriminant validity in our model, we assessed the cross-loadings between the constructs. In this regard, the correlation between privacy risks and trust should not exceed .7 (Fornell and Larcker 1981), which is not the case (.448). In sum, all scales for latent constructs possessed adequate reliability and validity.

As our main analysis, we evaluated the mental state's interaction and the model of privacy risks, trust, and data donation behavior. We created a model using SmartPLS representing privacy risks and trust as independent variables and data donation behavior as a dependent variable. The MGA is based on 5,000 bootstrapping iterations (Davison and Hinkley 1997), differentiating between the COVID-19 group (hot state) and the Ebola group (cold state; see Table 9). This approach is particularly suitable for analyzing heterogeneous data and two groups of equal sample size (Qureshi et al. 2009).

Table 9: Multi-Group-Analysis (* $P < .05$)

Path	β (Ebola)	β (COVID-19)	MGA (P)
Privacy Risk → Behavior	-.159*	-.176*	.87
Trust → Behavior	.269*	-.028	.01

All control variables are insignificantly related to data donation behavior. Moreover, neither did the CMB marker significantly correlate with our dependent variable ($\beta = -.063$; $P = .445$) nor did any regression path become insignificant. We, therefore, consider CMB not to be an issue in our data. Conducting the MGA showed that the effect of privacy risks on donation behavior is similarly strong in both groups ($P < .03$). Hence, we need to reject H1, as the mental states do not interact with the influence of privacy risks on data donation behavior. In contrast, the influence of trust on data donation behavior varied, depending on the participants' mental state. Specifically, trust significantly affects data donation behavior only for the Ebola group (cold state; $P = .03$), while the path-coefficient is close to zero and not significant for the COVID-19 group (hot state; $P = .73$). These results support H2 since the mental state mitigates the influence

of trust on data donation behavior. We report these results along with the path-coefficients of each group in Table 9.

5.6 Discussion

Our study aims to investigate individuals' data donation decisions in light of a current pandemic. An experimental survey study among 445 participants showed that trust is not significantly linked to individuals' data donation decisions for the COVID-19 group (hot state). Participants solely focus on perceived privacy risks. As a result, the intense negative emotions attached to the COVID-19 pandemic lead to a lower data donation extent than a less emotionally induced donation purpose (Ebola). This is concerning as gathering real-time health data has become a major surveillance factor in combatting COVID-19.

Our study implies several practical contributions that can help research institutes, politicians, and individuals to better handle the COVID-19 pandemic in specific and further potential crises in general. First of all, confirming the results of other studies (Weitzman et al. 2010), the majority of our study's respondents donated a high amount of health data for research purposes. At first sight, this result is surprising as health data is considered very private and highly sensitive (Bansal et al. 2016). However, this might be the case since the data was donated to a scientific institution to support research. Indeed, previous studies show that people's willingness-to-disclose data for a good cause is higher than for commercial purposes (Skatova and Goulding 2019). However, what is concerning is that the data donation extent was higher to combat Ebola, which is a disease already under control and not pervasive in Europe. In particular, our results show that this is the case since people in a hot mental state caused by the research purpose *COVID-19* are less able to decide rationally and narrow their attention to privacy risks when being asked to donate their health data. Based on strong negative emotions, they act rather egoistically than rationally. Even though trust might be rationally granted in the information recipient (e.g., a research institute) within a cold state, the actual willingness-to-donate data in a hot state is not affected by trust as a driver for disclosure, mitigating the data donation extent. This is critical as health data, and its analysis have become an effective instrument to combat the crisis, for example, by identifying virus hot spots (RKI 2020).

Since pandemics will continue to be possible in a globalized world, it is crucial to translate knowledge gained from crises experience and academic research into practice and start long-term preparations. To face possible future pandemics more effectively, we recommend establishing a persistent data donation platform on which people can provide their health data

continuously and repeatedly and with a broader, potentially less negatively affected research purpose (e.g., combating viruses in general). Like blood donation, a one-time registration with a subsequent continuing affiliation could be initiated (Cohen and Hoffner 2013). This way, data donation becomes more passive, and research institutes no longer need to actively encourage people to donate their data to combat an amid crisis, leading to a more defensive attitude or even reluctance regarding data donation. Furthermore, establishing a comprehensive dataset combining data from the pandemic state and from crisis-free periods might enable better analysis and respective countermeasures (e.g., early prediction of pandemic development) to prevent or combat pandemics. It is known from previous research that such data donation platforms have to comply with data protection regulations and be operated by a trustworthy institution (e.g., an independent research institution; Kaspar 2020). In this sense, trust-promoting seals of approval, similar to the seals used by online stores or service providers (Hui et al. 2007), could promote data donation platforms' dissemination in a non-pandemic state.

In addition to the long-term goals, our research can also help find short-term remedies to combat the COVID-19 crisis. Our results suggest that individuals within the hot state of the COVID-19 pandemic narrow their attention to potential privacy risks, leading to a reduced willingness-to-donate. A possible countermeasure against this could be to emphasize the privacy friendliness of the research institute to lower perceived privacy risks and eventually promote data donation. This could be done, for example, in the form of an official statement from the institute in which its strict data practices are disclosed. Besides, institutes could make potential data donors aware of their biased decision-making process within a hot state to broaden their narrowed attention to factors other than risks (specifically trust). In particular, institutes should openly explain to potential donors how intense and pressing emotions could influence them in a hot state. Individuals need to cool off to bridge the gap between their mental states and make stable decisions independently of its emotional momentum. Governments could actively support this process, which can help move the population from a hot state to a cold state in which rational decision-making is possible. Thus, marking *COVID-19* not only with negative emotions but also with positive feelings such as social cohesion, governmental resistance, and gratitude is necessary (Bavel et al. 2020). This could alleviate powerful negative emotions and keep individuals out of the hot state. In this regard, negative afflicted statements like "The situation is serious. Take it seriously, too." from the German chancellor Angela Merkel might fuel the high intensity of the population's negative emotions. In contrast, her statement, "I firmly believe that we can do this." might lower anxiety. As a result, even though fear appeals are common in crises, shaping the public's perception for good could be a key factor in

collecting health data—especially in times when people are actively seeking governmental guidance, for example, in the form of financial support or official information (Bavel et al. 2020).

Apart from practical contributions, our study's results provide several implications for theory on individuals' information disclosure behavior in general and privacy-related judgments in particular. The first overarching implication lies in investigating the interplay between individuals' actual willingness-to-donate and emotions attached to the donation purpose. Secondly, we contribute to the literature by extending the privacy calculus with a critical behavioral bias—namely, an empathy gap. In this vein, we respond to a call of Dinev et al. (2015) to test further nuances of the privacy calculus with regard to actual behavior and contribute to literature which questions the rational assumption of the privacy calculus (Acquisti et al. 2013; Acquisti and Grossklags 2005; Kehr et al. 2015). We show that users' privacy decisions are biased, but not regarding privacy risks as previously demonstrated in other studies (Brakemeier et al. 2017; Gerlach et al. 2019; Wagner and Mesbah 2019). Instead, individuals consider trust differently depending on their mental state present at the time of decision formation. When in a hot state characterized by negative feelings such as anxiety or frustration, individuals' emotions override rational trust analysis. People narrow their attention to losses and fully concentrate on their risks. They act rather egotistically and are, in turn, unable to account for trust in the data-collecting institute. Even though their long-term goal might be to combat the spread of the virus, they focus on their short-term goal of reducing privacy risks in this specific situation. This behavior is caused by an inability to empathize with their cold state, which would lead to deliberate information processing and thus rational decisions as traditionally assumed in privacy research. Therefore, we argue that the impact of trust is so far insufficiently studied and should be considered more comprehensively as an essential driver of self-disclosure. We support the notion that trust can be a key leveraging factor of self-disclosure online (Metzger 2004), particularly with regard to data donation in a cold state. However, our results explain why previous studies concerning the link between trust and self-disclosure have shown heterogeneous results. For instance, Norberg et al. (2007) find that trust is not linked to self-disclosure. However, most other studies building on the privacy calculus theory conclude that trust is a major driver of self-disclosure (Bansal et al. 2010; Dinev and Hart 2006; Metzger 2004). Our study sheds light on these contradicting results, which could stem from different situational factors associated with various mental states (Kehr et al. 2013; Norberg et al. 2007).

Even though we provide several implications for theory and practice, our results are not free from methodological limitations. First, we conducted an experimental survey study by manipulating the mental state caused by the purpose of data donation. Although we controlled for a successful manipulation, varying data collecting purposes could have led to different results. Our second limitation lies within the operationalization of data donation behavior. We decided to measure the magnitude of actual data donation based on a self-developed catalog of 14 questions inspired by available COVID-19 applications. This specific set of questions may have influenced participants' donation behavior. However, measuring actual disclosure is seldom in research and questions the transferability of study results to real-life behavior. We counteract this issue by testing individuals' actual data donation decisions instead of asking for the intended willingness-to-donate.

5.7 Conclusion

In these times, public health data has become a critical surveillance factor in combating the spread of COVID-19. Data analytics heavily rely on data to run algorithms and thus identify ideal countermeasures. However, research analysts depend on massive amounts of public health data for such endeavors to be successful. Therefore, understanding how people make decisions about data donation in a crisis is of utmost importance. In other words: It is not only the technology itself that will turn the tide in this pandemic, but rather people's prosocial behavior. Recently, we have seen enormous selfless behavior with high societal consequences: People went shopping for the elderly, supported local shops, and kept their distance from loved ones. It is now the task of politics, research institutes, and each individual to transfer this empathy to the area of data donation.

6 Paper 2.B: Opening the Black Box: Consumers' Willingness to Pay for Transparency of Intelligent Systems

Title

Opening the Black Box: Consumers' Willingness to Pay for Transparency of Intelligent Systems

Authors

Peters, Felix; Pumplun, Luisa; Buxmann, Peter

Publication Outlet

European Conference on Information Systems (ECIS)

Abstract

Artificial intelligence (AI) is becoming increasingly popular and intelligent systems are deployed for various use cases. However, as these systems typically rely on complex machine learning methods, they effectively exemplify black boxes. Thus, consumers are usually not informed about inner workings of these systems, e.g., data sources or feature importance. Public and private institutions have already called for fairness and transparency standards regarding intelligent systems. Although researchers develop mechanisms to ensure transparency of intelligent systems, it remains an open question how consumers perceive such transparency features. Consequently, our study examines to what extent consumers are willing to pay for these features, and what the underlying mechanisms of the purchase decision are. To answer these questions, we conduct an experiment and a subsequent survey in the context of credit scoring. We show that consumers exhibit significant willingness to pay for transparency. Furthermore, we observe that increased trust in the intelligent system caused by enhanced perceived transparency is the main driver for positive evaluation of transparency features. Our findings inform practitioners about the relevance of “fair AI” and manifest the importance of transparency research regarding intelligent systems in social sciences.

Keywords

Intelligent Systems, Machine Learning, Transparency, Willingness to Pay

6.1 Introduction

The field of artificial intelligence (AI) has seen a lot of technological breakthroughs in recent years: AI was able to defeat their human counterparts in strategic board (Silver et al. 2017) and online games (Vinyals et al. 2019) and has surpassed human-level performance on tasks like image recognition (He et al. 2015). Moreover, AI is increasingly being used for making predictions in high-stakes situations, e.g., medical diagnosis (Kourou et al. 2015) or credit scoring (Burrell 2015; Zhou 2017). Modern AI mainly uses machine learning (ML) to enable information systems to do tasks that we used to think were reserved for humans (Andrews et al. 2018; Brynjolfsson and Mitchell 2017; Elliot and Andrews 2017; Rai et al. 2019). These intelligent systems can be distinguished from conventional systems since they are able to learn based on incoming data, adapt their behavior over time without having to be explicitly reprogrammed, derive their results statistically, and thus will make mistakes (Andrews et al. 2018; Brynjolfsson and Mitchell 2017). High complexity of datasets and ML algorithms leads to what is commonly referred to as black box behavior: the lack of transparency in decision-making processes of intelligent systems. In this context, decisions which are inaccurate or not aligned with ethical standards due to biased incoming data could stay undetected (Kruse et al. 2019; Rudin 2019). Therefore, demands to open the black box of intelligent systems applied to high-stakes decisions (e.g., credit scoring) are increasingly claimed in research and practice (Burrell 2015; Chen et al. 2018; Google 2019; Rudin 2019; Shook et al. 2018; The Economist 2018). These requests coincide with the declared aim of the United Nations to support fairness, accountability and transparency of intelligent systems (ITU 2018).

The transparency issue regarding intelligent systems can be exemplified in the credit scoring context. So far, credit applicants could be sure that decisions about their creditworthiness were based on the evaluation of their credit history and the credit documents submitted. Furthermore, simple linear modeling approaches comprising a restricted number of well-known variables have been employed by financial institutions to guarantee the comprehensibility of the credit scoring process. This way, financial institutions are able to provide specific reasons to the credit applicant why one's credit was granted or denied (Martens et al. 2007). However, an increasing number of financial institutions additionally start to rely on tech companies such as Kreditech, Branch, or FICO to determine the creditworthiness of their applicants (Burrell 2015). These companies automatically collect applicant's personal data from online sources (e.g., corporate websites) and use complex non-linear models based on ML (Li et al. 2009; Zhou 2017). While these companies help consumers with limited financial history or past financial difficulties (e.g.,

in developing countries) getting access to credits (e.g., FICO 2019a), decisions about creditworthiness become less transparent than they were before. While the applicant has previously been granted the opportunity of clear reasoning, intelligent systems are limited in their comprehensibility and can hardly be validated to see if they discriminate applicants, e.g., due to gender, religion, or national origin (Shook et al. 2018).

To counteract this problem, a vast amount of research has been conducted over the past few years concerning the technical realization of transparency features for ML algorithms, e.g., neural networks (e.g., Diakopoulos 2016; Doshi-Velez and Kim 2017; B. Kim et al. 2018). In this context, more and more providers of intelligent systems for credit scoring are also striving to make their products more transparent (FICO 2019b). In contrast, information systems (IS) research has only recently begun to analyze the added value derived by transparency features and its effect on user's behavior in order to contribute to a responsible design of intelligent systems. In this context, Rzepka and Berger (2018) claim that a transparent system appearance with appropriate explanations shows positive impact on user's adoption decision. However, it is not yet known how consumers from the general public, who will increasingly depend on decisions from intelligent systems, assess enhanced transparency of intelligent systems since IS research on this topic is still in its infancy. In order to contribute to a more holistic picture of consumer's assessment of transparency features for intelligent systems, we examine not only their attitude towards these features but also go a step further and investigate their willingness to pay (WTP). This approach aims to demonstrate the interest of an emancipated consumer in more transparency while at the same time creating an incentive for companies to invest resources in improving the explainability of intelligent systems and offer them to customers. Thus, we pose the following research questions:

RQ1: *To what extent are consumers willing to pay for transparency features in the context of intelligent systems?*

RQ2: *What are the underlying mechanisms of the purchase decision for transparency features of intelligent systems?*

Aiming for an understanding of the value consumers assign to enhanced transparency, this study is the first to explore WTP for transparency features of intelligent systems. Therefore, we run an experiment in the credit scoring context, analyzing whether and why consumers would be willing to pay for features that explain rationale about how their creditworthiness was predicted by an intelligent system. Furthermore, the underlying effects of purchasing behavior are examined by conducting a subsequent survey with the 195 experiment participants. The

credit scoring context is chosen because it has been used to stress the importance of transparency of intelligent systems repeatedly in prior research (Burrell 2015; Kruse et al. 2019; Shook et al. 2018). We show that users exhibit a significant WTP for transparency features and that an increase in trust towards the intelligent credit scoring system is the main reason why consumers positively evaluate such features. As a result, we find a profound interest of consumers for fair and transparent AI, which is not only reflected in their attitude towards transparency features but is actually characterized by a significant WTP. In this regard, the positive evaluation of transparency features enables practitioners to assess their economical relevance for intelligent systems.

Thus, the remainder of this paper is structured as follows: In the beginning, an overview of related work concerning intelligent systems and possible transparency features is provided in order to mark out the research area. Here, we consider prior work from both information systems and computer science. Subsequently, an initial research model is derived and the quantitative research approach, including the experiment as well as the survey design, is presented. After introducing our sample, the results of the study are analyzed. In our discussion, we shed further light onto our contributions, limitations and opportunities for future research. Finally, we conclude our paper by summarizing the investigated problem, our research approach and findings.

6.2 Related Work

6.2.1 Intelligent Systems

AI as a research area covers aspects of mathematics, economics, computer engineering, cybernetics and linguistics, and is concerned with the development and understanding of intelligent systems (Russell and Norvig 2016). Intelligent systems such as speech-based assistants, search engines or credit scoring systems are already used by private persons in everyday life (Burrell 2015; von Krogh 2018). Not only consumers but also organizations are increasingly interested in intelligent systems, for example to recommend financial products, process transactions or schedule complex logistics (Bamberger 2018). Intelligent systems are generally classified into weak or strong AI, depending on the scope of their task. While the research area of strong AI aims to create a general human-like intelligence, weak AI refers to the solving of a specific problem (Kurzweil 2005). So far, all intelligent systems can be assigned to weak AI. Here, ML is increasingly used as the underlying technology, frequently combined with more classical approaches such as searching or planning algorithms (Russell and Norvig

2016; Silver et al. 2017). Intelligent systems based on ML are software programs that increase their performance with respect to a particular task by gaining more experience (Mitchell 1997). In this context, experience corresponds to the data (e.g., images, numbers, texts) that is used for training (Crowston and Bolici 2019; von Krogh 2018). Intelligent systems thus differ from conventional systems as they do not rely on prespecified instructions of developers, but develop internal representations based on patterns observed in the training data that are used for making predictions (Burrell 2015). As a result, intelligent systems have the ability to learn from user's behavior, deduce autonomously and react to their environment; characteristics that are usually reserved to humans (Rai et al. 2019). Interacting with intelligent systems can therefore lead users to perceive humaneness and be threatened by the system (Rzepka and Berger 2018). Furthermore, by learning progressively to provide users with predictions, intelligent systems can exhibit black box behavior, especially when relying on complex statistical ML methods (Adadi and Berrada 2018). Therefore, the attitude of users towards intelligent systems is significantly determined by the transparency of their underlying process for making predictions. In the following, we provide an overview of both computer science and information systems perspectives on transparency.

6.2.2 *Computer Science Perspective on Transparency*

In computer science (CS) literature, transparency research can be grouped under the larger research stream of intelligent systems alignment. Research in this area aims to ensure alignment between intelligent systems and human interests, mainly by developing mechanisms to facilitate algorithmic accountability (Chakraborti et al. 2019; Diakopoulos 2016). In general, alignment approaches can be categorized into cooperative and adversarial approaches. Cooperative approaches intend to facilitate alignment through understanding system behavior, in contrast to adversarial approaches that are more focused on avoiding misalignment, e.g., by introducing safety measures (Chakraborti et al. 2019). Transparency mechanisms are commonly categorized as a cooperative approach to create intelligent system alignment. Moreover, it is widely assumed that demand for transparency arises from a mismatch between the formal objectives of intelligent systems and the real-world costs occurring in a deployment setting (Doshi-Velez and Kim 2017; Lipton 2016). Studies in CS literature often presume that transparency serves as a proxy for adoption drivers such as perceived trust, fairness and user satisfaction (Diakopoulos 2016; Doshi-Velez and Kim 2017). However, these relationships are typically not empirically confirmed. Other theoretical work in technical literature also covers limitations of the transparency ideal in intelligent systems. According to these studies, technical

limitations especially occur in systems whose inherent complexity is naturally hard to grasp for humans (e.g., deep learning systems). Moreover, system predictions are commonly state- and thus time-dependent, requiring transparency features to store previous versions of the intelligent system in order to explain predictions at an earlier point in time (Ananny and Crawford 2018). Furthermore, some authors argue that system transparency might (1) inhibit user privacy and (2) enable gaming of the system (Ananny and Crawford 2018; Veale et al. 2018).

Further work in CS literature aims to establish guidelines about which parts of intelligent systems should be transparent and how the corresponding features should be designed. Diakopoulos (2016) names the following categories that might be disclosed: (1) human involvement (e.g., purpose of deploying system, responsibility inside organization), (2) utilized data (e.g., quality measures such as completeness and timeliness, conducted pre-processing steps, accessibility, privacy implications), (3) model characteristics (e.g., statistical assumptions, input features, weights), (4) inferencing (e.g., performance measures, confidence in predictions). Regarding the design of transparency features, current research in computer science mainly aims to develop explanation facilities for statistical ML algorithms. Transparency features typically serve one of two purposes: (1) visualizing the inner workings of learned ML models (e.g., intermediate representations), and/or (2) providing intuitions about the rationale behind individual predictions (local interpretability) or model workings (global interpretability; Guidotti et al. 2019). Recent studies concerning model visualization mainly aim to develop transparency facilities for deep neural networks which have achieved state-of-art performance on tasks such as image recognition or language modeling and are naturally hard to understand for humans because of their inherent complexity (often containing millions of learnable parameters). Common approaches intend to visualize the learned representations in hidden layers of neural networks, e.g., they provide intuitions about which specific objects are detected by kernels in a convolutional neural network (Carter et al. 2019). In natural language processing, similar techniques are applied to visualize how recurrent neural networks memorize characteristics of long input sequences for making predictions, e.g., about the next word in a sentence (Madsen 2019). Regarding the explanation of single predictions and model workings, most approaches use attribution techniques, which allow to measure the relative importance of input features as well as features detected in intermediate representations with regard to the model output (Olah et al. 2018). Techniques like LIME and SHAP explain predictions by quantifying the influence of each input feature and are well established in practice, especially for models working on tabular data (Lundberg and Lee 2017; Ribeiro et al. 2016). Both techniques work for any supervised ML algorithm and are thus called model-

agnostic approaches. Early stage research aims to provide similar mechanisms for other data types and more complex neural networks, e.g., by relating human-understandable concepts to features learned in opaque systems and quantifying their importance for predictions (B. Kim et al. 2018).

6.2.3 *Information Systems Perspective on Transparency*

Transparency of intelligent systems has also been a theme of information systems (IS) literature. Early work on transparency mechanisms dealt with Knowledge-Based Systems (KBS). Although these systems oftentimes do not rely on statistical ML methods and thus do not contain self-learning mechanisms, we include prior studies from this research area as some theoretical insights derived from transparency research in KBS should be relatable to modern ML-based systems. In KBS, transparency is mainly achieved via different types of explanations, e.g., rule traces or justifications (Ye and Johnson 1995). Implementation of these explanation facilities is based on established decision-making theory, e.g., Cognitive Effort Perspective and Toulmin's Model of Argumentation (Dhaliwal and Benbasat 1996; Gregor and Benbasat 1999). Furthermore, Dhaliwal and Benbasat (1996) develop a theoretical framework for empirical evaluation of user interaction with explanation facilities, incorporating user perceptions, learning effects and performance measures for judgmental decision-making. Subsequent studies rely on this framework and show that use of explanations improves performance on specific tasks in a cooperative problem solving setting (Gregor 2001). Moreover, differences between requirements of expert and novice users are established. Whereas novice users more often use explanation facilities for learning, expert users mainly employ these mechanisms in order to verify conclusions (Mao et al. 2019). More recent IS research examines explanations in recommender systems which more often utilize modern statistical ML models. For recommender systems, integration of explanation facilities is shown to increase decision efficiency and effectiveness, perceived transparency and user satisfaction (Gedikli et al. 2014). The authors also examine differences between multiple transparency mechanisms, varying degree of personalization and integration of content data. They find that non-personalized content-based explanations most positively influence decision effectiveness, whereas more efficient interfaces (e.g., average ratings) result in increased decision efficiency. Another study shows that perceived transparency of recommendation agents positively influences perceived usefulness and enjoyment, mainly through positive effects on affect- and cognition-based dimensions of trust (Wang et al. 2016). A positive relationship between perceived transparency and trusting beliefs is confirmed by Wang and Benbasat (2016). They

find that perceived transparency positively influences trust in competence, benevolence and integrity of recommendation agents. Although these findings are valuable for our research context, transparency research on intelligent systems is rare so far. Recently, Sidorova and Rafiee (2019) name lack of algorithmic transparency as an inhibitor of AI adoption. They explicitly list undefined data sources and the opaque, complex nature of modern algorithms as problematic aspects regarding this specific technology. Only a few studies explicitly address how transparency features for intelligent systems should be designed (Chai and Li 2019; Fernandez and Provost 2019; Martens and Provost 2014).

6.3 Theoretical Framework

Our study can be divided into two parts, as we aim to find evidence for (1) consumer's overall WTP for transparency features of intelligent systems (RQ1) and the underlying mechanism to establish their purchase intention (RQ2). Therefore, in the first subsection we discuss why we expect consumers to exhibit WTP for transparency features for intelligent systems. The second subsection then introduces our research model for explaining the mechanisms behind the purchase decision.

6.3.1 Willingness to Pay for Transparency Features

As mentioned in the introduction, consumers are oftentimes left in the dark regarding how predictions about personal characteristics (e.g., creditworthiness) are made by intelligent systems. In most cases, the utilized data and details about the applied algorithms are not disclosed to the consumer, thus creating an information gap between the consumer and the provider of the intelligent system (Burrell 2015; Martens and Provost 2014; Miller 2019). We argue that transparency features would constitute an added value for consumers because they provide a way to close this information gap. In detail, transparency features could include information about the intelligent system provider (e.g., certifications), utilized data (e.g., sources, pre-processing steps), model characteristics (e.g., applied algorithm, input features, feature weights) and inferencing (e.g., performance, confidence in predictions; Diakopoulos 2016). Having established the interest consumers should have in transparency features, the question remains why they would also be willing to pay for them. Here, we build on findings from privacy research, where examining how consumers value their personal data has been a theme for many years. Multiple studies show that consumers put meaningful monetary value on their personal information (Grossklags and Acquisti 2007; Krasnova et al. 2009; Schreiner and Hess 2015; Wagner et al. 2018). To the best of our knowledge there are no studies

examining WTP in the context of intelligent systems. However, as intelligent systems frequently process personal information and their predictions are oftentimes related to personal characteristics, we assume that consumers will exemplify comparable behavior as in privacy contexts. Based on these observations, we expect that consumers are willing to purchase transparency features for an intelligent system that judges their personal characteristics.

6.3.2 *Mechanisms Behind Purchase Decision*

Our study aims to explain whether and why consumers are willing to pay for transparency features in a finance context. In order to account for consumer's actual intention to purchase these features, we draw inspiration from Theory of Planned Behavior (TPB). TPB postulates that the behavior of individuals is induced by behavioral intentions, provided that the intention itself results from the individual's attitude towards the behavior as well as from subjective norms and perceived behavioral control regarding the according behavior (Ajzen 1991). We mainly use TPB as a starting point because this theoretical framework can be flexibly adapted and offers the possibility to add new variables (Venkatesh et al. 2003). With the goal of keeping our theoretical model concise we focus on relevant constructs from the TPB framework. To the best of our knowledge, this study is the first relating TPB to consumer's intention to purchase transparency features in intelligent systems. However, TPB has previously been used in WTP and purchasing contexts, making it suitable for our study as well (George 2004; Hansen et al. 2004; Schreiner and Hess 2015). Because of our focus on explaining purchase behavior, traditional theoretical models with a focus on adoption, e.g., the Technology Acceptance Model (TAM) or Unified Theory of Acceptance and Use of Technology (UTAUT), are not suitable for this study (Davis et al. 1989; Venkatesh et al. 2003).

Figure 11 shows our research model, including all relevant model constructs and control variables. In accordance with previous work on transparency in similar contexts, we define *perceived transparency of intelligent system (PT)* as the consumer's capability to understand the inner workings of an intelligent system, including assumptions and characteristics that determine its outputs (Wang and Benbasat 2016). Moreover, we examine influences of *trust in the intelligent system (TS)* onto *attitude towards transparency features (AT)*, which we define as the degree to which a consumer favorably or unfavorably assesses transparency features. Also adopted from TPB, we use *intention to purchase transparency features (IN)* as our main predicted variable. Finally, we include the two remaining TPB constructs as control variables: *subjective norm (SN)* resembles perceived social pressure to purchase transparency features and

perceived behavioral control (PBC) is defined as perceived ease or difficulty of purchasing transparency features.

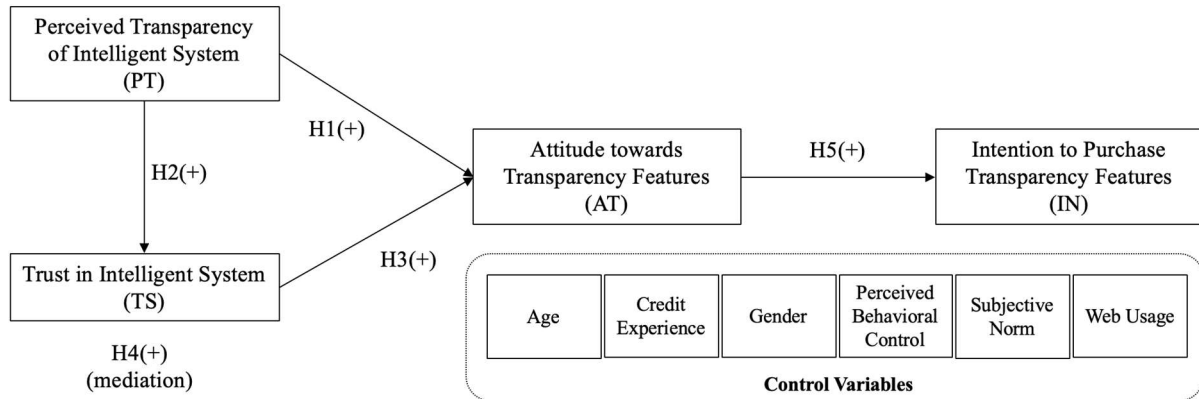


Figure 11: Research Model for Mechanisms Behind Purchase Decision

In order to examine the mechanisms that underlie this behavior, we rely on TPB and establish antecedents for consumer's attitude towards the transparency features. Increased perceived transparency in the form of additional insights about inner workings of a system constitute the main benefit for consumers when evaluating whether to buy transparency feature packages. Hence, we hypothesize that perceived transparency will have a positive influence on consumer's attitude toward transparency features. Furthermore, we expect this effect to be mediated by trust into the intelligent system as a whole. The more information consumers receive about how a system works, the more they trust outputs of the respective system. This assumption can be justified by previous studies in transparency research for related contexts (Wang et al. 2016; Wang and Benbasat 2016), and comparable work regarding the evaluation of privacy-enhancing features (Schreiner and Hess 2015). An increase in trust into the intelligent system should then lead to a more positive evaluation of the transparency features, as assessed by studies examining different contexts (George 2004; Schreiner and Hess 2015; Suh and Han 2002; Wu and Liu 2007). All in all, we postulate the following hypotheses:

H1: *PT will positively affect consumer's AT.*

H2: *PT will positively affect consumer's TS.*

H3: *TS will positively affect consumer's AT.*

H4: *PT will positively affect consumer's AT because of increased TS.*

Following the TPB framework, we expect that a positive attitude towards transparency features constitutes the main predictor for their intention to purchase these features. Previous studies that apply TPB to the related context of internet purchases confirm this assumption (George

2004; Hansen et al. 2004). Furthermore, we control for subjective norm and perceived behavioral control, which are additional predictors for consumer's intentions according to the theory (Ajzen 1991). Therefore, we assume that:

H5: *AT will positively affect consumer's IN.*

6.4 Research Design and Data Collection

6.4.1 Study A: Measuring Willingness to Pay

Following the recommendations of Karahanna et al. (2018), we conducted an online lab experiment since transparency of intelligent systems has not been monetarized yet and new insights into the evaluation of this technology feature are still necessary. By using this method, we ensured the highest degree of internal validity possible as the boundary conditions of an experiment can be determined precisely (Karahanna et al. 2018). In the beginning, the respondent was briefed that she/he needs to apply for credit of 50,000 € (corresponds to mean credit amount in Germany). Furthermore, respondents were told that their financial institution employs an intelligent system, i.e., AI which collects and evaluates online data from sources like corporate websites, civil registers and credit agencies. This approach is in line with real rating providers such as Branch, Kreditech or FICO, that use ML algorithms to specify credit scores (Branch 2019; FICO 2019a; Kreditech 2018). All respondents received an initial credit score of 65 points (scale 1-100, 100 = highest creditworthiness). This score was based on a pre-study and ensured that participants were not certain about the final decision of their financial institution, i.e., credit commitment or rejection (see Table 10).

Table 10: Overview of Pre-Studies and Experiment Parameters

	N	Parameter Name	Measure	Parameter Value
Pre-study 1	56	Initial credit score	Median of stated uncertainty range	65 points
Pre-study 2	60	Threshold price	Median of stated WTP	15 €
		Budget	90 th percentile of stated WTP	50 €

Once the initial score had been provided, participants were offered to purchase additional transparency features. In order to create a realistic setting, the participants received a mock-up containing an example output of the transparency features. The features were designed according to established practices from CS literature (esp., Diakopoulos 2016; Lundberg and Lee 2017; Ribeiro et al. 2016) and findings from social sciences about human perception of explanations (Gedikli et al. 2014; Gregor and Benbasat 1999; Miller 2019). In detail, the transparency features contained information about collected features based on online data from

various sources (e.g., job status based on data from corporate website and professional networks) as well as the relative importance of each feature for the final prediction on creditworthiness (e.g., 29%). All participants received a fictional budget of 50 €, which was determined based on a second pre-study (see Table 10). We simulated a real-world purchase decision, i.e., respondents had to evaluate costs compared to the personal value they attach to the transparency package. In order to precisely approximate consumer's real WTP, we employed the Name-Your-Own-Price (NYOP) method with multiple bidding rounds, which is an established method for measuring WTP (Breidert et al. 2006) and has been applied in various contexts (e.g., Chernev 2003; Hann and Terwiesch 2003; Spann et al. 2004). Therefore, the applicability in our credit scoring setting could be assumed. According to the method, respondents had to bid for the additional package against an unknown, but fixed threshold price of 15€ (see Table 10) by submitting bids between 0€ and 50€. Moreover, respondents of the study were informed about a total of three bidding rounds that were interrupted as soon as the threshold was surpassed (experiment process is summarized in Figure 12). By including a repeated bidding option, participants were able to adjust to the situation and feel more comfortable as they could include feedback in their bidding behavior and thus overcome the lack of experience concerning the monetization of transparency features (Liu et al. 2016). As NYOP is not an incentive-compatible method on its own, we created an incentive for the respondents by promising them a variable payout based on their choices during the experiment (Breidert et al. 2006).

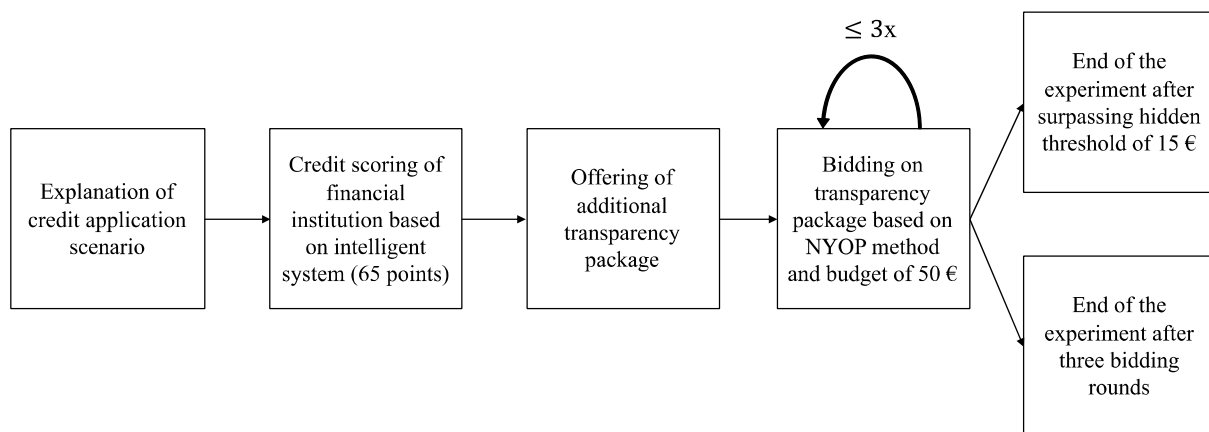


Figure 12: Experiment Procedure

In order to set all experiment parameters in the first place, two pre-studies were performed (see Table 10). The first pre-study was conducted to determine a credit score, where it is not obvious whether a credit is granted or not. Therefore, respondents were asked to indicate two values between 0 and 100 points, for which they feel (1) confident that they will be granted a credit

and (2) confident that they will be rejected for a credit. This procedure was intended to prevent possible distortions within the experiment that could arise from participant's knowledge about credit acceptance or rejection. In the second pre-study, we employed a direct survey to determine the fictitious budget communicated to the experiment participants as well as the unknown and fixed threshold to bid against. In this regard, we asked respondents of the pre-study to directly report their WTP for a transparency package within a credit scoring scenario. As described above, we assume that the participants of pre-study 2 did not reveal their true WTP (e.g., due to prestige effects). Therefore, the results provide a first guidance to parameter settings but should be critically examined by conducting the experiment and employing the NYOP method.

6.4.2 Study B: Explaining Purchase Decision

After the WTP had been measured by the experiment described above, participants were tasked with completing a survey. Here, we assessed constructs related to our presented research model regarding the purchase decision. Only adapted standard scales were used that are known from and validated in extant TPB literature or that could be derived from the literature concerning system transparency. Suitability of the applied constructs was positively evaluated using a subset of items in the second pre-study conducted with 60 students. In the resulting survey, reflective latent measures were used, operationalized on a seven-point Likert scale: IN, PBC, PT, SN and TS ranged from 1 (strongly disagree) to 7 (strongly agree), while AT was implemented as a semantic differential evaluated from 1 (bad, foolish, unpleasant, dislike) to 7 (good, wise, pleasant, like). Table 11 shows exemplary items of our survey constructs.

Table 11: Examples of Construct Operationalization

Construct	No of Items	Example of Items	Source
PT	6	With the additional transparency features it gets readily apparent to me how the algorithm generates its prediction.	Wang and Benbasat (2016)
TS	5	The intelligent system keeps my best interests in mind.	Koufaris and Hampton-Sosa (2004)
AT	4	Purchasing the additional transparency features is a bad/good idea.	Taylor and Todd (1995)
IN	3	I intent to purchase the additional transparency features.	Venkatesh et al. (2003)
PBC	3	I would be able to purchase the additional transparency features.	Taylor and Todd (1995)
SN	2	People who influence my behavior would think that I should purchase the additional transparency features.	Taylor and Todd (1995)

6.4.3 Sample

The previously described studies were conducted within Germany in March 2019 in cooperation with a market research institute. A total of 223 participants completed the experiment and the survey, resulting in an overall response rate of 76.9%. While selecting the participants, quotas were taken into account in order to reflect the age and gender distribution of consumers using online applications (Eurostat 2018). The sample consisted of 46.2% female and 53.8% male participants and included a wide range of age groups (18-68 years) resulting in a mean of 38.8 years (SD: 12.6). The majority of respondents were salaried employees (58.5%) and had experience in applying for a credit (64.1%). 79.5% of participants indicated that they use the internet very frequently. Therefore, an appropriate demographic distribution was achieved concerning the setting of both experiment and survey. Hence, our sample allows us to draw inferences concerning our hypotheses for consumers from the general public, at least for the country under study. From this sample, we had to filter 28 participants that failed to pass two included attention checks or showed unengaged behaviour, resulting in a final sample of 195 cases for our analyses.

6.5 Data Analysis and Results

6.5.1 Results of Study A

In order to answer RQ1, i.e., to what extent consumers exhibit WTP for transparency features for intelligent systems, we examined the experiment outcomes. We determined WTP for each respondent by extracting her/his maximum bid out of all three rounds (descriptive statistics are listed in Table 12). We observed a mean WTP value of 21.57€, with a standard deviation of 16.11. Out of 195 respondents, 142 crossed the threshold price of 15€ (72.8%). Thereof, 95 participants surpassed the threshold price in round 1 (48.7%), 38 in round 2 (19.5%) and 9 in round 3 (4.6%). Only 43 participants (22.1%) exhibited bids of 0€ in all three rounds, thus not demonstrating WTP for the transparency feature package.

Table 12: Descriptive Statistics Regarding Measured WTP

	N	Mean	Std. Dev.	25%	Median	75%	Skewness	Kurtosis
WTP	195	21.57	16.11	10.00	20.00	30.00	.31 (.17)	-.80 (.35)

6.5.2 Results of Study B

Having established that WTP exists for transparency features, we aimed to examine mechanisms in consumer's purchase decisions using our previously hypothesized research model.

Measurement Model. In order to establish that our measurement model was suitable for causal analyses, we performed exploratory factor analysis (EFA) in SPSS 26 and confirmatory factor analysis (CFA) using AMOS 25. During EFA, we dropped item PT6 due to weak factor loadings, all other items were retained. We assessed reliability and validity of constructs by examining factor loadings, Cronbach's alphas, composite reliability (CR) and average variances extracted (AVE; see Table 13 and Table 14). All model constructs showed highly positive loadings, cross-loadings were smaller than .30 (omitted here for the sake of brevity). Reliability was given, as Cronbach's alphas and CR were above .7 for all constructs, and convergent validity was ensured by AVE greater than .5. Discriminant validity was assured, as square roots of AVE were greater than factor correlations and cross-loadings were smaller than factor loadings. Model fit for the measurement model was also above accepted thresholds, CFI = .98 (> .95), SRMR = .03 (< .08), RMSEA = .05 (< .06). All thresholds are taken from Hu and Bentler (1999).

Table 13: Factor Loadings and Reliability

Construct	Items	Factor Loadings	Cronbach's Alpha
PT	PT1	.797	.942
	PT2	.791	
	PT3	.910	
	PT4	.920	
	PT5	.925	
TS	TS1	.885	.961
	TS2	.774	
	TS3	.809	
	TS4	.853	
	TS5	.911	
AT	AT1	.931	.936
	AT2	.923	
	AT3	.744	
	AT4	.754	
IN	IN1	.966	.971
	IN2	.842	
	IN3	.875	
PBC	PBC1	.931	.809
	PBC2	.631	
	PBC3	.644	
SN	SN1	1.036	.894
	SN2	.558	

Table 14: Factor Correlations, Reliability and Validity Measures

	CR	AVE	PT	TS	AT	IN	PBC	SN
PT	.943	.767	.876					
TS	.962	.834	.621	.913				
AT	.937	.787	.512	.794	.887			
IN	.972	.919	.514	.746	.785	.959		
PBC	.817	.603	.670	.443	.343	.426	.777	
SN	.900	.818	.442	.683	.737	.750	.323	.905

We tested for common method bias using collinearity diagnostics and Harman's single factor test. Regarding collinearity statistics, we looked at variable inflation factors (VIFs) for each examined relationship in our structural model. No method bias was found, as all VIFs were below the threshold of 3.3 (Hair et al. 2009). In addition, we ran Harman's single factor test and observed that no single factor accounted for more than 50% of total variance. This further indicated that occurrence of common method bias was very unlikely in this study (Fuller et al. 2016; Podsakoff et al. 2003). Finally, we imputed factor scores for subsequent path analysis using AMOS.

Structural Model. Following our evaluation of the measurement model, we tested our hypotheses by examining path coefficients and their significance (see Figure 13) using AMOS 25. In order to ensure our model's predictive relevance, we again assessed model fit. Model fit was found to be adequate, according to widely accepted thresholds, CFI = .99, SRMR = .05, RMSEA = .06 (Hu and Bentler 1999). Furthermore, we were able to account for 64 to 75% of the variance in the endogenous constructs TS, AT and IN. Regarding path coefficients, we found no support for H1 as there was no observed direct effect between PT and AT. However, we found that PT significantly influenced TS, with TS significantly affecting AT (supporting H2 and H3). Moreover, H4 is supported, as an indirect effect of PT on AT was confirmed via mediation analysis using bootstrapping (5000 samples), resulting in an statistically significant indirect-only effect over TS, $\beta = .21$, $p < .001$ (Hayes 2009; Zhao et al. 2010). We could further establish a significant effect of AT on IN, as proposed by TPB (Ajzen 1991). Consequently, we can also support H5. Figure 13 also displays statistically significant paths related to our control variables. As expected, PBC and SN positively affected IN (Ajzen 1991). SN also influenced both TS and AT in a statistically significant way. Furthermore, we found a statistically significant effect of gender on TS. Table 15 sums up our analysis of study B.

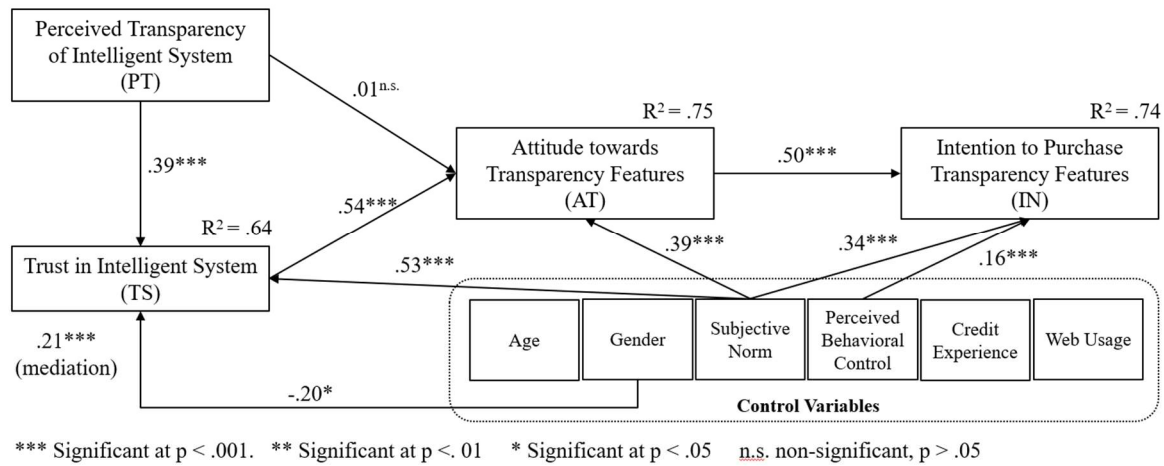


Figure 13: SEM Results

Table 15: Overview of Tested Hypotheses

Hypothesis	Relationship	Support
H1	PT (+) → AT	No, $p = .89$
H2	PT (+) → TS	Yes, $p < .001$
H3	TS (+) → AT	Yes, $p < .001$
H4	PT → TS → AT (Mediation)	Yes, $p < .001$
H5	AT (+) → IN	Yes, $p < .001$

6.6 Discussion

This study examined to what extent consumers would be willing to pay for transparency features in the context of intelligent systems (RQ1), and what the mechanisms behind the purchase decision looks like (RQ2). We conducted an online lab experiment with a subsequent survey in order to examine both research questions adequately and increase robustness of results. To the best of our knowledge, we were the first to study WTP for transparency features in intelligent systems. In addition, we developed a research model for the mechanisms regarding the purchase prediction, drawing inspiration from TPB and prior work in related areas like knowledge-based and recommender systems. For the experiment, we developed a real-world scenario in which respondents had to choose whether to purchase an additional feature package that offers insights into algorithmic predictions about their creditworthiness. WTP was measured using the established Name-Your-Own-Price procedure with a total of three bidding rounds and a hidden threshold price of 15€. All experiment parameters for the study were rigorously determined, based on results of two conducted pre-studies. Our experiment unveiled that a large majority of participants exhibited meaningful WTP (median 20€) for the presented transparency feature package. Moreover, our research model for explaining the purchase decision was shown to have high measurement quality and explanation capabilities.

Our study offers significant theoretical contributions regarding research into consumer interactions with intelligent systems. First and foremost, we established that significant WTP among consumers exists for transparency features in intelligent systems through an experiment that simulated actual purchase behavior. Furthermore, we are the first to transfer transparency research to a context different from knowledge-based and recommender systems. Our findings indicate that increasing perceived transparency leads to a positive attitude towards feature packages with this purpose. We showed that an increase in trust towards the intelligent system is the main driver behind this effect, which confirms prior research in the area of recommender systems (Wang et al. 2016; Wang and Benbasat 2016). This finding also counters the possible fear that consumers might use the derived insights from transparency features to game the intelligent system. In accordance with TPB, a more positive evaluation of transparency features positively affected the actual intention to purchase these. Additional factors influencing purchase intentions were subjective norms, i.e., social pressure, and perceived behavioral control, in accordance with TPB (Ajzen 1991). Moreover, transparency could be important to consumers not only in the context of intelligent systems, but for a variety of digital customer experiences. Thus, further transparency research should be conducted in related digital contexts.

Our results also have significant implications for practitioners. First of all, we have shown that WTP for transparency features exists for the real-world scenario of credit scoring. This constitutes a valuable insight for providers of intelligent systems and services when evaluating whether to offer comparable features related to algorithmic transparency. Our findings reveal that inclusion of transparency features might have two main benefits. First, they can be monetized separately from the intelligent service and thus constitute an additional revenue stream in the form of a premium service. Here, results from this study could serve as an entry point for price determination. We found that respondents were willing to spend on average 20€ for such a package which constitutes meaningful WTP. Although these findings are limited to the context of credit scoring, they can still be meaningful anchor points for decision-makers, e.g., product managers. Second, we showed that increasing perceived transparency has significant positive effects on trust into the intelligent system. This insight can be meaningful to practitioners, because trust has been shown to positively influence important business metrics such as customer retention (e.g., Chiu et al. 2012; Gounaris 2005; Han and Hyun 2015). Beyond WTP, our results point towards the relevance of “fair AI” for consumers. This finding is in line with previous research, as well as company and governmental reports that all name transparency as a key driver for establishing fairness in intelligent systems (Doshi-Velez and

Kim 2017; Google 2019; ITU 2018). As algorithmic transparency is seen as particularly important in the public sector (Diakopoulos 2016), results from this study can also be used for intelligent systems that are employed by public institutions.

Our study is subject to some limitations. First, we focused on examining WTP in the financial context of credit scoring. Although we expect our results to be transferable to related contexts involving intelligent systems (esp., for high-stakes decisions concerning personal characteristics), these contexts might have different outcomes associated with increasing perceived transparency. For example, WTP in these contexts might vary from what we found in this study. Second, we conducted our study within one country. Thus, we are not taking potential cultural differences regarding transparency and WTP into account but would not expect large deviations from our findings. Third, we were not able to establish a direct effect between perceived transparency and attitude towards transparency features. This is probably due to the fully mediated effect of perceived transparency on attitude over trust. Fourth, data for our structural equation model was obtained through a single method of data collection for both independent and dependent variables. Thus, a common method bias could reduce our contributions. Although we conducted statistical tests, namely Harmon's single factor test and collinearity diagnostics, that did not point towards the occurrence of common method bias, we cannot completely rule out its existence.

Based on the results from this study, we see the following opportunities for future research. First, our findings have to be examined in other contexts. Here, we expect transparency of intelligent systems to have similar importance across both B2C and B2B scenarios. In a B2B context we believe algorithmic transparency to have comparable relevance, especially in regulated industries such as banking. Second, further research into transparency approaches for intelligent systems is needed. From a technical perspective, it is important to investigate how complex ML models (e.g., deep neural networks) can be queried for human-understandable explanations. Future research should also draw from previous work in the area of knowledge-based and recommender systems in order to establish best practices for the design of interfaces for transparency features (e.g., Gedikli et al. 2014; Ye and Johnson 1995). Additionally, consumer behavior regarding transparency of intelligent systems should be observed in relation to other service attributes, e.g., predictive performance. Conjoint analysis could be a suitable research methodology for this, as it still allows to integrate WTP. Privacy research has involved similar studies which future transparency studies could draw inspiration from (Krasnova et al.

2009). Case studies with credit scoring providers, potentially in conjunction with field experiments, would be a further alternative to strengthen external control of our study.

6.7 Conclusion

Advances in AI technology have led to widespread use of intelligent systems for a variety of use cases. Oftentimes, these systems rely on modern ML models, e.g., deep neural networks, making them effectively black boxes. Thus, system characteristics like data sources, input features, statistical models and feature importance for predictions are not revealed to consumers of the system. Consequently, public and private institutions have called for transparency and fairness standards regarding intelligent systems (Google 2019; ITU 2018; The Economist 2018). However, how transparency features can be designed and are perceived by consumers remains largely an open question in both research and practice.

In this study, we investigated whether and why consumers would be willing to pay for transparency features of intelligent systems. Therefore, we conducted an online lab experiment and a subsequent survey with 195 participants in a European country, placed in the context of credit scoring. This allowed us to (1) measure WTP experimentally and (2) develop a research model examining mechanisms of the purchase decision with respect to a real-world scenario. We found that consumers exhibited meaningful WTP (median 20€) for the offered transparency feature package. Furthermore, we observed that perceived transparency of the intelligent system positively influenced the trust consumers have into it, which led to a more positive evaluation of the offered transparency feature package. Our results have significant implications for research and practice. To the best of our knowledge, we are the first to study WTP for transparency features with regard to intelligent systems. Since our theoretical framework shows high explanation capabilities, it can inform further transparency research in various B2B or B2C contexts. Moreover, practitioners can use our results as they offer a new perspective on how intelligent systems providers can monetize their services. On the consumer side, our findings also indicate the relevance of transparency in the context of digital services.

7 Paper 2.C: “Hello, I’m Here to Help You” – Medical Care Where It Is Needed Most: Seniors’ Acceptance of Health Chatbots

Title

“Hello, I’m Here to Help You” – Medical Care Where It Is Needed Most: Seniors’ Acceptance of Health Chatbots

Authors

Mesbah, Neda; Pumplun, Luisa

Publication Outlet

European Conference on Information Systems (ECIS)

Abstract

Demographic change is resulting in an increasing number of people getting older and needing adequate medical care. In order to solve this problem, easily accessible health chatbots could be established, which are capable of identifying diseases on the basis of symptoms. Such applications already exist today, but usage is primarily restricted to younger generations. Therefore, this study examines which factors affect the acceptance of health chatbots by seniors. With the help of 21 qualitative interviews within the respective target group and the extended Unified Theory of Acceptance and Use of Technology (UTAUT2) a comprehensive model for the adoption of health chatbots is developed. Additionally to being interviewed, the participants were able to test an exemplary health chatbot to better comprehend the technology. Thus, a practically oriented overview as well as 18 propositions were established, which could be used as a basis for further research regarding seniors’ acceptance of health chatbots.

Keywords

Senior, Adoption, mHealth, Health Chatbot

7.1 Introduction

The ongoing demographic change and the associated decline in comprehensive healthcare provision urgently require new developments in medical technology (Krys and Fuest 2017). A progress, that could help solve this social problem, can be found in the so-called mobile health (mHealth) area. M-health technologies “empower individuals to manage their personal health” (Kenny and Connolly 2017, p. 1129) and therefore make healthcare more efficient. Consequently, the market of mHealth is constantly growing. It is predicted to reach USD 151.57 billion by the year 2025 (Grand Review Research 2018). This is also reflected by the number of health apps available on iOS and Android. Over 100,000 apps with a medical focus can be downloaded to respective devices (Taylor 2015). A particularly promising progress in the mHealth area is the use of artificial intelligence (AI). Panetta (2018) state that AI-enhanced virtual care will help offer a more accessible, comfortable and cost-effective medical supply.

In this context, chatbots are a possible application form of AI. Chatbots are intelligent software programs that are able to communicate with users in text form by emulating natural language (Shawar et al. 2005). They are being used increasingly for health purposes, such as self-diagnosis. A particular chatbot is provided by Ada Health. Their app Ada works as a personal health guide and assesses symptoms by leading users through a directed dialogue while collecting relevant information. With the help of a medical knowledge base, patterns of symptoms are then examined for possible causes by applying intelligent algorithms (Ada Health 2019).

Especially elderly citizens can vastly benefit from such self-diagnosis tools based on AI. The term “elderly”, as defined by the United Nations, includes all persons aged 60 years and older and is used accordingly in this study (United Nations 2017). This age group is particularly affected by chronic illnesses, such as osteoporosis, cardiovascular diseases or diabetes and therefore needs comprehensive medical care (Sanyal 2018). In order to facilitate the use of digital technologies by older people, much research is already being done concerning human-computer interaction (HCI) with the elderly as a target group. This target group has highly specific requirements and drivers regarding technology usage and adoption. Various factors, such as the perceived ability to live alone, play a major role concerning the adoption decision of the elderly (Vichitvanichphong et al. 2017). Unfortunately, according to Vichitvanichphong et al. (2017), older citizens are commonly resistant to changes.

In order to help overcome this obstacle and provide the elderly with sufficient medical care, the present study attempts to determine which specific factors influence the adoption of intelligent

chatbots in healthcare (called health chatbots) by older people and integrate the findings into a clear overview. Our approach thus aims to identify:

RQ: *Which factors influence the decision and ability of the elderly to adopt chatbots for self-diagnosis?*

To answer this research question, an overview of related literature (i.e., chatbots in healthcare) is provided first. Afterwards, we present the theoretical background of our study, which we base on the extended Unified Theory of Acceptance and Use of Technology (UTAUT2). Afterwards, our study and the sample comprising 21 interviews with 23 interviewees is introduced. Our study is based on the application Ada, which can be used for self-diagnosis. Based on a first guided experience with Ada, participants are questioned and the empirical results are discussed and integrated into a model to reflect the specific needs of the elderly. The results of our paper provide a first overview of the age-specific factors (e.g., need for emotional support) that determine the adoption of chatbots in the healthcare sector. Thereby, our findings contribute to current research on the adoption of chatbots and show which factors must be taken into account, especially considering the continued progression of current demographic changes. Finally, we conclude our paper by pointing out possible limitations of the study and presenting directions for future research.

7.2 Related Research

Recent technological advances in the field of AI, machine learning and natural language processing (NLP) have renewed the interest of research and business in conversational interfaces (Seeger et al. 2018). Chatbots, as machine conversation systems that interact with human users through natural language, represent a specific application form of conversational interfaces (Shawar et al. 2005). Chatbots are text based systems that emulate human-to-human chat conversations and are therefore often perceived as anthropomorphic (Seeger et al. 2018). They can be differentiated in general-purpose or domain-specific bots according to their scope of application (Gnewuch et al. 2017). Possible specific purposes include the use in museums (e.g., Kopp et al. 2005), e-commerce (e.g., Qiu and Benbasat 2009), or in the healthcare sector (e.g., Minutolo et al. 2017).

Especially health chatbots (HC) have a long history, starting with Eliza, an early therapeutic chatbot based on a keyword search and stated rules (Weizenbaum 1966). In accordance with Wang and Siau (2018), we define HC as programs that are able to conduct intelligent conversations regarding health issues. Within the healthcare sector, chatbots can help solve a

variety of problems. For example, they may be used to recommend preventive health measures (Amato et al. 2017). However, they can also assist in the follow-up care of patients by reminding users to take their medication, tracking users’ health trend or finding the nearest pharmacy (Wang and Siau 2018). In this paper, we focus on HC which are able to identify possible diseases based on entered symptoms (Minutolo et al. 2017). In a society facing demographic change and the obligation to deal with high life-expectancies as well as urbanization, medical supply is becoming increasingly costly (Calero Valdez and Ziefle 2019). Medical processes must therefore increase in efficiency. One solution could be to examine patients by intelligent machines in advance and decide which diseases require medical treatment (Hoermann et al. 2017). This allows clinicians to spend more time on real emergencies (Bibault et al. 2019). Therefore, the usage of self-diagnosis HC could “likely become the first point of contact for primary care” (Wang and Siau 2018, p. 1). Real products like Ada, Babylon Health or Your.Md prove a first feasibility of such systems (Ada Health 2019; Babylon 2019; Your.MD 2019). Nevertheless, various factors could affect the extensive adoption of HC. First, the functionality of intelligent chatbots relies heavily on provided data. Therefore, the quality as well as privacy of this data is crucial (Wang and Siau 2018). Another aspect could be the lack of trust in the recommendations of HC that could arise from communicating with a dehumanized entity (Amato et al. 2017). Additionally, Fadhil and Schiavo (2019) state that HC should take into account users’ demographics. Although these factors clearly influence users’ decision of adoption, little is known about the respective needs of different user groups interacting with HC. To the best of our knowledge, there is only one contribution that considers age specific adoption factors of HC from a more theoretical perspective and integrates them into a comprehensive model. This study concentrates on students with a mean age of 24.8 years. This age group may more likely belong to the group of early adopters, but has a much lower need for and easier access to medical care than the elderly. Therefore, the authors call for further studies, in which senior citizens are considered (Laumer et al. 2019).

In order to meet this demand, we are building our study regarding the technology adoption of elderly people on the extended Unified Theory of Acceptance and Use of Technology for the consumer context (UTAUT2). UTAUT2 is the most recently developed and discussed model to study technology acceptance from a consumer’s point of view. As we investigate seniors’ acceptance regarding HC, we consider UTAUT2 as particularly helpful in order to examine the individual adoption of a specific consumer group, i.e., the elderly (Venkatesh et al. 2012). Moreover, we chose to use UTAUT2 as the underlying model for our research as it was used

in pertinent contexts, namely healthcare and seniors’ technology acceptance (Chen et al. 2014; Laumer et al. 2019). In the following, we thus explain the UTAUT2 model in more detail (see Figure 14).

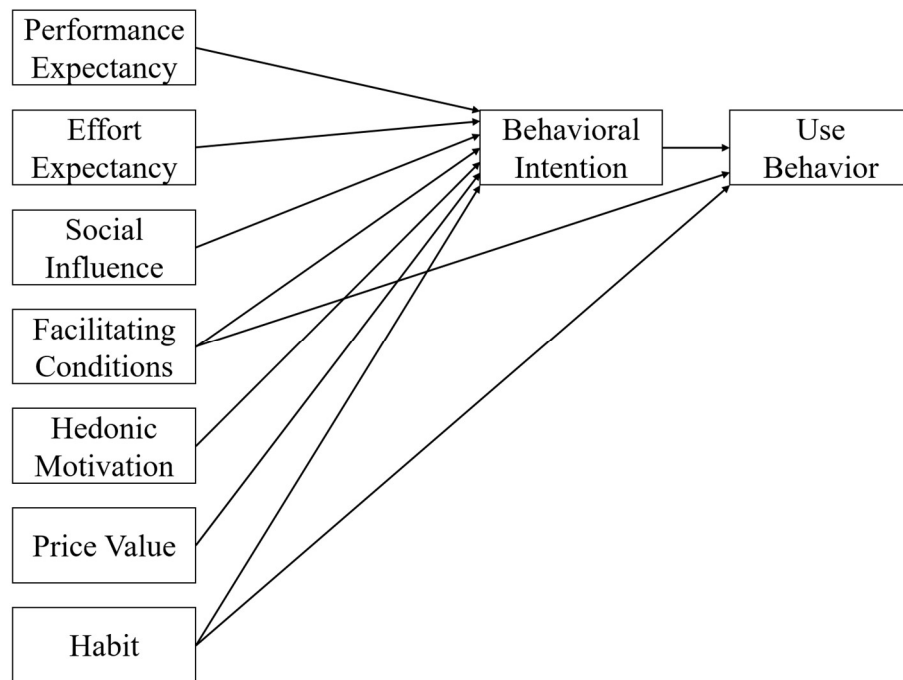


Figure 14: UTAUT2 Model (based on Venkatesh et al. 2012)

Performance expectancy is “the degree to which using a technology will provide benefits to consumers in performing certain activities” (Venkatesh et al. 2012, p. 159), whereas effort expectancy is defined as “the degree of ease associated with consumers’ use of technology” (Venkatesh et al. 2012, p. 159) – or in other words: how easy is it for someone to use the technology studied? Social influence describes the impact of people important to the deciding person (e.g., family, friends) on their technology usage. Facilitating conditions “refer to consumers’ perceptions of the resources and support available to perform a behavior” (Venkatesh et al. 2012, p. 159). Hedonic motivation refers to the enjoyment or pleasure a person receives from using a technology, whereas the price value represents the result of the cognitive trade-off between the perceived benefits and monetary costs of a technology. Habit is a construct that indicates the extent to which the beliefs and behavior of a consumer is described as automatic, since they have become a routine over time. The factors described influence the behavioral intention to use a technology, which serves in turn as a key predictor of the actual use behavior. The dependent variable use behavior is furthermore directly affected by the facilitating conditions and habit (Venkatesh et al. 2012).

The basic UTAUT2 model describes different factors that determine the acceptance and use of technologies by consumers in general. Since HC are applied in a highly sensitive context (i.e., healthcare) and require a very specific and new use behaviour compared to other technologies, it is crucial for researchers to reflect on these particular artefacts (Davison and Martinsons 2015). Moreover, various studies have shown that seniors pose their very own needs in terms of complex technologies (Vichitvanichphong et al. 2017). Both, the novelty of technology and specific requirements of seniors make it necessary to review, adapt, and extend the basic model (i.e., UTAUT2) according to the specific research context.

7.3 Qualitative Research Methodology

The aim of the study is to close the current research gap regarding the acceptance of HC by older people. Since the use of HC is a new, little researched topic, an explorative approach was chosen in which the target group of seniors is interviewed qualitatively (Flick 2004). To do so, a three-stage qualitative approach was used (Elliott and Timulak 2015). At the beginning, the current state of literature was searched with regard to technology acceptance of older people as well as adoption in a HC context (see Section 7.2). Based on the achieved conceptual foundation and the principles of Sarker et al. (2013), a semi-structured interview guideline was established, which was used to conduct in-depth qualitative interviews with participants. The research process was concluded by transcribing, coding and analyzing the interviews with the aim of showing senior-specific adoption factors for HC.

7.3.1 Research Design, Sample, and Data Collection

The semi-structured interview process comprised three different phases which were accompanied by the interview guide. The first phase included general questions on demographic factors, personal technology affinity, and health conditions. During the second phase, participants were encouraged to use the Ada app, which served as a representative example to test the handling of HC. While using the application, participants were provided help if necessary and were animated to share their experiences for the recording. By actually using the app, it was ensured that participants could better assess the advantages and disadvantages of HC. In the last and most comprehensive part of the interview process, questions about the acceptance of such chatbots were asked. Only open questions were demanded (e.g., “Would you use Ada and why?”) in order to give the participants sufficient freedom to describe specific acceptance factors not yet included in UTAUT2. Due to the nature of the semi-structured approach, interviewers were able to make further inquiries and take up

new aspects mentioned by the participants (Bryman 2016; Myers 2013). In this context, we used the laddering technique to ask successive “why” questions (Reynolds and Olson 2001).

The qualitative data was collected over a two-week period and took place in May 2019. This timeframe was chosen to ensure that all participants used a consistent version of the app. In total 21 interviews with 23 participants from our target group “senior citizens” (60+ years) were conducted within Germany. Among the seniors were 12 female and 11 male participants, whose ages were between 60 and 96 years (mean = 71 years). 20 respondents were already retired, while 3 participants were still working. Half of the participants would describe themselves as technology affine, although none of them had used a chatbot before. Nevertheless, the relevance of HC was high for the selected sample, as many of the participants already had health restrictions and were therefore forced to see the physician frequently. After the 21st interview, data collection was discontinued as only redundant aspects arose in the interviews (i.e., theoretical saturation was assumed; Flick 2004). The interviews lasted an average of 38 minutes and were exclusively conducted face-to-face to enable the active use of the app and to take into account the complexity and sensitivity of the topic. The interviews took place in closed rooms, guaranteeing freedom from disturbance and external influences. All participants took part in the interviews on a voluntary, non-paid basis and have been assured of their anonymity.

7.3.2 *Content Analysis*

As recommended by Weber (1990), the assessment of interviews was based on content analysis, which is particularly applicable to the investigation of open-ended questions. Following the steps of content analysis, UTAUT2 was used as a foundation to conduct the study. To include both the known factors as well as new aspects mentioned by seniors, a combination of directed and conventional analysis was used. While the directed approach accounts for constructs known from relevant literature (i.e., UTAUT2), conventional analysis enables to critically evaluate, extend or subdivide initial factors based on new findings from the qualitative data (Hsieh and Shannon 2005). To make this information available, all interviews were recorded in consultation with participants and transcribed immediately after conducting the interviews in order to ensure that no relevant content was lost. Afterwards, transcripts were coded using the NVivo 12 software. As recommended by Saldaña (2009), coding was conducted via two constitutive coding cycles. The first cycle included a combination of attribute, hypothesis and descriptive coding. Attribute coding was conducted to gain insights about demographic data (e.g., participant’s age, gender). The hypothesis coding took into account the initial factors of UTAUT2. Afterwards, descriptive coding was used to summarize relevant passages of the

derived qualitative data, which were specifically reflecting the needs of seniors in context of HC. During the second coding cycle, pattern coding was used to evaluate the previously generated codes critically and combine them into a smaller number of categories. In the sense of an investigator triangulation, achieved factors were discussed in a group of four Information Systems (IS) researchers and students.

7.4 Results and Discussion

Within our analysis, we were able to confirm the applicability of the UTAUT2. Nevertheless, the factors as defined in the seminal work by Venkatesh et al. (2012) are not sufficient to explain the acceptance and use of HC by seniors. Therefore, we have extended the model by redefining existing factors, adding new ones, and excluding factors which could not be identified as relevant within the interviewing process (i.e., hedonic motivation). Among the most mentioned new factors were the need for emotional support (14 of 21 interviews), technology self-efficacy (13 of 21 interviews) and medical history (11 of 21 interviews). Furthermore, the relevance of the factor price value could not be examined based on the Ada app, as it is made available to users free of charge. The result of our study is shown in Figure 15. In the following we will explain and discuss our findings in more detail.

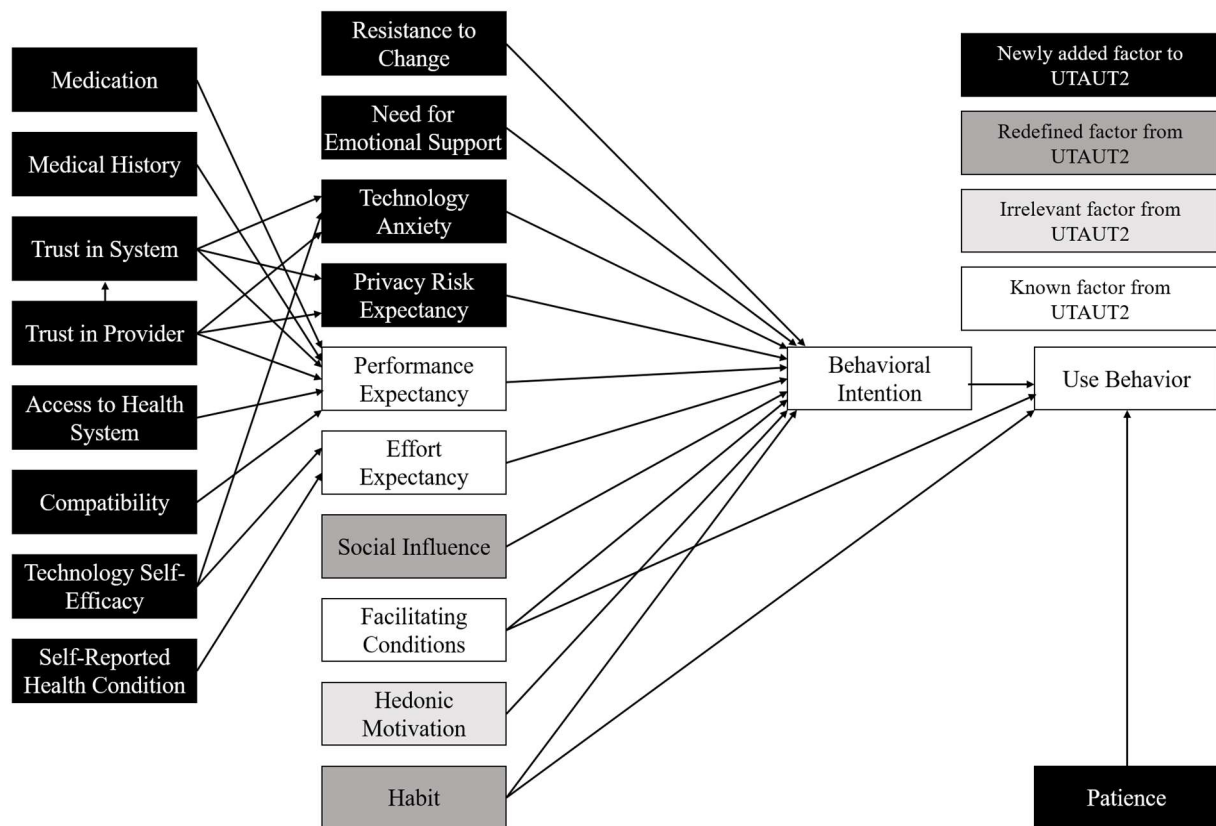


Figure 15: Seniors' Acceptance of HC – An Extended UTAUT2 Model

7.4.1 Known UTAUT2 Factors

In the following, the factors adopted from the UTAUT2 model are discussed in more detail.

7.4.1.1 Performance Expectancy

In line with existing literature on technology acceptance, the interviews have shown that the usage intention of seniors is strongly dependent on the *performance expectancy* regarding the HC. Seniors said that the usage of HC is time-saving compared to visiting the physician and allows them to quickly obtain information to assess how urgent a visit to the physician is or to alleviate worries of a serious condition. However, it was further stated that from their point of view, this cannot replace the first visit to the physician. In the eyes of the elderly, HC are a kind of supporting tool in case a physician is not immediately available or to follow up their past visit to the doctor. In this context, seniors would rather use the tool to gain more knowledge about the diagnosed disease and to use this information as a basis for further conversations with the physician. This is illustrated by the following statement:

“Also to be able to understand the statement of the physician is quite nice [...], because the visit at the physician happens quickly in a relatively short time. It is not necessarily always comprehensible [...]. For this it [i.e., Ada] can help very well.” #5

Another aspect that is part of the performance expectancy of seniors towards HC is the ability of the chatbot to respond to their individual needs, similar to what a physician would have done:

“No, I don’t want to be informed about a machine first. I want to see a physician. [...] Because it [i.e., Ada] should be more individualized, but I don’t know how that would work. As I said, because there are so many different people out there.” #11

There is a wide variety of research in the field of HCI that shows that technology should fit to the individual needs of seniors in order to encourage a usage intention (e.g., M. Kim et al. 2018): For example, their medical demands change over time and technology should be flexible enough to adhere to these new circumstances (Gao and Koronios 2010). Hence, we have formulated the following proposition (**P1**):

The higher seniors’ performance expectancy, the higher is seniors’ intention to use HC.

7.4.1.2 Effort Expectancy

Another crucial factor confirmed by the interviews is *effort expectancy*. It is crucial for seniors to be able to use the HC with little effort. For example, the handling of the HC should be

preferably simple and the answer design within the dialogue should be sufficient, precise, and easy to understand. In this regard, some of the possible answers and explanations of the respective HC were not comprehensive enough for the participants to be able to understand their intention and to answer appropriately, as can be seen from the following comment:

“The explanation was very poor. For example, ‘at the hollow of the knee’. I wanted to know if I understood that correctly. The hollow of the knee is what is directly below the knee, isn’t it? [...] And then he just told me again, ‘yes, pain in the hollow of the knee’. But I would have liked [...] to receive more information.” #12

All in all, the use should be simple and answers should be easily understandable for seniors in order to increase their willingness to use the HC. Therefore, our second proposition is **(P2)**:

The higher seniors’ effort expectancy, the higher is seniors’ intention to use HC.

7.4.1.3 Facilitating Conditions

Another challenge to accept and use HC, which is critical for seniors, is to own a smartphone and have access to the internet. These requirements were mentioned by several interviewees, showing the relevance of this topic for seniors. Therefore, *facilitating conditions*, as already shown in the literature (Venkatesh et al. 2012) have a particularly high relevance for the usage intention and behavior of seniors regarding HC. However, the technological requirements are not the only challenge that seniors face when intending or actually using HC. This segment of the population often needs appropriate and long-term assistance in order to use digital systems and reap their full benefits (Chen et al. 2014). This is particularly the case when it comes to novel and complex technologies, such as HC. With the assistance of other people, seniors feel more confident to accept and finally use an HC:

“So now, as you are here and explain it to me, I think it’s easy. But I don’t [...] know what it would be like, if I was left on my own.” #4

Based on the pertinent literature and our interviews we thus state **(P3)**:

a: The better seniors’ facilitating conditions, the higher is seniors’ intention to use HC.

b: The better seniors’ facilitating conditions, the higher is seniors’ actual use of HC.

7.4.2 Redefined UTAUT2 Factors

The factors social influence and habit known from the original UTAUT2 had to be slightly redefined.

7.4.2.1 Social Influence

In agreement with the research on technology acceptance, our interviews show that the social environment has a large influence on seniors’ acceptance of HC. So far, *social influence* has been interpreted as consumer’s perception of how much the social environment believes a technology should be used by the consumer and demands her or him to do so (Venkatesh et al. 2012). However, our interviews show that besides the active call of others to use a technology, the personal assessment and recommendation of the social environment play a crucial role for seniors’ intention to use an HC as can be seen from this statement:

“It’s a basis of trust when you say it’s great. [...] And then it [i.e., Ada] will be accepted.” #13

We thus redefine the factor in accordance with Laumer et al. (2019) “as the extent to which consumers perceive that important others believe they should use a particular technology, that important others say they should use a particular technology, that important others recommend to use a technology or that individuals observe that others are using a specific technology” (Laumer et al. 2019, p. 7).

There are several other studies which show that social influence plays a key role for seniors’ technology acceptance (e.g., Guan et al. 2017). Based on the new definition of the construct, we conclude our third proposition (**P4**):

The higher seniors’ social influence, the higher is seniors’ intention to use HC.

7.4.2.2 Habit

In the interviews, it turned out that none of the seniors has knowingly used a chatbot so far. Therefore, in contrast to Laumer et al. (2019), we found no support for the proposition that the *habit* to use chatbots in general, has an influence on the acceptance and use of HC in particular. Instead, it has become apparent that many seniors are used to visit the physician in person (e.g., because of chronic illnesses) and unlike the younger generation do not have the habit to search for their symptoms online in order to diagnose themselves. Although the usage of HC is quite new for younger persons, too, they have grown up using current technologies such as search engines. Therefore, the first step towards using a machine to assess illnesses is less unfamiliar for younger people than it is for older ones:

“The challenge is even if you’re technology affine or have a good knowledge of digital media, but you’ve been visiting a physician for 60, 70 years [...] it’s a challenge to not deal with a human.” #12

Our results confirm the findings of researchers that seniors are struggling to adapt to the rapid changes entailed by current technologies and to alter their habits accordingly (Holgersson and Söderström 2019). We thus redefined habit as the habit of seniors to visit physicians instead of searching online or consulting HC, resulting in the following proposition (**P5**):

a: The higher seniors’ habit of visiting physicians, the lower seniors’ intention to use HC.

b: The higher seniors’ habit of visiting physicians, the lower seniors’ actual use of HC.

7.4.3 Newly Added Factors

In addition to the original factors, the interviewees also addressed factors that are specific to the seniors’ acceptance and use of HC. These factors will be explained below.

7.4.3.1 Patience

The interviews revealed *patience* of the seniors as a factor that influences seniors’ usage behavior directly. As already described, the handling of HC is new for seniors and demands that they have patience while learning to use the technology. Accordingly, dealing with HC requires more patience from seniors than from younger generations, who are more used to deal with related technologies such as chatbots in general. However, seniors often do not have the patience to learn how to use a new app (Carlsson and Walden 2015). The inherent characteristic of HC to guide their users through a dialogue with several questions and query loops about symptoms and health conditions amplifies this problem further:

“When you forgot to enter an answer during the usage of the application, maybe you can’t enter it anymore, then you get confused and impatient.” #11

In other words, the patience of seniors with HC plays a crucial role in the final usage or rejection of HC. Accordingly, we propose (**P6**):

The higher seniors’ patience, the higher seniors’ actual use of HC.

7.4.3.2 Resistance to Change

Since senior citizens mostly have limited experience with technologies, even small adjustments of known app design and user interfaces can cause major defensive reactions. In this context,

the inherent characteristic of HC to guide the user through a longer dialogue has led to difficulties for the seniors. For example, the navigation of the HC Ada was different from other known apps, which makes handling less intuitive for them. In this regard, in order to return to the previous page, participants have to click on the greyed answer to the earlier question instead of using a familiar back button. This *resistance to change* can be explained by resorting to a theory known from psychology. According to the theory of fluid and crystallized intelligence, humans own two kinds of intelligence: i.e., fluid and crystallized intelligence. Crystallized intelligence refers to abilities gained from experience, whereas fluid intelligence describes whether people are able to adapt and solve new situations as they occur. With age, crystallized intelligence increases while fluid intelligence decreases (Cattell 1963). However, this makes it more difficult for seniors to adapt to unknown procedures. Consequently, seniors are more resistant to changes (Hoque and Sorwar 2017). As a result, we formulated the following proposition (**P7**):

The higher seniors’ resistance to change, the lower seniors’ intention to use HC.

7.4.3.3 Need for Emotional Support

During the interviews it was most frequently mentioned that the diagnosis by a machine as provided by the HC is not sufficient for seniors. The participants emphasized the importance of a human who knows you, has cared for you for years, who you are familiar with, who you trust, with whom you can have a dialogue and who you can ask questions if necessary. This is pointed out in the following statements:

“But at the physician’s, I could also ask more in-between questions. And the whole process [in Ada] is planned, or rather preset, and the app doesn’t allow interposed questions.” #23

“But I think it’s the personal relationship. This feeling, you are taken seriously, that you are valued, that people listen to you. The other person is trying, to get to the root of your pain or complaints and find out more. [...] And this patient-physician relationship, that is also unique, I think. So in my case [...] my physician, who made the first diagnosis breast cancer, he conveyed that we will get it all done [...] And that’s what I mean when I am speaking about personal relationship. You don’t have to die of cancer anymore. We can manage that. And because of this charisma of the physician, I’m sure that was a giant healing factor for me.” #8

The *need for emotional support* can be of great relevance for seniors in particular, since they often have less social contact and more severe illnesses than the younger generation (Wilson 2018). In particular, the dialogue with the HC cannot replace the need for human empathy. Therefore, we conclude the following proposition (**P8**):

The higher seniors’ need for emotional support, the lower is seniors’ intention to use HC.

7.4.3.4 Technology Anxiety

Few interviews have shown that working with HC is a big and daunting challenge for seniors due to their lack of digital literacy. Seniors did not grow up with technology in their everyday lives. Rather, they were confronted with it at an advanced age. As a result, they have a lower level of technological experience compared to the younger generation. This leads to greater skepticism and fear regarding technology in general, as can be seen in the following:

“But I just think [...] that many people are being cheated and fooled through the internet and modern forms of communication. So we have to be extremely careful with everything we do, right?” #15

Our interviews have shown that among seniors not only the fear of technology itself, but also the fear of making mistakes while using it is severe. Furthermore, seniors state the anxiety of being cheated through a technological application. Hence, we call *technology anxiety* not only the fear of using technologies, but also the fear of being deceived while using it. HC, which are based on complex AI and NLP algorithms, can be judged even less by senior citizens than conventional technologies. Therefore, technology anxiety significantly affects the intention of seniors to use HC. In this context our ninth proposition is defined as follows (**P9**):

The higher seniors’ technology anxiety, the lower is seniors’ intention to use HC.

7.4.3.5 Privacy Risk Expectancy

The interviews have shown that *privacy risks* and data protection regarding the entered symptoms are a factor which seniors take into account. For example, the following consideration is expressed during the interviews:

“Who can guarantee me that I can work with this system safely? On the one hand, I know it cannot work, if I do not give qualified information. On the other hand, I would be worried depending on the type of the disease. I do not necessarily want everyone to know it.” #19

However, opinions on privacy were controversial among the participants. In this regard, a lot of participants did not consider the protection of data as necessary and preferred a well-founded diagnosis:

“My protected data will no longer be useful to me if I am no longer alive.” #13

“Anybody can hack my data if they want, I don’t care. For me, a well-founded diagnosis is more important. That’s what I would use this app for, because I need information, a well-founded diagnosis.” #3

Compared to the younger generation, privacy risks could not be identified as a decisive factor (Laumer et al. 2019). We thus conclude that the importance of a well-founded diagnosis becomes higher with age and therefore exceeds the fear of privacy risks. Nevertheless, seniors tend to consider privacy issues as relevant. Based on these insights, we have concluded the following propositions (**P10**):

a: The higher seniors’ privacy risks expectancy, the lower is seniors’ intention to use HC.

b: The higher the need for an informed diagnosis, the lower is the negative impact of privacy risks expectancy on seniors’ intention to use HC.

7.4.3.6 Medication

Many seniors need medication to treat chronic diseases (Charlesworth et al. 2015). Thus, participants expressed the desire that Ada should consider prior *medication* and potential side effects to achieve a well-founded diagnosis. However, the current version of Ada is not able to consider pre-medication and their potential effects on health for diagnosis. Furthermore, a few seniors mentioned that they need to see a physician to get new prescriptions for medication. Since these seniors, who need medical treatment, visit the physician anyway, the perceived performance expectancy of HC decreases:

“I wouldn’t get any medication through this app [Ada] [...]” #3

Due to the current technological limitations, we have formulated the following proposition (**P11**):

The higher the amount of seniors’ requirement for medication, the lower is seniors’ performance expectancy of HC.

7.4.3.7 Medical History

The older people get, the longer their *medical history* is and the more likely they are to already suffer from a disease (Charlesworth et al. 2015). This history of chronic illnesses can have a crucial influence on current symptoms entered to the HC. Therefore, the use of HC only makes sense if they take into account the individual anamnesis of seniors, as the following statement shows:

“But for my case, it doesn’t fit. [...] I had a stroke. Naturally I have strange sensations and things that feel like pain or cold or something and it [the answer of the HC] doesn’t fit right.” #7

Due to their longer lifespan, seniors are more likely to have had contact with the topic of diseases and medical treatment (Charlesworth et al. 2015; World Health Organization 2019). Therefore, more than half of the interviews have shown, that seniors’ performance expectancy is strongly dependent on the complexity of their diseases:

“She [Ada] can identify simple symptoms and diagnose them correctly, but as it becomes complex, I would no longer trust Ada, because there can be many different causes for one symptom.” #8

The interviews showed that the medical issues of seniors cannot be neglected and that they are decisive for the performance expectancy of HC. As a result, we conclude the following proposition (**P12**):

The larger seniors’ medical history, the lower is seniors’ performance expectancy of HC.

7.4.3.8 Trust in System

As in the previous study on HC (Laumer et al. 2019) the interviews showed that *trust in system* has an influence on the extent to which the HC is perceived as privacy friendly, how skeptical the seniors are towards the technology, and how the performance of the HC is expected to be. In order to increase the trust of seniors in HC, it is necessary to provide them the possibility to test the system. Interviewed seniors have not yet gained any experience with HC. Therefore, in many interviews participants stated that they could only build trust by using the system and checking whether it provides correct diagnoses. This is stated in the following example:

“Well, I’d always have it cleared up, I guess. If she [Ada] now tells me, for example, that you have osteoarthritis, [...] then I would go to the physician first and have it checked. And if he also says that I have osteoarthritis, then I would trust the system

more. Next time, if I [...] have diarrhea or something and she gives me a diagnosis, then maybe I would say: Last time, she was right. So maybe it’s okay this time, so it’s right.”
#8

Accordingly, we propose **(P13)**:

a: The higher seniors’ trust in system, the higher is seniors’ performance expectancy of HC.

b: The higher seniors’ trust in system, the lower is seniors’ technology anxiety of HC.

c: The higher seniors’ trust in system, the lower is seniors’ privacy risk expectancy of HC.

7.4.3.9 Trust in Provider

As shown for younger age groups, *trust in provider*, similar to trust in system, is a key factor in the adoption decision of the elderly (Laumer et al. 2019). In the case of seniors’ usage, however, it was mainly important to the participants that the HC is supported and co-developed by physicians and that the data is generated and approved by physicians. This can be seen in the following statement:

“Yeah, well, is it written by physicians? Do physicians stand behind it?” #11

The high relevance of this factor can be explained by the fact that this age group has a very high respect for the profession and the knowledge of doctors (Marcinowicz et al. 2014). Therefore, seniors call for transparent explanations of provider’s product development process of the HC and of the creation of the utilized database. The interviews have also shown that trust in the provider increases seniors’ performance expectancy towards the HC, lowers the perception of privacy risks and technology anxiety, and builds trust in the system, and. Therefore, we propose **(P14)**:

a: The higher seniors’ trust in provider, the higher is seniors’ performance expectancy of HC.

b: The higher seniors’ trust in provider, the lower is seniors’ privacy risk expectancy of HC.

c: The higher seniors’ trust in provider, the lower is seniors’ technology anxiety regarding HC.

d: The higher seniors’ trust in provider, the higher is seniors’ trust in the system.

7.4.3.10 Access to Health System

Access to public healthcare services is limited for many people, especially the elderly (Steele et al. 2009). This is the case as seniors are often not mobile enough to autonomously access public healthcare services. Furthermore, the physicians are often very busy and thus hard to

contact. Technologies like HC can partly counteract this problem. In nearly half of our interviews various scenarios are described, where access to a medical professional was difficult:

“The physicians are overloaded. You can’t reach the physicians, sometimes for days.”
#11

HC could make it easier and faster for seniors to get access to medical advice, increasing their performance expectancy regarding this technology. Thus, we conclude the following proposition (**P15**):

The lower seniors’ perceived access to health systems, the higher is seniors’ performance expectancy of HC.

7.4.3.11 Compatibility

Our interviews have shown that HC hold the potential to serve as a central repository for medical records (as symptoms are entered regularly) and as a shortcut to a wide variety of medical knowledge (as it is based on a growing database). If HC would be integrated in the healthcare system, healthcare processes thus could be optimized, for example by avoiding unnecessary repetitive examinations, and seniors’ performance expectancy would increase. This is illustrated by the following statement:

“When you can simply say: everything [diagnoses and test results] is available and the physician [...] can access all necessary information just by a click, I would always use it immediately.” #21

Accordingly, we have concluded the following propositions (**P16**):

The higher the compatibility of HC with seniors’ healthcare system, the higher is seniors’ performance expectancy of HC.

7.4.3.12 Technology Self-Efficacy

During the interviews the seniors very often felt that they do not have the competences to make use of the HC properly on their own. This had an influence on seniors’ perception of how much effort the technology usage requires. This aspect is illustrated in the following statement:

“I somehow have no understanding for the technology, I have become more cumbersome” #14

In this context, Chen et al. (2014) also have shown that an increased technology self-efficacy has a positive impact on seniors’ perceived ease of use regarding gerontechnology. As a result,

technology self-efficacy has to be increased in order to increase the seniors’ effort expectancy. In this regard, studies have shown that trainings are helpful to introduce seniors to information and communication technologies (Nguyen et al. 2014). Furthermore, the interviews have shown that more technologically competent seniors feel less anxious in dealing with technology. As a result, we formulate the following propositions (**P17**):

- a: The higher seniors’ technology self-efficacy, the higher is seniors’ effort expectancy of HC.*
- b: The higher seniors’ technology self-efficacy, the lower is seniors’ technology anxiety regarding HC.*

7.4.3.13 Self-Reported Health Condition

The interviews showed that the more seniors were restricted in their health, for example in terms of visual or auditory ability, the harder it was for them to use the HC. The restrictive effect of a poor *self-reported health condition* on effort expectancy can also be confirmed by the findings from other senior technology acceptance research (Chen et al. 2014). As a consequence, it is especially important to seniors to have an age-appropriate visualization of the HC. In this regard, it was mentioned several times that the font size as well as the distance between the buttons of Ada were too small and that the visual contrast of the app should be higher. This can easily be explained by the fact that eyesight decreases with age (Glasser and Campbell 1998). Therefore, we proposed (**P18**):

The better seniors’ self-reported health condition, the higher is seniors’ effort expectancy of HC.

7.5 Conclusion, Limitations, and Future Research

In summary, we conducted an exploratory qualitative study by interviewing 23 participants in 21 interviews with a mean age of 71 years. This procedure allowed us to develop senior-specific factors regarding the acceptance and use of HC. As a result, we were able to confirm the applicability of the UTAUT2 model as a theoretical foundation. However, besides the already known factors of UTAUT2, the interviews showed that two factors had to be redefined and thirteen new factors had to be added. Furthermore, it turned out that the factor hedonic motivation did not seem to be relevant in the context of seniors’ acceptance and use of HC. Based on our extended model it is now possible to evaluate the factors affecting the usage intention and subsequent behavior of seniors regarding HC in order to be able to fully exploit the advantages of HC for elderly people.

Although a study of Laumer et al. (2019) has identified factors specific to the acceptance and use of HC by the younger generation before, it cannot be assumed that the results are transferable to the context of seniors as this target group has highly specific requirements for the use of technologies. Furthermore, existing research on technology acceptance of seniors is also not sufficient to explain the intention and actual use of HC by older people, as HC is a novel technology in a highly sensitive context (i.e., healthcare) that was not investigated in detail before. Therefore, we contribute to the IS research of seniors’ technology acceptance theoretically by identifying senior-specific factors that influence the intention and usage behavior of this specific mHealth service. In general, our results show that the participants not only showed a high need for technological assistance in order to use HC, but also expressed a greater demand for emotional support than younger people. Furthermore, seniors might have problems using currently available HC solutions with regard to their medication or medical history, which are not yet considered adequately by HC. In addition, it cannot be assumed that seniors are used to handle novel and complex chatbot technologies and therefore have issues trusting them.

Beside these general findings, the relevance of some factors is particularly noteworthy. First of all, the interviews illustrated the relevance of context-specific factors (e.g., medication, medical history) that should be considered for HC solutions used by elderly people. This finding could also be transferred to technologies other than HC (e.g., intelligent pillbox, wearable medical devices) in order to increase seniors’ performance expectancy of the respective technology (Solís et al. 2017; Zhang and Li 2017). Among all identified context-specific factors, need for emotional support should be particularly emphasized. When it comes to seniors’ health issues, emotions play a major role. In this context, studies have shown that a positive attitude can promote recovery (Glass and Maddox 1992; Shyu et al. 2006; Tsouna-Hadjis et al. 2000). Especially for seniors, who have a higher risk for serious illnesses and a higher probability of being less socially involved than the younger generation, emotional support is crucially important to deal with serious diagnoses. In particular, the support of a physician, they have known and trusted for many years, can be required in the event of a serious illness. As a consequence, the fact that machines do not offer emotional support, but rather perform very rational, short and concise diagnoses, can seriously hinder the acceptance and use of HC depending on the health condition of the senior citizen. Therefore, the questions arising are, for which types of diseases, according symptoms, and sub-groups of seniors the use of HC makes sense and how should the diagnoses of HC be designed in order to give users a feeling of emotional support (e.g., appropriate presenting of diagnosis, digital follow-up discussion with

physician, references to help-offerings)? Given the high relevance of the factor need for emotional support, future research could furthermore investigate whether all identified factors influencing the acceptance and usage of HC should be equally prioritized and how context and specific environmental situations could affect this weighting.

Our study also offers some practical implications. In order to leverage the high potential of HC for the elderly, HC providers should design their offerings according to the specific needs of the elderly. For example, in order to ensure seniors’ acceptance and continuous use of HC, it is essential that their medication and medical history is taken into account when they are guided through the dialogue with the machine. Since older people in particular are more likely to have pre-existing conditions, consideration of medication and medical history is particularly important to ensure a higher performance expectancy exhibited by seniors. In addition, the visual design of HC should be kept as simple as possible, based on known procedures, and appropriate for persons with weaker eyesight. Furthermore, the answer design should be easy to understand, but comprehensive enough in case of inquiries. Finally, as seniors need appropriate facilitating conditions and an enhanced technology self-efficacy to accept and use HC, procedures should be created to explain the advantages and usage of HC step by step. This could be done by introductory descriptions within the system itself or through trainings in cooperation with insurances and physicians.

Due to the explorative nature of this study, only 21 interviews are considered, which could result in a potential generalizability problem. However, we applied various criteria to ensure rigor and trustworthiness to encounter this problem (i.e., theoretical saturation, investigator triangulation). In addition, the study was only conducted within one European country (i.e., Germany). In order to examine possible cultural influences, further studies should be executed in other countries and continents. In addition, the context-specific factors identified within this study could be investigated in more detail. In this regard, researchers could examine in which situations and under what conditions HC are used by seniors (e.g., for which symptoms). The interviews showed that there is a general willingness to use HC among seniors, but it does not replace a visit to a physician for them. Therefore, future studies should consider that HC will be used in parallel with a visit to the doctor and represent an additional offer for senior citizens. Overall, we have established a model that identifies the factors that contribute to senior citizens’ acceptance and use of HC. This model could serve as a basis for further research, e.g., concerning seniors’ acceptance of other related AI technologies (e.g., humanoid robots).

8 Paper 3.A: Coordinating Human and Machine Learning for Effective Organizational Learning

Title

Coordinating Human and Machine Learning for Effective Organizational Learning

Authors

Sturm, Timo^a; Gerlach, Jin^a; Pumplun, Luisa^b; Mesbah, Neda^b; Peters, Felix^b; Tauchert, Christoph^b; Nan, Ning^c; Buxmann, Peter^c

Authors with the same superscript contributed equally to this paper

Publication Outlet

MIS Quarterly

Abstract

With the rise of machine learning (ML), humans are no longer the only ones capable of learning and contributing to an organization's stock of knowledge. We study how organizations can coordinate human learning and ML in order to learn effectively as a whole. Based on a series of agent-based simulations, we find that, first, ML can reduce an organization's demand for human explorative learning that is aimed at uncovering new ideas; second, adjustments to ML systems made by humans are largely beneficial, but this effect can diminish or even become harmful under certain conditions; and third, reliance on knowledge created by ML systems can facilitate organizational learning in turbulent environments, but this requires significant investments in the initial setup of these systems as well as adequately coordinating them with humans. These insights contribute to rethinking organizational learning in the presence of ML and can aid organizations in reallocating scarce resources to facilitate organizational learning in practice.

Keywords

Artificial Intelligence, Machine Learning, Human-Machine Coordination, Organizational Learning, Simulation, Agent-Based Modeling

License

License Agreement with Copyright Clearance Center (Publisher: Society for Management Information Systems and Management Information Systems Research Center of the University of Minnesota) on the use of the abstract and Table 16.

Table 16: Summary of Results Regarding Organizational Learning Effectiveness

Research Questions	Findings and Propositions	Implications
RQ1: The Role of Human Exploration in the Presence of ML Systems	ML systems with a high initial learning capability reduce the need for human exploration (see P1).	<ul style="list-style-type: none"> ML systems' ability to take over explorative tasks counters learning myopia, allowing humans to learn at their preferred pace. Organizations should consider the reallocation of R&D resources to the initial setup of ML systems.
RQ2: Reconfiguration of ML Systems by Humans	<p>Humans' learning behavior moderates the non-linear effect of reconfiguration intensity on organizational learning effectiveness. For ML systems with a</p> <ul style="list-style-type: none"> low initial learning capability: If humans engage in exploitation (exploration), this effect is positive and decreases (increases) in strength with increasing reconfiguration intensity (see P2a). high initial learning capability: If humans engage in exploitation, this effect decreases in strength with increasing reconfiguration intensity. If humans engage in exploration, the reconfiguration intensity has an inverted U-shaped effect (see P2b). 	<ul style="list-style-type: none"> Acquiring high levels of organizational knowledge requires at least a moderate amount of reconfiguration effort. Humans should never be completely taken "out of the loop," even if tasks are largely automated. As the deep problem understanding of domain experts is required for reconfiguration efforts, leaving reconfiguration of ML systems to the IT department alone is not sufficient.
RQ3: Coordinating Human Learning and ML Systems in Turbulent Environments	In turbulent environments, effective organizational learning with ML systems requires human exploration and a rapid codification of knowledgeable humans' beliefs. The more turbulent the environment, the more beneficial the rapid codification of beliefs offered by ML systems with a high initial learning capability will be (see P3).	<ul style="list-style-type: none"> Reliance on knowledge created by ML systems can be beneficial for organizations in turbulent environments, reducing the need for more radical measures (e.g., forced personnel turnovers). Significant investments in the initial setup of ML systems and appropriate coordination of humans and ML systems are required to materialize these beneficial effects.

9 Discussion of Contributions and Concluding Remarks

To realize the full potential of ML systems in organizations, certain obstacles must be overcome and driving forces must be effectively harnessed to ensure the successful adoption of these systems. Especially with ML systems, there is a high risk of misadoption and thus of unintended negative consequences due to their unique characteristics (Rana et al. 2021).

This thesis indicates how the risk of an abortive adoption can be minimized, while simultaneously responding to the calls of Benbya et al. (2021), Kane et al. (2021), and Shaw et al. (2019) for research on effective integration of ML systems into organizations in general and in the field of healthcare in particular. It advances our understanding of what factors influence the organizational adoption of ML systems and reveals that a purely technical perspective is not sufficient to understand the complex socio-technical adoption process. Rather, this dissertation takes a holistic view, and illustrates how several factors beyond those of a strictly technical nature can be addressed. It further demonstrates the positive impact this integrative approach can have on organizational key performance indicators, especially on the organizational knowledge level.

The contributions of this thesis to theory and practice are discussed in detail below. Limitations and propositions for future research resulting from these findings are included in the respective publications.

9.1 Overarching Theoretical Contributions

In the following, the theoretical contributions of this dissertation are outlined. These are organized and presented according to their respective contribution to answering the three RQs (see Table 1).

In response to RQ1, this thesis contributes to providing an overview of potential factors influencing the successful adoption of ML systems in organizations. For this purpose, the basic TOE and NASSS frameworks served as a basis and were supplemented in the first two papers (1.A and 1.B) with empirically derived factors specific to the integration of ML systems in (healthcare) organizations (DePietro et al. 1990; Greenhalgh et al. 2017). In this way, an inclusive overview could be obtained that applies in general, as well as in the field of healthcare. Newly discovered factors include not only technical, but also organizational and human-centered aspects, as well as influences from an organization's environment on the organizational adoption of ML systems, such as:

- the unique characteristics of ML systems, including their lack of transparency and ability to adapt
- the availability of data that is of high quality, anonymized, representative, and adequately formatted
- the existence of an innovative organizational culture and an ML strategy
- the readiness of an organization's employees and customers/patients
- legal, structural, or ethical requirements, for example due to the GDPR, a works council, or medical ethics

These factors show how closely technology and humans are intertwined in the context of ML systems and that a holistic perspective that extends beyond existing theoretical frameworks is necessary to ensure successful adoption (DePietro et al. 1990; Greenhalgh et al. 2017). Furthermore, building on the identified factors, a maturity model was derived in paper 1.B, which provides researchers with a tool to operationalize empirical studies on ML systems adoption. This instrument can be applied, for instance, to measure an (healthcare) organization's status quo in the adoption of ML systems, and to include a corresponding variable in structural equation modeling.

In addition, paper 2.A, 2.B, and 2.C contribute to the fulfillment of the derived factors, thus answering RQ2. In particular, paper 2.A shows how the necessary database for training ML systems in organizations can be generated by referring to the concept of data donation. On a theoretical level, the paper draws on the privacy calculus and extends the theoretical concept to include a critical behavioral bias, namely an empathy gap (Van Boven and Loewenstein 2005; Culnan and Armstrong 1999; Dinev and Hart 2006; Krasnova et al. 2010; Loewenstein 2005; Metcalfe and Mischel 1999). The study, which examines actual data donation behavior using a self-developed catalog of 14 questions, demonstrates that individuals in an agitated mental state are less prone to base their data donation decision on established trust in the receiving organization, but are significantly influenced by perceived privacy risks when deciding whether to disclose data. These findings show the strong influence of situational factors, such as the COVID-19 pandemic and associated mental states of individuals on their donation behavior, and thus help to explain the previously contradictory results on the effect of trust on data-sharing behavior (e.g., Bansal et al. 2010; Dinev and Hart 2006; Kehr et al. 2013; Metzger 2004; Norberg et al. 2007). They also underline the high relevance of perceived privacy risks, the influence of which is constant across different situational factors and individuals' emotional states, and should therefore always be accounted for by organizations in order to encourage data donation. Paper 2.B further shows that the TPB can be used and adapted to study the impact

of a more transparent design of ML systems on customers' attitudes and willingness to pay (Ajzen 1991). In this case, the positive impact of increased transparency on customers' attitudes is fully mediated by trust in the ML system. This result shows the high relevance of trust in the context of transparent ML systems. The importance of trust in ML systems is reaffirmed in paper 2.C as well. It identifies the factors by which the UTAUT2 model needs to be extended or adjusted to predict seniors' use of a specific ML system in the healthcare context (Venkatesh et al. 2012). Among them are trust in the system and its provider organization. Situational factors that vary from person to person depending on the context are perceived to be particularly important for seniors' usage intentions. These include factors, such as previous medication, which the ML system should take into account, or the need for emotional support, especially in the case of severe diagnoses, which is still lacking in provided ML systems. It thus extends the existing literature in the field of gerontology, which has so far only dealt with the influence of situational factors on seniors' technology use in a rudimentary and rather generalistic way (e.g., Chen et al. 2014; Renaud and Biljon 2008; Vichitvanichphong et al. 2017). However, ML systems in particular hold the potential to closely adapt to the needs of seniors and adequately respond to emerging situational factors (Goldenberg et al. 2021).

When ML systems are successfully adopted by an organization, they can have a major impact on various performance metrics and especially on organizational learning. Paper 3.A describes the derivation of a model of organizational learning with ML systems based on March's (1991) model. The established model allows for new insights into how to alleviate the long-standing problem of learning myopia (Levinthal and March 1993). While resource-intensive countermeasures such as R&D, employee training, and incentives had to be taken in the past to foster new ideas (e.g., Bushee 1998; Kretkowski 1998; Levinthal and March 1993), the introduction of ML systems reveals a new effective way to balance exploitative and exploratory learning in organizations. The results contribute further to the discussion revolving around automating or augmenting processes with ML systems (e.g., Brynjolfsson and Mitchell 2017; Rai et al. 2019; Raisch and Krakowski 2021). They demonstrate the relevance of humans in the loop for organizational learning, reconfiguring ML systems on a regular basis. Such reconfiguration is only possible if a certain level of domain expertise is available and being developed further over time. Indeed, it is unlikely that pure technicians such as data scientists can fully assume this role. As a result, humans and especially domain experts will continue to have a vital part to play in the organizational learning process, even as ML systems work side-by-side with them.

From a more abstract perspective, this dissertation contributes to a deeper understanding of the adoption of ML systems in organizations, especially in healthcare. It illustrates which specific contextual conditions the healthcare sector imposes for the integration of ML systems and which situational contingencies contribute to the systems' successful adoption. We draw on and substantially extend several existing generic frameworks from IS and healthcare research to deepen our understanding of the key predictors of ML system adoption in organizations. Furthermore, we elaborate on these findings by showing at a more in-depth level how these predictors of ML system adoption can be achieved and what impact these systems might have on the organization. Thereby, we take a socio-technical perspective and illustrate how closely humans and machines are intertwined in harnessing the full potential of ML systems. In this way, we are contributing to advancing research on the adoption of ML systems and fostering the maturing of this line of research (Venkatesh et al. 2007).

9.2 Overarching Practical Contributions

In addition to the theoretical contributions, this thesis offers insights into practical aspects arising from the RQs (see Table 1). This allows decision-makers in (healthcare) organizations to gain practical insights into what is required to adopt ML systems, what the concrete levers for integration are, and what impact their efforts might have.

Addressing RQ1, papers 1.A and 2.B provide a guide for organizations that intend to or are in the process of adopting ML systems. For this purpose, various factors influencing adoption and represent potential pitfalls for the integration of ML systems in organizations in general and healthcare organizations in particular are presented. Not only do they allow (healthcare) organizations to identify challenges early on that might stand in the way of successful ML system adoption, but they also highlight drivers that could be enforced in organizations to actively facilitate the adoption of ML systems. In addition, the systematic examination of the provided factors and the application of the maturity model allow the evaluation of the status quo of ML system adoption in organizations. This allows an organization to identify areas for improvement and to compare its adoption status with competitors in a structured way. Widespread use of the maturity model would also permit entire industries to be compared with one another, such as the healthcare and the financial sectors, to derive insights and arguments for policy support measures.

Papers 2.A, 2.B, and 2.C examine the above factors at a deeper level to show how they can be met and thus answer RQ2. Paper 2.A examines the factors that influence potential donors'

willingness to disclose their data. In this context, data donation represents a valid concept for providing a suitable database, especially for healthcare organizations to train their ML systems, as there is a high societal willingness to donate health data. Such a data donation could be realized via a data donation platform. Care should be taken to ensure that the data donation platform is not only targeted at crisis situations and the prevailing need for data at that time, but promotes a continuous donation process. Not only could this encourage donation behavior, but it could also create a larger, more meaningful database that allows for a combination of non-crisis and crisis data. An important, persistent factor influencing donation behavior is the privacy-friendliness of the organizations involved in the data donation process. It therefore makes sense as an organization to pursue a privacy-friendly image, for example by transparently communicating pursued privacy measures. In addition, it may be useful to help donors to overcome their agitated state in order for them to regain trust, consider this in their decision-making, and eventually increase their willingness to donate. To counteract ML systems' inscrutability, paper 2.B provides an example of how a more transparent ML system could be designed. Such a design is shown to have a positive effect on customer trust, which in turn creates a positive attitude among customers. Beyond that, the trust resulting from the transparent design of ML systems may have other desired effects on business metrics not explored in this paper, such as customer loyalty or technology acceptance (e.g., Leninkumar 2017; Suh and Han 2002). Overall, it has been demonstrated that a strong need for transparent ML systems exists among customers and can actually be monetized separately. Consequently, a transparent design is a lever that can be harnessed for several positive outcomes and should be targeted by organizations offering ML systems. Paper 2.C takes this a step further by examining a variety of factors beyond transparency that influence customers' intentions to use ML systems in a healthcare context. It thus provides a guideline for designing ML systems and their usage process to maximize acceptance by potential customers or patients. In particular, context specific factors peculiar to the individuals and their given situation must be considered in the creation of ML systems; for instance, when potential customers might wish for emotional support if the ML system were to provide a particularly severe diagnosis. This might be addressed, for example, by displaying a selection of mental health support services or an immediate telephone referral to the primary care physician. Overall, the individual person and her or his specific needs ought to be the focus of the ML system's development. Especially for an application in the healthcare sector, the functional scope of the ML system should be carefully determined.

With the advent of ML systems in organizations, a number of practical benefits can be achieved, as enquired by RQ3 and answered in paper 3.A. As these systems participate in organizational learning, employees are no longer the only ones adding knowledge to the organization. This frees up employees to learn more according to their own preferences. They are now no longer forced to constantly acquire new knowledge to drive, for example, innovation, but can also spend more of their time applying and refining what they have already learned. This could not only help employees develop in-depth knowledge in a specific domain, but also increase employee well-being. Furthermore, the organization no longer needs to invest to the same extent in costly activities such as employee incentives, free time for ideation, or R&D to foster new knowledge (e.g., Bushee 1998; Kretkowski 1998; Levinthal and March 1993), as ML systems already take over some part of this. The slack resources thereby released can be spent on other measures and investments, such as the reconfiguration of ML systems. Moderate reconfiguration of these systems can, in turn, have a positive impact on the level of knowledge in the organization and should therefore be actively managed. This required reconfiguration is one of the reasons why humans will continue to play an essential role alongside machines. Since employees will need to update ML systems at some point in time, for example by providing new data, they should not be excluded from the organizational learning process, but instead learn alongside, with, and from the machines. Since ML systems often take into account very different criteria than humans, it is indeed useful to examine the machines more closely. For instance, physicians who gain new insights into why a lung nodule is deemed malignant by interacting with an ML system can improve their ability to make diagnoses. This requires a certain level of transparency of the ML system in order to understand its fundamentals, functioning, and output, and to compare it with one's own mental models (e.g., Letham et al. 2015; Xu et al. 2015). Furthermore, since ML systems possess the ability to swiftly discover patterns in large data sets reflecting an organization's environment (Ransbotham et al. 2020), they play a role in the organization's ability to deal with volatile contexts. Today's economy is characterized by constant, rapid changes, for example triggered by crises such as the COVID-19 pandemic (Ancona et al. 2020; Benbya et al. 2020; Nan and Tanriverdi 2017). For organizations to survive in such unstable environments, they should quickly absorb the beliefs of both employees and ML systems. Yet ML systems, despite their ability to introduce new, independent beliefs into an organization, are not a panacea to counteracting environmental turbulence. Properly managed, however, they can be a support to organizational learning, obviating the need for other more radical measures, such as forced turnovers of employees.

More abstractly, this work highlights how a business problem that will become increasingly important in the future, namely the adoption of ML systems, can be concretely addressed in organizations. It provides practical guidance on how decision makers, designers, and instructors can influence the integration of ML systems into organizations by creating the technical, organizational, and personnel prerequisites, and actively managing potential outcomes for the business (Venkatesh et al. 2007).

9.3 Concluding Remarks

ML systems offer the potential to solve today's pressing societal problems by promising to increase effectiveness and efficiency in virtually all organizational forms and industries, and especially in the healthcare sector (Brynjolfsson and McAfee 2017a; Lynch 2017; Shaw et al. 2019). In particular, the successful adoption of ML systems in (healthcare) organizations determines whether these potentials can be leveraged or whether negative side effects occur, as they can in the case of a failed introduction (Rana et al. 2021). This work draws on three qualitative studies, two experimental survey studies, and a simulation to investigate, disentangle, and clarify complex adoption of ML systems and to derive applicable guidance for organizations thriving to integrate ML systems into their processes. It reveals not only the relevance of the context under study, in this case the healthcare sector, but also the necessity of including technical, organizational, and human aspects in conjunction, to enable fruitful adoption of ML systems. The result is a more nuanced framework for the factors influencing organizational adoption of ML systems in general and in the context of healthcare, concrete insights into the path to successful integration of these systems into the organization, and an outlook on potential organizational impact.

These findings pave the way for further prospective research. Some ideas for subsequent research projects have already been named in the respective research papers. Beyond the potential follow-up projects described therein, other research avenues are eligible for IS scholars. First, the studies presented in this thesis take a holistic, abstract perspective on the adoption process of ML systems. Follow-up studies, however, could focus particularly on the final phase of the adoption process, during which the technology is ultimately to be implemented and embraced for routine use (Damanpour and Schneider 2006). In this phase, a particular focus lies on the collaboration between human users and ML systems, raising further questions for future research, such as: which decisions should be made by human users and which by ML systems? What aspects of human-machine interaction affect decision effectiveness and efficiency, and how can these outcomes be measured? What training is

needed to educate human users on the handling of ML systems? What are the individual differences of users in applying ML systems?

Furthermore, in this thesis, some factors influencing the adoption of ML systems were considered in more detail, to identify feasible paths to their fulfillment. Future research could not only address the remaining prerequisites that were identified but not explored in depth in this dissertation (e.g., developing an ML strategy, finding a suitable business case, gaining top management support, creating a shared ethical foundation), but also investigate alternative valid avenues that might contribute to harnessing the drivers and overcoming the barriers to the adoption of ML systems. For instance, other notions apart from data donation, such as publicly available data sets, platforms for data vending, or data received in exchange for services could be employed to obtain a high quality database for training ML systems. The circumstances and ways in which these concepts can be beneficially applied remain the subject of further research. In addition, other potential demands on and design alternatives for transparent ML systems could be systematically investigated, for example, depending on the type of training data or the criticality of the context in which the systems are to be applied. Likewise, IS researchers might explore further creative approaches to increase acceptance of ML systems among potential users. A promising stream of research in this regard concerns the anthropomorphic design of ML systems (e.g., Pfeuffer et al. 2019; Seeger et al. 2021).

Moreover, other effects of ML systems on the organization could be investigated in the future. Even though organizational learning is a key factor in organizational survival (e.g., Bushee 1998; Cohen and Levinthal 1989), the introduction of ML systems entails other organizational consequences beyond knowledge acquisition. As an example, scholars may investigate the influence of these systems on organizational structures and hierarchies in general, and on team dynamics, employee job satisfaction, or decision-making effectiveness in particular (e.g., Benbya et al. 2021). Adopting ML systems in organizations will also yield consequences that exceed the scope of a singular organization. Future research could therefore investigate the linkages between the growing organizational application of ML systems and associated automation and the potential consequences for professions and the labor market (e.g., Strich et al. 2021; Willcocks 2020).

It is my hope that this dissertation will sharpen the understanding of what challenges and success factors are essential in the integration of ML systems and provide fertile ground for conducting further research in the area of organizational adoption of ML systems.

References

- Acquisti, A. 2004. "Privacy in Electronic Commerce and the Economics of Immediate Gratification," in *Proceedings of the 5th ACM Conference on Electronic Commerce*, New York, New York, United States, pp. 21–29. (<https://doi.org/10.1145/988772.988777>).
- Acquisti, A., and Grossklags, J. 2005. "Privacy and Rationality in Individual Decision Making," *IEEE Security and Privacy* (3:1), pp. 26–33. (<https://doi.org/10.1109/MSP.2005.22>).
- Acquisti, A., John, L. K., and Loewenstein, G. 2013. "What Is Privacy Worth?," *Journal of Legal Studies* (42:2), pp. 249–274. (<https://doi.org/10.1086/671754>).
- Ada Health. 2019. "Your Personal Health Guide." (<https://ada.com>, accessed May 28, 2019).
- Adadi, A., and Berrada, M. 2018. "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access* (6), pp. 52138–52160. (<https://doi.org/10.1109/ACCESS.2018.2870052>).
- Ågerfalk, P. J., Conboy, K., and Myers, M. D. 2020. "Information Systems in the Age of Pandemics: COVID-19 and Beyond," *European Journal of Information Systems* (29:3), pp. 203–207. (<https://doi.org/10.1080/0960085X.2020.1771968>).
- Agrawal, A., Gans, J., and Goldfarb, A. 2018. "Is Your Company's Data Actually Valuable in the AI Era?," *Harvard Business Review*, Harvard Business Review. (<https://hbr.org/2018/01/is-your-companys-data-actually-valuable-in-the-ai-era>).
- Ajzen, I. 1991. "The Theory of Planned Behavior," *Organizational Behavior and Human Decision Processes* (50:2), pp. 179–211. ([https://doi.org/10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T)).
- Akcay, S., Kundegorski, M. E., Willcocks, C. G., and Breckon, T. P. 2018. "Using Deep Convolutional Neural Network Architectures for Object Classification and Detection within X-Ray Baggage Security Imagery," *IEEE Transactions on Information Forensics and Security* (13:9), pp. 2203–2215. (<https://doi.org/10.1109/TIFS.2018.2812196>).
- Akkiraju, R., Sinha, V., Xu, A., Mahmud, J., Gundecha, P., Liu, Z., Liu, X., and Schumacher, J. 2020. "Characterizing Machine Learning Processes: A Maturity Framework," in *Proceedings of 18th International Conference on Business Process Management*, Seville, Spain, pp. 17–31. (https://doi.org/10.1007/978-3-030-58666-9_2).

- Alami, H., Lehoux, P., Denis, J. L., Motulsky, A., Petitgand, C., Savoldelli, M., Rouquet, R., Gagnon, M. P., Roy, D., and Fortin, J. P. 2021. "Organizational Readiness for Artificial Intelligence in Health Care: Insights for Decision-Making and Practice," *Journal of Health Organization and Management* (35:1), pp. 106–114. (<https://doi.org/10.1108/JHOM-03-2020-0074>).
- Alashoor, T., Al-Maidani, N., and Al-Jabri, I. 2018. "The Privacy Calculus under Positive and Negative Mood States," in *Proceedings of 39th International Conference on Information Systems*, San Francisco, California, United States.
- Alsheibani, S. A., Cheung, Y., and Messom, C. 2019a. "Factors Inhibiting the Adoption of Artificial Intelligence at Organizational-Level: A Preliminary Investigation," in *Proceedings of 25th Americas Conference of Information Systems*, Cancun, Mexico.
- Alsheibani, S. A., Cheung, Y., and Messom, C. 2019b. "Towards an Artificial Intelligence Maturity Model: From Science Fiction to Business Facts," in *Proceedings of the 23rd Pacific Asia Conference on Information Systems*, Xi'an, China.
- Alsheibani, S., Cheung, Y., and Messom, C. 2018. "Artificial Intelligence Adoption: AI-Readiness at Firm-Level," in *Proceedings of the 22nd Pacific Asia Conference on Information Systems*, Yokohama, Japan.
- Amato, F., Marrone, S., Moscato, V., Piantadosi, G., Picariello, A., and Sansone, C. 2017. "Chatbots Meet Ehealth: Automatizing Healthcare," in *Workshop on Artificial Intelligence with Application in Health*, Bari, Italy.
- Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., and Zimmermann, T. 2019. "Software Engineering for Machine Learning: A Case Study," in *Proceedings of the 41st International Conference on Software Engineering*, Montreal, Quebec, Canada, pp. 291–300. (<https://doi.org/10.1109/icse-seip.2019.00042>).
- Ananny, M., and Crawford, K. 2018. "Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability," *New Media & Society* (20:3), pp. 973–989. (<https://doi.org/10.1177/2F1461444816676645>).
- Ancona, D., Williams, M., and Gerlach, G. 2020. "The Overlooked Key to Leading Through Chaos," *MIT Sloan Management Review* (62:1), pp. 34–39.
- Anderson, C. L., and Agarwal, R. 2011. "The Digitization of Healthcare: Boundary Risks, Emotion, and Consumer Willingness to Disclose Personal Health Information,"

- Information Systems Research* (22:3), pp. 469–490.
(<https://doi.org/10.1287/isre.1100.0335>).
- Andrews, W., Sau, M., Dekate, C., Mullen, A., Brant, K. F., Revang, M., and Plummer, D. C. 2018. “Predicts 2018: Artificial Intelligence,” *Gartner*, Stamford, p. 21.
(<https://www.gartner.com/en/documents/3827163>, accessed November 21, 2019).
- Ardabili, S. F., Mosavi, A., Ghamisi, P., Ferdinand, F., Varkonyi-Koczy, A. R., Reuter, U., Rabczuk, T., and Atkinson, P. M. 2020. “COVID-19 Outbreak Prediction with Machine Learning,” *Algorithms* (13:10). (<https://doi.org/10.3390/a13100249>).
- Argote, L., Lee, S., and Park, J. 2021. “Organizational Learning Processes and Outcomes: Major Findings and Future Research Directions,” *Management Science* (67:9), pp. 5399–5429. (<https://doi.org/10.1287/mnsc.2020.3693>).
- Arora, A. 2020. “Conceptualising Artificial Intelligence as a Digital Healthcare Innovation: An Introductory Review,” *Medical Devices: Evidence and Research* (13), pp. 223–230. (<https://doi.org/10.2147/MDER.S262590>).
- Asatiani, A., Malo, P., Nagbøl, P. R., Penttinen, E., Rinta-Kahila, T., and 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), pp. 325–352. (<https://doi.org/10.17705/1jais.00664>).
- Awa, H. O., and Ojiabo, O. U. 2016. “A Model of Adoption Determinants of ERP within T-O-E Framework,” *Information Technology and People* (29:4), pp. 901–930.
(<https://doi.org/10.1108/ITP-03-2015-0068>).
- Babylon. 2019. “Mission.” (<https://www.babylonhealth.com/about>, accessed May 28, 2019).
- Bagozzi, R. P., and Yi, Y. 2012. “Specification, Evaluation, and Interpretation of Structural Equation Models,” *Journal of the Academy of Marketing Science* (40:1), pp. 8–34.
(<https://doi.org/10.1007/s11747-011-0278-x>).
- Bagozzi, R. P., Yi, Y., and Phillips, L. W. 1991. “Assessing Construct Validity in Organizational Research,” *Administrative Science Quarterly* (36:3), pp. 421–458.
(<https://doi.org/10.2307/2393203>).
- Bamberger, P. A. 2018. “AMD—Clarifying What We Are about and Where We Are Going,” *Academy of Management Discoveries* (4:1), pp. 1–10.
(<https://doi.org/10.5465/amd.2018.0003>).

- Bansal, G., Zahedi, F. M., and Gefen, D. 2010. "The Impact of Personal Dispositions on Information Sensitivity, Privacy Concern and Trust in Disclosing Health Information Online," *Decision Support Systems* (49:2), pp. 138–150. (<https://doi.org/10.1016/j.dss.2010.01.010>).
- Bansal, G., Zahedi, F. M., and Gefen, D. 2016. "Do Context and Personality Matter? Trust and Privacy Concerns in Disclosing Private Information Online," *Information and Management* (53:1), pp. 1–21. (<https://doi.org/10.1016/j.im.2015.08.001>).
- Bardhan, I., Chen, H., and Karahanna, E. 2020. "Connecting Systems, Data, and People: A Multidisciplinary Research Roadmap for Chronic Disease Management," *MIS Quarterly* (44:1), pp. 185–201. (<https://doi.org/10.25300/MISQ/2020/14644>).
- Bavel, J. J. V., Baicker, K., Boggio, P. S., Capraro, V., Cichocka, A., Cikara, M., Crockett, M. J., Crum, A. J., Douglas, K. M., Druckman, J. N., Drury, J., Dube, O., Ellemers, N., Finkel, E. J., Fowler, J. H., Gelfand, M., Han, S., Haslam, S. A., Jetten, J., Kitayama, S., Mobbs, D., Napper, L. E., Packer, D. J., Pennycook, G., Peters, E., Petty, R. E., Rand, D. G., Reicher, S. D., Schnall, S., Shariff, A., Skitka, L. J., Smith, S. S., Sunstein, C. R., Tabri, N., Tucker, J. A., Linden, S. van der, Lange, P. van, Weeden, K. A., Wohl, M. J. A., Zaki, J., Zion, S. R., and Willer, R. 2020. "Using Social and Behavioural Science to Support COVID-19 Pandemic Response," *Nature Human Behaviour* (4:5), pp. 460–471. (<https://doi.org/10.1038/s41562-020-0884-z>).
- Becker, J., Knackstedt, R., and Pöppelbuß, J. 2009. "Developing Maturity Models for IT Management," *Business & Information Systems Engineering* (1:3), pp. 213–222. (<https://doi.org/10.1007/s12599-009-0044-5>).
- Benbasat, I., Goldstein, D. K., and Mead, M. 1987. "The Case Research Strategy in Studies of Information Systems," *MIS Quarterly* (11:3), pp. 369–386. (<https://doi.org/10.2307/248684>).
- Benbya, H., Nan, N., Tanriverdi, H., and Yoo, Y. 2020. "Complexity and Information Systems Research in the Emerging Digital World," *MIS Quarterly* (44:1), pp. 1–17. (<https://doi.org/10.25300/misq/2020/13304>).
- Benbya, H., Pachidi, S., and 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), pp. 281–303. (<https://doi.org/10.17705/1jais.00662>).

- Berente, N., Gu, B., Recker, J., and Santhanam, R. 2021. "Managing Artificial Intelligence," *Mis Quarterly* (45:3), pp. 1433–1450. (<https://doi.org/10.25300/MISQ/2021/16274>).
- Bibault, J.-E., Chaix, B., Nectoux, P., Pienkowski, A., Guillemasé, A., and Brouard, B. 2019. "Healthcare Ex Machina: Are Conversational Agents Ready for Prime Time in Oncology?," *Clinical and Translational Radiation Oncology* (16), pp. 55–59. (<https://doi.org/10.1016/j.ctro.2019.04.002>).
- Bietz, M., Patrick, K., and Bloss, C. 2019. "Data Donation as a Model for Citizen Science Health Research," *Citizen Science: Theory and Practice* (4:1), pp. 1–11. (<https://doi.org/10.5334/cstp.178>).
- Bogner, A., Littig, B., and Menz, W. 2009. "Introduction: Expert Interviews – an Introduction to a New Methodological Debate," in *Interviewing Experts* (1st ed.), A. Bogner, B. Littig, and W. Menz (eds.), London, United Kingdom: Palgrave Macmillan, pp. 1–13.
- Van Boven, L., and Loewenstein, G. 2005. "Empathy Gaps in Emotional Perspective Taking," in *Other Minds: How Humans Bridge the Divide between Self and Others* (1st ed.), B. F. Malle and S. D. Hodges (eds.), New York, New York, United States: Guilford Press, pp. 284–297.
- Brakemeier, H., Wagner, A., and Buxmann, P. 2017. "When Risk Perceptions Are Nothing but Guesses – An Evaluability Perspective on Privacy Risks," in *Proceedings of the 38th International Conference on Information Systems*, Seoul, South Korea.
- Branch. 2019. "The Power of Financial Access." (<https://branch.co/about>, accessed October 4, 2019).
- Brant, K. F., Sau, M., Mullen, A., Revang, M., Dekate, C., Plummer, D. C., and Andrews, W. 2017. "Predicts 2018: Artificial Intelligence," Gartner. (<https://www.gartner.com/en/documents/3827163/predicts-2018-artificial-intelligence>).
- Breidert, C., Hahsler, M., and Reutterer, T. 2006. "A Review of Methods for Measuring Willingness-to-Pay," *Innovative Marketing* (2:4), pp. 8–32.
- Bremser, C. 2018. "Starting Points for Big Data Adoption," in *Proceedings of 26th European Conference on Information Systems*, Portsmouth, United Kingdom.
- Bremser, C., Piller, G., and Rothlauf, F. 2017. "Strategies and Influencing Factors for Big Data Exploration," in *Proceedings of the 23rd Americas Conference on Information Systems*, Boston, Massachusetts, United States.

- Briganti, G., and Le Moine, O. 2020. "Artificial Intelligence in Medicine: Today and Tomorrow," *Frontiers in Medicine* (7:February), pp. 1–6. (<https://doi.org/10.3389/fmed.2020.00027>).
- Brinati, D., Campagner, A., Ferrari, D., Locatelli, M., Banfi, G., and Cabitza, F. 2020. "Detection of COVID-19 Infection from Routine Blood Exams with Machine Learning: A Feasibility Study," *Journal of Medical Systems* (44:8), Journal of Medical Systems. (<https://doi.org/10.1007/s10916-020-01597-4>).
- vom Brocke, J., Maaß, W. M., Buxmann, P., Maedche, A., Leimeister, J. M., and Pecht, G. 2018. "Future Work and Enterprise Systems," *Business & Information Systems Engineering* (60:4), pp. 357–366. (<https://doi.org/10.1007/s12599-018-0544-2>).
- Brooks, S. K., Webster, R. K., Smith, L. E., Woodland, L., Wessely, S., Greenberg, N., and Rubin, G. J. 2020. "The Psychological Impact of Quarantine and How to Reduce It: Rapid Review of the Evidence," *The Lancet* (395:10227), pp. 912–920. ([https://doi.org/10.1016/S0140-6736\(20\)30460-8](https://doi.org/10.1016/S0140-6736(20)30460-8)).
- de Bruin, T., Rosemann, M., Freeze, R., and Kulkarni, U. 2005. "Understanding the Main Phases of Developing a Maturity Assessment Model," in *Proceedings of the 16th Australasian Conference on Information Systems*, Sydney, Australia, pp. 8–19.
- Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitzoff, T., Filar, B., Anderson, H., Roff, H., Allen, G., Steinhardt, J., Flynn, C., O hEigearthaigh, S., Beard, S., Belfield, H., Farquhar, S., Lyle, C., Crootof, R., Evans, O., Page, M., Bryson, J., Yampolskiy, R., and Amodei, D. 2018. "The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation," Future of Humanity Institute; University of Oxford; Centre for the Study of Existential Risk; University of Cambridge; Center for a New American Security; Electronic Frontier Foundation; OpenAI.
- Bryman, A. 2016. *Social Research Methods*, (5th ed.), Oxford: Oxford University Press.
- Brynjolfsson, E., and McAfee, A. 2017a. "The Business of Artificial Intelligence," *Harvard Business Review*. (<https://hbr.org/2017/07/the-business-of-artificial-intelligence>).
- Brynjolfsson, E., and McAfee, A. 2017b. "What's Driving the Machine Learning Explosion?," *Harvard Business Review*2. (<https://hbr.org/2017/07/whats-driving-the-machine-learning-explosion?ab=seriesnav-bigidea>).

- Brynjolfsson, E., and Mitchell, T. 2017. "What Can Machine Learning Do? Workforce Implications," *Science* (358:6370), pp. 1530–1534. (<https://doi.org/10.1126/science.aap8062>).
- Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlström, P., Henke, N., and Trench, M. 2017. "Artificial Intelligence: The Next Digital Frontier?," McKinsey Global Institute.
- Burrell, J. 2015. "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* (3:1), pp. 1–12. (<https://doi.org/10.2139/ssrn.2660674>).
- Bushee, B. J. 1998. "The Influence of Institutional Investors on Myopic R&D Investment Behavior," *The Accounting Review* (73:3), pp. 305–333.
- Calero Valdez, A., and Ziefle, M. 2019. "The Users' Perspective on the Privacy-Utility Trade-Offs in Health Recommender Systems," *International Journal of Human-Computer Studies* (121), pp. 108–121. (<https://doi.org/10.1016/j.ijhcs.2018.04.003>).
- Carlsson, C., and Walden, P. 2015. "Digital Wellness for Young Elderly: Research Methodology and Technology Adaptation," in *Proceedings of the 28th Bled EConference*, Bled, Slovenia, pp. 239–250.
- Carter, N., Bryant-Lukosius, D., Dicenso, A., Blythe, J., and Neville, A. J. 2014. "The Use of Triangulation in Qualitative Research," *Oncology Nursing Forum* (41:5), pp. 545–547. (<https://doi.org/10.1177/14687941149830077>).
- Carter, S., Armstrong, Z., Schubert, L., Johnson, I., and Olah, C. 2019. "Activation Atlas," *Distill* (4:3). (<https://doi.org/10.23915/distill.00015>).
- Carvalho, J. V., Rocha, Á., and Abreu, A. 2016. "Maturity Models of Healthcare Information Systems and Technologies: A Literature Review," *Journal of Medical Systems* (40:6), *Journal of Medical Systems*, pp. 1–10. (<https://doi.org/10.1007/s10916-016-0486-5>).
- Cattell, R. B. 1963. "Theory of Fluid and Crystallized Intelligence: A Critical Experiment," *Journal of Educational Psychology* (54:1), pp. 1–22. (<https://doi.org/10.1037/h0046743>).
- Chai, Y., and Li, W. 2019. "Towards Deep Learning Interpretability : A Topic Modeling Approach," in *Proceedings of 40th International Conference on Information Systems*, Munich, Germany.
- Chakraborti, T., Kulkarni, A., Sreedharan, S., Smith, D. E., and Kambhampati, S. 2019. "Explicability? Legibility? Predictability? Transparency? Privacy? Security? The

- Emerging Landscape of Interpretable Agent Behavior,” in *Proceedings of the 29th International Conference on Automated Planning and Scheduling*, Berkley, California, United States, pp. 86–96.
- Charlesworth, C. J., Smit, E., Lee, D. S. H., Alramadhan, F., and Odden, M. C. 2015. “Polypharmacy Among Adults Aged 65 Years and Older in the United States: 1988-2010,” *Journals of Gerontology: Medical Sciences* (70:8), pp. 989–995. (<https://doi.org/10.1093/gerona/glv013>).
- Chen, C., Lin, K., Rudin, C., Shaposhnik, Y., Wang, S., and Wang, T. 2018. “An Interpretable Model with Globally Consistent Explanations for Credit Risk,” in *Proceedings of the NIPS 2018 Workshop on Challenges and Opportunities for AI in Financial Services: The Impact of Fairness, Explainability, Accuracy, and Privacy*, Montréal, Quebec, Canada, pp. 1–10.
- Chen, K., Hoi, A., and Chan, S. 2014. “Gerontechnology Acceptance by Elderly Hong Kong Chinese: A Senior Technology Acceptance Model (STAM),” *Ergonomics* (57:5), pp. 635–652. (<https://doi.org/10.1080/00140139.2014.895855>).
- Chernev, A. 2003. “Reverse Pricing and Online Price Elicitation Strategies in Consumer Choice,” *Journal of Consumer Psychology* (13:1–2), pp. 51–62. (https://doi.org/10.1207/s15327663jcp13-1&2_05).
- Chiu, C.-Y., Chen, S., and Chen, C.-L. 2017. “An Integrated Perspective of TOE Framework and Innovation Diffusion in Broadband Mobile Applications Adoption by Enterprises,” *International Journal of Management, Economics and Social Sciences* (6:1), pp. 14–39.
- Chiu, C. M., Hsu, M. H., Lai, H., and Chang, C. M. 2012. “Re-Examining the Influence of Trust on Online Repeat Purchase Intention: The Moderating Role of Habit and Its Antecedents,” *Decision Support Systems* (53:4), pp. 835–845. (<https://doi.org/10.1016/j.dss.2012.05.021>).
- Cho, A. 2020. “AI Systems Aim to Sniff out Coronavirus Outbreaks,” *Science* (368:6493), pp. 810–811. (<https://doi.org/10.1126/science.368.6493.810>).
- Chong, A. Y. L., and Chan, F. T. S. 2012. “Structural Equation Modeling for Multi-Stage Analysis on Radio Frequency Identification (RFID) Diffusion in the Health Care Industry,” *Expert Systems with Applications* (39:10), Elsevier Ltd, pp. 8645–8654. (<https://doi.org/10.1016/j.eswa.2012.01.201>).

- Choudhury, P., Allen, R. T., and Endres, M. G. 2021. "Machine Learning for Pattern Discovery in Management Research," *Strategic Management Journal* (42:1), pp. 30–57. (<https://doi.org/10.1002/smj.3215>).
- Cohen, E. L., and Hoffner, C. 2013. "Gifts of Giving: The Role of Empathy and Perceived Benefits to Others and Self in Young Adults' Decisions to Become Organ Donors," *Journal of Health Psychology* (18:1), pp. 128–138. (<https://doi.org/10.1177/1359105311433910>).
- Cohen, W. M., and Levinthal, D. A. 1989. "Innovation and Learning: The Two Faces of R & D," *The Economic Journal* (99:397), pp. 569–596. (<https://doi.org/10.2307/2233763>).
- Coombs, C., Hislop, D., Taneva, S. K., and Barnard, S. 2020. "The Strategic Impacts of Intelligent Automation for Knowledge and Service Work: An Interdisciplinary Review," *Journal of Strategic Information Systems* (29:4), Elsevier, p. 101600. (<https://doi.org/10.1016/j.jsis.2020.101600>).
- Cresswell, K., and Sheikh, A. 2013. "Organizational Issues in the Implementation and Adoption of Health Information Technology Innovations: An Interpretative Review," *International Journal of Medical Informatics* (82:5), Elsevier Ireland Ltd, pp. e73–e86. (<https://doi.org/10.1016/j.ijmedinf.2012.10.007>).
- Creswell, J. W. 2007. *Qualitative Inquiry and Research Design: Choosing among Five Traditions*, (2nd ed.), Thousand Oaks, California, United States: SAGE Publications.
- Crowston, K., and Bolici, F. 2019. "Impacts of Machine Learning on Work," in *Proceedings of the 52nd Hawaii International Conference on System Sciences*, Grand Wailea, Maui, Hawaii, United States, pp. 5961–5970. (<https://doi.org/10.24251/HICSS.2019.719>).
- Cruz-Jesus, F., Pinheiro, A., and Oliveira, T. 2019. "Understanding CRM Adoption Stages: Empirical Analysis Building on the TOE Framework," *Computers in Industry* (109), pp. 1–13. (<https://doi.org/10.1016/j.compind.2019.03.007>).
- Culnan, M. J., and Armstrong, P. K. 1999. "Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation," *Organization Science* (10:1), pp. 104–115. (<https://doi.org/10.1287/orsc.10.1.104>).
- Damanpour, F., and Schneider, M. 2006. "Phases of the Adoption of Innovation in Organizations: Effects of Environment, Organization and Top Managers," *British Journal of Management* (17:3), pp. 215–236. (<https://doi.org/10.1111/j.1467->

8551.2006.00498.x).

- Damanpour, F., and Schneider, M. 2009. "Characteristics of Innovation and Innovation Adoption in Public Organizations: Assessing the Role of Managers," *Journal of Public Administration Research and Theory* (19:3), pp. 495–522. (<https://doi.org/10.1093/jopart/mun021>).
- Davis, F. D., Bagozzi, R. P., and Warshaw, P. R. 1989. "User Acceptance of Computer Technology: A Comparison of Two Theoretical Models," *Management Science* (35:8), pp. 982–1003. (<https://doi.org/10.1287/mnsc.35.8.982>).
- Davis, M. H. 1980. "A Multidimensional Approach to Individual Differences in Empathy," *Journal of Personality and Social Psychology* (10:85). (<https://doi.org/10.1037/0022-3514.44.1.113>).
- Davison, A. C., and Hinkley, D. V. 1997. "Bootstrap Methods and Their Application," *Technometrics* (42:2), pp. 216–217. (<https://doi.org/10.2307/1271471>).
- Davison, R. M., and Martinsons, M. G. 2015. "Context Is King! Considering Particularism in Research Design and Reporting," *Journal of Information Technology* (31:3), pp. 241–249. (<https://doi.org/10.1057/jit.2015.19>).
- DePietro, R., Wiarda, E., and Fleischer, M. 1990. "The Context for Change: Organization, Technology and Environment," in *The Process of Technological Innovation* (4th ed.), L. Tornatzky and M. Fleischer (eds.), Lexington, US: Lexington Books, pp. 152–175.
- Deutsch, M. 1958. "Trust and Suspicion," *Journal of Conflict Resolution* (2:4), pp. 265–279. (<https://doi.org/10.1177/002200275800200401>).
- Dhaliwal, J. S., and Benbasat, I. 1996. "The Use and Effects of Knowledge-Based System Explanations: Theoretical Foundations and a Framework for Empirical Evaluation," *Information Systems Research* (7:3), pp. 342–363. (<https://doi.org/10.1287/isre.7.3.342>).
- Diakopoulos, N. 2016. "Accountability in Algorithmic Decision Making," *Communications of the ACM* (59:2), pp. 56–62. (<https://doi.org/10.1145/2844110>).
- Dickert, S., and Slovic, P. 2009. "Attentional Mechanisms in the Generation of Sympathy," *Judgment and Decision Making* (4:4), pp. 297–306.
- Dinev, T., and Hart, P. 2006. "An Extended Privacy Calculus Model for E-Commerce Transactions," *Information Systems Research* (17:1), pp. 61–80. (<https://doi.org/10.1287/isre.1060.0080>).

- Dinev, T., McConnell, A. R., Smith, H. J., Dinev, T., Raton, B., and Smith, H. J. 2015. "Informing Privacy Research through Information Systems, Psychology, and Behavioral Economics: Thinking Outside the 'APCO' Box," *Information Systems Research* (26:4), pp. 639–655. (<https://doi.org/10.1287/isre.2015.0600>).
- Dinev, T., Xu, H., Smith, J. H., and Hart, P. 2013. "Information Privacy and Correlates: An Empirical Attempt to Bridge and Distinguish Privacyrelated Concepts," *European Journal of Information Systems* (22:3), pp. 295–316. (<https://doi.org/10.1057/ejis.2012.23>).
- Dodgson, M. 1993. "Organizational Learning: A Review of Some Literatures," *Organization Studies* (14:3), pp. 375–394. (<https://doi.org/10.1177/017084069301400303>).
- Domingos, P. 2012. "A Few Useful Things to Know about Machine Learning," *Communications of the ACM* (55:10), pp. 78–87. (<https://doi.org/10.1145/2347736.2347755>).
- Doshi-Velez, F., and Kim, B. 2017. "Towards a Rigorous Science of Interpretable Machine Learning," *ArXiv*.
- Drew, D. A., Nguyen, L. H., Steves, C. J., Menni, C., Freydin, M., Varsavsky, T., Sudre, C. H., Jorge Cardoso, M., Ourselin, S., Wolf, J., Spector, T. D., and Chan, A. T. 2020. "Rapid Implementation of Mobile Technology for Real-Time Epidemiology of COVID-19," *Science* (368:6497), pp. 1362–1367. (<https://doi.org/10.1126/science.abc0473>).
- Elliot, B., and Andrews, W. 2017. "A Framework for Applying AI in the Enterprise," Stamford. (<https://www.gartner.com/ngw/globalassets/en/information-technology/documents/insights/a-framework-for-applying-ai-in-the-enterprise.pdf>).
- Elliott, R., and Timulak, L. 2015. "Descriptive and Interpretive Approaches to Qualitative Research," in *A Handbook of Research Methods for Clinical and Health Psychology*, J. Miles and P. Gilbert (eds.), Oxford University Press, pp. 147–159. (<https://doi.org/10.1093/med:psych/9780198527565.003.0011>).
- European Commission. 2003. "Commission Recommendation of 6 May 2003 Concerning the Definition of Micro, Small and Medium-Sized Enterprises (Notified under Document Number C (2003) 1422)," *Official Journal of the European Union* (L 124), pp. 36–41.
- European Medicines Agency. 2020. "EMA Regulatory Science to 2025: Strategic Reflection." (<https://www.ema.europa.eu/en/documents/regulatory-procedural->

- guideline/ema-regulatory-science-2025-strategic-reflection_en.pdf).
- European Parliament. 2016. "Regulation (EU) 2016/679 of the European Parliament and of the Council," *Official Journal of the European Union*. (<https://eur-lex.europa.eu/eli/reg/2016/679/oj>).
- Eurostat. 2018. "Internet Access and Use Statistics - Households and Individuals." (https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Archive:Internet_access_and_use_statistics_-_households_and_individuals&oldid=379591, accessed September 29, 2020).
- Fadhil, A., and Schiavo, G. 2019. "Designing for Health Chatbots," *ArXiv*.
- Fagherazzi, G., Goetzinger, C., Rashid, M. A., Aguayo, G. A., and Huiart, L. 2020. "Digital Health Strategies to Fight COVID-19 Worldwide," *Journal of Medical Internet Research* (22:6), JMIR Publications Inc. (<https://doi.org/10.2196/19284>).
- Fernandez, C., and Provost, F. 2019. "Counterfactual Explanations for Data-Driven Decisions," in *Proceedings of 40th International Conference on Information Systems*, Munich, Germany.
- FICO. 2019a. "FICO® Score X Data." (<https://www.fico.com/en/products/fico-score-x-data>, accessed November 21, 2019).
- FICO. 2019b. "Explainable Machine Learning Challenge." (<https://community.fico.com/s/explainable-machine-learning-challenge>, accessed October 23, 2019).
- Flam, F. 2018. "IBM's Watson Hasn't Beaten Cancer, But A.I. Still Has Promise." (<https://www.bloomberg.com/opinion/articles/2018-08-24/ibm-s-watson-failed-against-cancer-but-a-i-still-has-promise>, accessed November 20, 2018).
- Flick, U. 2004. "Triangulation in Qualitative Research," in *A Companion to Qualitative Research*, U. Flick, E. von Kardoff, and I. Steinke (eds.), London, United Kingdom: SAGE Publications, pp. 178–183.
- Flick, U., von Kardoff, E., and Steinke, I. 2004. "What Is Qualitative Research? An Introduction to the Field," in *A Companion to Qualitative Research*, U. Flick, E. Kardoff, and I. Steinke (eds.), London, United Kingdom: SAGE Publications, pp. 3–11.
- Food and Drug Administration. 2021. "Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan."

- (<https://www.fda.gov/media/145022/download>).
- Fornell, C., and Larcker, D. F. 1981. "Evaluating Structural Equation Models with Unobservable Variables and Measurement Error," *Journal of Marketing Research* (18:1), pp. 39–50. (<https://doi.org/10.2307/3151312>).
- Fox, G., and James, T. L. 2020. "Toward an Understanding of the Antecedents to Health Information Privacy Concern: A Mixed Methods Study," *Information Systems Frontiers*, *Information Systems Frontiers*. (<https://doi.org/10.1007/s10796-020-10053-0>).
- Frambach, R. T., and Schillewaert, N. 2002. "Organizational Innovation Adoption: A Multi-Level Framework of Determinants and Opportunities for Future Research," *Journal of Business Research* (55:2), pp. 163–176. ([https://doi.org/10.1016/S0148-2963\(00\)00152-1](https://doi.org/10.1016/S0148-2963(00)00152-1)).
- France, C. R., and France, J. L. 2018. "Fear of Donation-Related Stimuli Is Reported across Different Levels of Donation Experience," *Transfusion* (58:1), pp. 113–120. (<https://doi.org/10.1111/trf.14382>).
- Fraser, P., Moultrie, J., and Gregory, M. 2002. "The Use of Maturity Models I Grids as a Tool in Assessing Product Development Capability," in *IEEE International Engineering Management Conference*, Cambridge, United Kingdom, pp. 244–249.
- Fuller, C. M., Simmering, M. J., Atinc, G., Atinc, Y., and Babin, B. J. 2016. "Common Methods Variance Detection in Business Research," *Journal of Business Research* (69:8), pp. 3192–3198. (<https://doi.org/10.1016/j.jbusres.2015.12.008>).
- Furneaux, B., and Wade, M. R. 2011. "An Exploration of Organizational Level Information Systems Discontinuance Intentions," *MIS Quarterly* (35:3), pp. 573–598. (<https://doi.org/10.2307/23042797>).
- Gao, J., and Koronios, A. 2010. "Mobile Application Development for Senior Citizens," in *Proceedings of the 14th Pacific Asia Conference on Information Systems*, Taipei, Taiwan, pp. 214–223.
- Gao, S. P. L. M. R. A. J. 2008. "Non-Normality of Data in Structural Equation Models," *Journal of the Transportation Research Board* (2082:1), pp. 116–124. (<https://doi.org/doi.org/10.3141/2082-14>).
- Gedikli, F., Jannach, D., and Ge, M. 2014. "How Should I Explain? A Comparison of Different Explanation Types for Recommender Systems," *Journal of Human-Computer*

- Studies* (72:4), pp. 367–382. (<https://doi.org/10.1016/j.ijhcs.2013.12.007>).
- George, J. F. 2004. “The Theory of Planned Behavior and Internet Purchasing,” *Internet Research* (14:3), pp. 198–212. (<https://doi.org/10.1108/10662240410542634>).
- Gerlach, J. P., Buxmann, P., Dinev, T., and Lowry, P. 2019. “‘They’re All the Same!’ Stereotypical Thinking and Systematic Errors in Users’ Privacy-Related Judgments about Online Services,” *Journal of the Association for Information Systems* (20:6), pp. 787–823. (<https://doi.org/10.17705/1jais.00551>).
- Giermindl, L. M., Strich, F., Christ, O., Leicht-Deobald, U., and Redzepi, A. 2021. “The Dark Sides of People Analytics: Reviewing the Perils for Organisations and Employees,” *European Journal of Information Systems*, Taylor & Francis, pp. 1–26. (<https://doi.org/10.1080/0960085X.2021.1927213>).
- Glass, T. A., and Maddox, G. L. 1992. “The Quality and Quantity of Social Support: Stroke Recovery as Psycho-Social Transition,” *Social Science and Medicine* (34:11), pp. 1249–1261. ([https://doi.org/10.1016/0277-9536\(92\)90317-J](https://doi.org/10.1016/0277-9536(92)90317-J)).
- Glasser, A., and Campbell, M. C. W. 1998. “Presbyopia and the Optical Changes in the Human Crystalline Lens With Age,” *Vision Research* (38:2), pp. 209–229. ([https://doi.org/10.1016/S0042-6989\(97\)00102-8](https://doi.org/10.1016/S0042-6989(97)00102-8)).
- Gnewuch, U., Morana, S., and Maedche, A. 2017. “Towards Designing Cooperative and Social Conversational Agents for Customer Service,” in *Proceedings of 38th International Conference on Information Systems*, Seoul, South Korea.
- Golden, B. 2006. “Change: Transforming Healthcare Organizations,” *Healthcare Quarterly* (10), pp. 10–19. (<https://doi.org/10.12927/hcq.18490>).
- Goldenberg, D., Kofman, K., Albert, J., Mizrachi, S., Horowitz, A., and Teinemaa, I. 2021. “Personalization in Practice: Methods and Applications,” in *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, Virtual Event, Israel, pp. 1123–1126. (<https://doi.org/10.1145/3437963.3441657>).
- Google. 2019. “Perspectives on Issues in AI Governance.” (<https://ai.google/static/documents/perspectives-on-issues-in-ai-governance.pdf>, accessed November 25, 2019).
- Gounaris, S. P. 2005. “Trust and Commitment Influences on Customer Retention: Insights from Business-to-Business Services,” *Journal of Business Research* (58:2), pp. 126–140.

- ([https://doi.org/10.1016/S0148-2963\(03\)00122-X](https://doi.org/10.1016/S0148-2963(03)00122-X)).
- Grand Review Research. 2018. "MHealth Market Size Worth \$151.57 Billion By 2025," *Grand View Research*. (<https://www.grandviewresearch.com/press-release/global-mhealth-market>, accessed May 29, 2019).
- Greener, S. 2008. *Business Research Methods*, London, United Kingdom: Ventus Publishing ApS.
- Greenhalgh, T., Shaw, S., Wherton, J., Vijayaraghavan, S., Morris, J., Bhattacharya, S., Hanson, P., Campbell-Richards, D., Ramoutar, S., Collard, A., and Hodkinson, I. 2018. "Real-World Implementation of Video Outpatient Consultations at Macro, Meso, and Micro Levels: Mixed-Method Study," *Journal of Medical Internet Research* (20:4). (<https://doi.org/10.2196/jmir.9897>).
- Greenhalgh, T., Wherton, J., Papoutsis, C., Lynch, J., Hughes, G., A'Court, C., Hinder, S., Fahy, N., Procter, R., and Shaw, S. 2017. "Beyond Adoption: A New Framework for Theorizing and Evaluating Nonadoption, Abandonment, and Challenges to the Scale-up, Spread, and Sustainability of Health and Care Technologies," *Journal of Medical Internet Research* (19:11). (<https://doi.org/10.2196/jmir.8775>).
- Gregor, S. 2001. "Explanations from Knowledge-Based Systems and Cooperative Problem Solving: An Empirical Study," *International Journal of Human-Computer Studies* (54:1), pp. 81–105. (<https://doi.org/10.1006/ijhc.2000.0432>).
- Gregor, S., and Benbasat, I. 1999. "Explanations from Intelligent Systems: Theoretical Foundations and Implications for Practice," *MIS Quarterly* (23:4), pp. 497–530. (<https://doi.org/10.2307/249487>).
- Grossklags, J., and Acquisti, A. 2007. "When 25 Cents Is Too Much: An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information," in *Proceedings of 6th Workshop on the Economics of Information Security*, Pittsburgh, Pennsylvania, United States.
- Guan, C., Wei, J., Chan, C. M. L., and Chia, A. B. H. 2017. "Senior Citizens' Self-Efficacy for ICT Use: The Influence of Gender, Social Influence and Social Support," in *Proceedings of the 17th International Conference on Electronic Business*, Dubai, United Arab Emirates, pp. 234–240.
- Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., and Pedreschi, D. 2019. "A

- Survey of Methods for Explaining Black Box Models,” *ACM Computing Surveys* (51:5), p. 93. (<https://doi.org/10.1145/3236009>).
- Hair, J. F., Black, W. C., Babin, B. J., and Anderson, R. E. 2009. *Multivariate Data Analysis*, (7th ed.), Upper Saddle River, New Jersey, United States: Prentice Hall.
- Han, H., and Hyun, S. S. 2015. “Customer Retention in the Medical Tourism Industry: Impact of Quality, Satisfaction, Trust, and Price Reasonableness,” *Tourism Management* (46), pp. 20–29. (<https://doi.org/10.1016/j.tourman.2014.06.003>).
- Hann, I.-H., and Terwiesch, C. 2003. “Measuring the Frictional Costs of Online Transactions: The Case of a Name-Your-Own-Price Channel,” *Management Science* (49:11), pp. 1563–1579. (<https://doi.org/10.1287/mnsc.49.11.1563.20586>).
- Hanoch, Y. 2002. “‘Neither an Angel nor an Ant’: Emotion as an Aid to Bounded Rationality,” *Journal of Economic Psychology* (23:1), pp. 1–25. ([https://doi.org/10.1016/S0167-4870\(01\)00065-4](https://doi.org/10.1016/S0167-4870(01)00065-4)).
- Hansen, T., Jensen, J. M., and Solgaard, H. S. 2004. “Predicting Online Grocery Buying Intention: A Comparison of the Theory of Reasoned Action and the Theory of Planned Behavior,” *International Journal of Information Management* (24:6), pp. 539–550. (<https://doi.org/10.1016/j.ijinfomgt.2004.08.004>).
- Hatta, N. N. M., Miskon, S., and Abdullah, N. S. 2017. “Business Intelligence System Adoption Model for SMEs,” in *Proceedings of the 21st Pacific Asia Conference on Information Systems*, Langkawi, Malaysia.
- Hayes, A. F. 2009. “Beyond Baron and Kenny: Statistical Mediation Analysis in the New Millennium,” *Communication Monographs* (76:4), pp. 408–420. (<https://doi.org/10.1080/03637750903310360>).
- He, J., Baxter, S. L., Xu, Jie, Xu, Jiming, Zhou, X., and Zhang, K. 2019. “The Practical Implementation of Artificial Intelligence Technologies in Medicine,” *Nature Medicine* (25:1), pp. 30–36. (<https://doi.org/10.1038/s41591-018-0307-0>).
- He, K., Zhang, X., Ren, S., and Sun, J. 2015. “Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification,” in *Proceedings of the IEEE International Conference on Computer Vision*, Santiago, Chile, pp. 1026–1034. (<https://doi.org/10.1109/ICCV.2015.123>).
- He, M., Li, Z., Liu, C., Shi, D., and Tan, Z. 2020. “Deployment of Artificial Intelligence in

- Real-World Practice: Opportunity and Challenge,” *Asia-Pacific Journal of Ophthalmology* (9:4), pp. 299–307. (<https://doi.org/10.1097/APO.0000000000000301>).
- Hermanns, H. 2004. “Interviewing as an Activity,” in *A Companion to Qualitative Research*, U. Flick, E. von Kardoff, and I. Steinke (eds.), London, United Kingdom: SAGE Publications, pp. 209–213.
- Hevner, A. R., March, S. T., Park, J., and Ram, S. 2004. “Design Science in Information Systems Research,” *MIS Quarterly* (28:1), pp. 75–105. (<https://doi.org/10.2307/25148625>).
- Hikmet, N., Bhattacharjee, A., Menachemi, N., Kayhan, V. O., and Brooks, R. G. 2008. “The Role of Organizational Factors in the Adoption of Healthcare Information Technology in Florida Hospitals,” *Health Care Management Science* (11:1), pp. 1–9. (<https://doi.org/10.1007/s10729-007-9036-5>).
- Hilaris. 2021. “Healthcare Organizations and Systems,” *Journal of Advanced Practices in Nursing*. (<https://www.hilarispublisher.com/scholarly/healthcare-organizations-and-systems-journals-articles-ppts-list-876.html>, accessed October 5, 2021).
- Hoermann, S., McCabe, K. L., Milne, D. N., and Calvo, R. A. 2017. “Application of Synchronous Text-Based Dialogue Systems in Mental Health Interventions: Systematic Review,” *Journal of Medical Internet Research* (19:8), pp. 1–10. (<https://doi.org/10.2196/jmir.7023>).
- Hofmann, P., Rust, P., and Urbach, N. 2019. “Machine Learning Approaches along the Radiology Value Chain – Rethinking Value Propositions,” in *Proceedings of the 27th European Conference on Information Systems*, Stockholm & Uppsala, Sweden.
- Holgersson, J., and Söderström, E. 2019. “Bridging the Gap - Exploring Elderly Citizens’ Perceptions of Digital Exclusion,” in *Proceedings of the 27th European Conference on Information Systems*, Stockholm & Uppsala, Sweden.
- Hoque, R., and Sorwar, G. 2017. “Understanding Factors Influencing the Adoption of MHealth by the Elderly: An Extension of the UTAUT Model,” *International Journal of Medical Informatics* (101), pp. 75–84. (<https://doi.org/10.1016/j.ijmedinf.2017.02.002>).
- Horan, T. A., and Schooley, B. 2005. “Inter-Organizational Emergency Medical Services: Case Study of Rural Wireless Deployment and Management,” *Information Systems Frontiers* (7:2), pp. 155–173. (<https://doi.org/10.1007/s10796-005-1476-1>).

- Hosanagar, K., and Saxena, A. 2017. "The First Wave of Corporate AI Is Doomed to Fail," *Harvard Business Review*, Harvard Business Review. (<https://hbr.org/2017/04/the-first-wave-of-corporate-ai-is-doomed-to-fail>).
- Hsieh, H.-F., and Shannon, S. E. 2005. "Three Approaches to Qualitative Content Analysis," *Qualitative Health Research* (15:9), pp. 1277–1288. (<https://doi.org/10.1177/1049732305276687>).
- Hu, L. T., and Bentler, P. M. 1999. "Cutoff Criteria for Fit Indexes in Covariance Structure Analysis: Conventional Criteria versus New Alternatives," *Structural Equation Modeling* (6:1), pp. 1–55. (<https://doi.org/10.1080/10705519909540118>).
- Huang, M. H., and Rust, R. T. 2018. "Artificial Intelligence in Service," *Journal of Service Research* (21:2), pp. 155–172. (<https://doi.org/10.1177/1094670517752459>).
- Hufnagl, C., Doctor, E., Behrens, L., Buck, C., and Eymann, T. 2019. "Digitisation along the Patient Pathway in Hospitals," in *Proceedings of the 27th European Conference on Information Systems*, Stockholm & Uppsala, Sweden.
- Hui, K., Kong, H., Teo, H. H., and Lee, S. T. 2007. "The Value of Privacy Assurance: An Exploratory Field Experiment," *MIS Quarterly* (31:1), pp. 19–33. (<https://doi.org/10.2307/25148779>).
- Humphrey, W. S. 1988. "Characterizing the Software Process: A Maturity Framework," *IEEE Software* (5:2), pp. 73–79. (<https://doi.org/10.1109/52.2014>).
- ITU. 2018. "United Nations Activities on Artificial Intelligence (AI)." (<http://handle.itu.int/11.1002/pub/8120d5d5-en>).
- James, H. M., Papoutsis, C., Wherton, J., Greenhalgh, T., and Shaw, S. E. 2021. "Spread, Scale-up, and Sustainability of Video Consulting in Health Care: Systematic Review and Synthesis Guided by the NASSS Framework," *Journal of Medical Internet Research* (23:1). (<https://doi.org/10.2196/23775>).
- Jammal, A. A., Thompson, A. C., Mariottoni, E. B., Berchuck, S. I., Urata, C. N., Estrela, T., Wakil, S. M., Costa, V. P., and Medeiros, F. A. 2020. "Human versus Machine: Comparing a Deep Learning Algorithm to Human Gratings for Detecting Glaucoma on Fundus Photographs," *American Journal of Ophthalmology* (211), pp. 123–131. (<https://doi.org/10.1016/j.ajo.2019.11.006.Human>).
- Jin, C., Chen, W., Cao, Y., Xu, Z., Tan, Z., Zhang, X., Deng, L., Zheng, C., Zhou, J., Shi, H.,

- and Feng, J. 2020. "Development and Evaluation of an Artificial Intelligence System for COVID-19 Diagnosis," *Nature Communications* (11:1). (<https://doi.org/10.1038/s41467-020-18685-1>).
- Joachim, N., Beimborn, D., and Weitzel, T. 2011. "An Instrument for Measuring SOA Maturity," in *Proceedings of the 32nd International Conference on Information Systems*, Shanghai, China, pp. 405–415.
- Jöhnk, J., Weißert, M., and Wyrski, K. 2020. "Ready or Not, AI Comes— An Interview Study of Organizational AI Readiness Factors," *Business and Information Systems Engineering* (63), pp. 5–20. (<https://doi.org/10.1007/s12599-020-00676-7>).
- Jordan, M. I., and Mitchell, T. M. 2015. "Machine Learning: Trends, Perspectives, and Prospects," *Science* (349:6245), pp. 255–260. (<https://doi.org/10.1126/science.aaa8415>).
- Kaggle. 2020. "State of Machine Learning and Data Science 2020." (<https://www.kaggle.com/kaggle-survey-2020>).
- Kane, G. C., Young, A. G., Majchrzak, A., and Ransbotham, S. 2021. "Avoiding an Oppressive Future of Machine Learning: A Design Theory for Emancipatory Assistants," *MIS Quarterly: Management Information Systems* (45:1), pp. 371–396. (<https://doi.org/10.25300/MISQ/2021/1578>).
- Kankanhalli, A., Hahn, J., Tan, S., and Gao, G. 2016. "Big Data and Analytics in Healthcare: Introduction to the Special Section," *Information Systems Frontiers* (18:2), pp. 233–235. (<https://doi.org/10.1007/s10796-016-9641-2>).
- Karahanna, E., Benbasat, I., Bapna, R., and Rai, A. 2018. "Opportunities and Challenges for Different Types of Online Experiments," *MIS Quarterly* (42:4), iii–xi.
- Kaspar, K. 2020. "Motivation for Social Distancing and App Use as Complementary Measures to Combat the COVID-19 Pandemic: Quantitative Survey Study," *Journal of Medical Internet Research* (22:8). (<https://doi.org/10.2196/21613>).
- Kehr, F., Kowatsch, T., Wentzel, D., and Fleisch, E. 2015. "Blissfully Ignorant: The Effects of General Privacy Concerns, General Institutional Trust, and Affect in the Privacy Calculus," *Information Systems Journal* (25:6), pp. 607–635. (<https://doi.org/10.1111/isj.12062>).
- Kehr, F., Wentzel, D., and Mayer, P. 2013. "Rethinking the Privacy Calculus: On the Role of Dispositional Factors and Affect," in *Proceeding of the 34th International Conference on*

- Information Systems*, Milan, Italy, pp. 3355–3364.
- Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., and King, D. 2019. “Key Challenges for Delivering Clinical Impact with Artificial Intelligence,” *BMC Medicine* (17:1), pp. 1–9. (<https://doi.org/10.1186/s12916-019-1426-2>).
- Kenny, G., and Connolly, R. 2017. “Towards an Inclusive World: Exploring M-Health Adoption Across Generations,” in *Proceedings of the 25th European Conference on Information Systems*, Guimarães, Portugal, pp. 1129–1144.
- Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., and Sayres, R. 2018. “Interpretability beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV),” in *Proceedings of 35th International Conference on Machine Learning*, Stockholm, Sweden, pp. 2673–2682.
- Kim, M., Park, Jihye, Jeong, D., Park, Jaehyun, and Han, S. 2018. “TRAVEL READER:Free Independent Travel Application for Active Senior by Design Thinking Approach,” in *Proceedings of 22nd Pacific Asia Conference on Information Systems*, Yokohama, Japan.
- Komiak, S. X., and Benbasat, I. 2004. “Understanding Customer Trust in Agent-Mediated Electronic Commerce, Web-Mediated Electronic Commerce, and Traditional Commerce,” *Information Technology and Management* (5:1–2), pp. 181–207. (<https://doi.org/10.1023/B:ITEM.0000008081.55563.d4>).
- Kopfman, J. E., and Smith, S. W. 1996. “Understanding the Audiences of a Health Communication Campaign: A Discriminant Analysis of Potential Organ Donors Based on Intent to Donate,” *Journal of Applied Communication Research* (24:1), pp. 33–49. (<https://doi.org/10.1080/00909889609365438>).
- Kopp, S., Gesellensetter, L., Krämer, N. C., and Wachsmuth, I. 2005. “A Conversational Agent as Museum Guide - Design and Evaluation of a Real-World Application,” in *Proceedings of the 5th International Working Conference on Intelligent Virtual Agents*, Kos, Greece, pp. 329–343. (https://doi.org/10.1007/11550617_28).
- Koufaris, M., and Hampton-Sosa, W. 2004. “The Development of Initial Trust in an Online Company by New Customers,” *Information & Management* (41:3), pp. 377–397. (<https://doi.org/10.1016/j.im.2003.08.004>).
- Kourou, K., Exarchos, T. P., Exarchos, K. P., Karamouzis, M. V., and Fotiadis, D. I. 2015.

- “Machine Learning Applications in Cancer Prognosis and Prediction,” *Computational and Structural Biotechnology Journal* (13), pp. 8–17.
(<https://doi.org/10.1016/j.csbj.2014.11.005>).
- Krasnova, H., Hildebrand, T., and Guenther, O. 2009. “Investigating the Value of Privacy on Online Social Networks: Conjoint Analysis,” in *Proceedings of the 30th International Conference on Information Systems*, Phoenix, Arizona, United States.
- Krasnova, H., Spiekermann, S., Koroleva, K., and Hildebrand, T. 2010. “Online Social Networks: Why We Disclose,” *Journal of Information Technology* (25:2), pp. 109–125.
(<https://doi.org/10.1057/jit.2010.6>).
- Kreditech. 2018. “Through a Unique Data Scoring Algorithm Based on Machine Learning.” (<https://www.kreditech.com/>, accessed March 20, 2019).
- Kretkowski, P. D. 1998. “The 15 Percent Solution,” *WIRED*.
(<https://www.wired.com/1998/01/the-15-percent-solution/>, accessed February 5, 2021).
- von Krogh, G. 2018. “Artificial Intelligence in Organizations: New Opportunities for Phenomenon-Based Theorizing,” *Academy of Management Discoveries* (4:4), pp. 404–409. (<https://doi.org/10.5465/amd.2018.0084>).
- Kruse, L., Wunderlich, N., and Beck, R. 2019. “Artificial Intelligence for the Financial Services Industry: What Challenges Organizations to Succeed,” in *Proceedings of the 52nd Hawaii International Conference on System Sciences*, Wailea, Hawaii, United States, pp. 6408–6417. (<https://doi.org/10.24251/HICSS.2019.770>).
- Krys, C., and Fuest, K. 2017. “Roland Berger Trend Compendium 2030.” (https://www.rolandberger.com/publications/publication_pdf/roland_berger_trend_compendium_2030___trend_3_scarcity_of_resources_1.pdf).
- Kuan, R. 2019. “Adopting AI in Health Care Will Be Slow and Difficult,” *Harvard Business Review*. (<https://hbr.org/2019/10/adopting-ai-in-health-care-will-be-slow-and-difficult>).
- Kurzweil, R. 2005. *The Singularity Is near: When Humans Transcend Biology*, London, United Kingdom: Penguin.
- Labres Mallmann, G., and Gastaud Maçada, A. C. 2018. “Adoption of Cloud Computing: A Study with Public and Private Hospitals in a Developing Country,” *International Journal of Innovation and Technology Management* (15:5).
(<https://doi.org/10.1142/S021987701850044X>).

- Laufer, R. S., and Wolfe, M. 1977. "Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory," *Journal of Social Issues* (33:3), pp. 22–42. (<https://doi.org/10.1111/j.1540-4560.1977.tb01880.x>).
- Laumer, S., Maier, C., and Gubler, F. T. 2019. "Chatbot Acceptance in Healthcare: Explainin User Adoption of Conversational Agents for Disease Diagnosis," in *Proceedings of 27th European Conference on Information Systems*, Stockholm & Uppsala, Sweden.
- Lavie, D., Stettner, U., and Tushman, M. L. 2010. "Exploration and Exploitation within and across Organizations," *The Academy of Management Annals* (4:1), pp. 109–155. (<https://doi.org/10.1080/19416521003691287>).
- Law, C. C. H., and Ngai, E. W. T. 2007. "ERP Systems Adoption: An Exploratory Study of the Organizational Factors and Impacts of ERP Success," *Information and Management* (44:4), pp. 418–432. (<https://doi.org/10.1016/j.im.2007.03.004>).
- Lebovitz, S. 2019. "Diagnostic Doubt and Artificial Intelligence: An Inductive Field Study of Radiology Work," in *Proceedings of the 40th International Conference on Information Systems*, Munich, Germany.
- Lebovitz, S., Levina, N., and 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: Management Information Systems* (45:3), pp. 1501–1525. (<https://doi.org/10.25300/MISQ/2021/16564>).
- Lee, J., Liu, C., Shang, N., Jiang, X., Chen, K., Kalluri, K. S. D., Pang, C., Natarajan, K., Ryan, P., and Weng, C. 2019. "Generate the Concept Representation Using OMOP Ontology Graph," New York, New York, United States.
- Leninkumar, V. 2017. "The Relationship between Customer Satisfaction and Customer Trust on Customer Loyalty," *International Journal of Academic Research in Business and Social Sciences* (7:4), pp. 450–465. (<https://doi.org/10.6007/ijarbss/v7-i4/2821>).
- Leonard-Barton, D. 1992. "Core Capabilities and Core Rigidities: A Paradox in Managing New Product Development," *Strategic Management Journal* (13:1 S), pp. 111–125. (<https://doi.org/10.1002/smj.4250131009>).
- Letham, B., Rudin, C., McCormick, T. H., and Madigan, D. 2015. "Interpretable Classifiers Using Rules and Bayesian Analysis: Building a Better Stroke Prediction Model," *Annals*

- of Applied Statistics* (9:3), pp. 1350–1371. (<https://doi.org/10.1214/15-AOAS848>).
- Levinthal, D. A., and March, J. G. 1993. “The Myopia of Learning,” *Strategic Management Journal* (14:S2), pp. 95–112. (<https://doi.org/10.1002/smj.4250141009>).
- Lewnard, J. A., Liu, V. X., Jackson, M. L., Schmidt, M. A., Jewell, B. L., Flores, J. P., Jentz, C., Northrup, G. R., Mahmud, A., Reingold, A. L., Petersen, M., Jewell, N. P., Young, S., and Bellows, J. 2020. “Incidence, Clinical Outcomes, and Transmission Dynamics of Severe Coronavirus Disease 2019 in California and Washington: Prospective Cohort Study,” *The BMJ* (369:9). (<https://doi.org/10.1136/bmj.m1923>).
- Li, F.-C., Wang, P.-K., and Liu, Y. C. 2009. “Credit Scoring Based on Hybrid Data Mining Classification,” in *Proceedings of the 15th Americas Conference on Information Systems*, San Francisco, California, United States.
- Li, H., Luo, X. (Robert), Zhang, J., and Xu, H. 2017. “Resolving the Privacy Paradox: Toward a Cognitive Appraisal and Emotion Approach to Online Privacy Behaviors,” *Information and Management* (54:8), pp. 1012–1022. (<https://doi.org/10.1016/j.im.2017.02.005>).
- Li, J., Li, M., Wang, X., and Thatcher, J. B. 2021. “Strategic Directions for AI: The Role of Cios and Boards of Directors,” *MIS Quarterly: Management Information Systems* (45:3), pp. 1603–1643. (<https://doi.org/10.25300/MISQ/2021/16523>).
- Li, T., Zhang, Y., Gong, C., Wang, J., Liu, B., Shi, L., and Duan, J. 2020. “Prevalence of Malnutrition and Analysis of Related Factors in Elderly Patients with COVID-19 in Wuhan, China,” *European Journal of Clinical Nutrition* (74), Springer US, pp. 871–875. (<https://doi.org/10.1038/s41430-020-0642-3>).
- Lian, J.-W., Yen, D. C., and Wang, Y.-T. 2014. “An Exploratory Study to Understand the Critical Factors Affecting The decision to Adopt Cloud Computing in Taiwan Hospital,” *International Journal of Information Management* (34:1), pp. 28–36. (<https://doi.org/10.1016/j.ijinfomgt.2013.09.004>).
- Lichtenthaler, U. 2020. “Five Maturity Levels of Managing AI: From Isolated Ignorance to Integrated Intelligence,” *Journal of Innovation Management* (8:1), pp. 39–50. (https://doi.org/10.24840/2183-0606_008.001_0005).
- Lin, A., and Chen, N. C. 2012. “Cloud Computing as an Innovation: Percepation, Attitude, and Adoption,” *International Journal of Information Management* (32:6), pp. 533–540.

- (<https://doi.org/10.1016/j.ijinfomgt.2012.04.001>).
- Lipton, Z. C. 2016. “The Mythos of Model Interpretability,” *ArXiv.Org* (cs.LG), pp. 1–9.
- Liu, J., Dai, R., Wei, X. (David), and Li, Y. 2016. “Information Revelation and Customer Decision-Making Process of Repeat-Bidding Name-Your-Own-Price Auction,” *Decision Support Systems* (90), pp. 46–55. (<https://doi.org/10.1016/j.dss.2016.06.018>).
- Loewenstein, G. 1996. “Out of Control: Visceral Influences on Behavior,” *Organizational Behavior and Human Decision Processes* (65:3), pp. 272–292. (<https://doi.org/10.1006/obhd.1996.0028>).
- Loewenstein, G. 2005. “Hot-Cold Empathy Gaps and Medical Decision Making,” *Health Psychology* (24:4), pp. 49 – 56. (<https://doi.org/10.1037/0278-6133.24.4.S49>).
- Lou, B., and Wu, L. 2021. “AI on Drugs: Can Artificial Intelligence Accelerate Drug Development? Evidence from a Large-Scale Examination of Bio-Pharma Firms,” *MIS Quarterly* (45:3), pp. 1451–1482. (<https://doi.org/10.25300/MISQ/2021/16565>).
- Lowry, P. B., D’Arcy, J., Hammer, B., and Moody, G. D. 2016. “‘Cargo Cult’ Science in Traditional Organization and Information Systems Survey Research: A Case for Using Nontraditional Methods of Data Collection, Including Mechanical Turk and Online Panels,” *Journal of Strategic Information Systems* (25:3), pp. 232–240. (<https://doi.org/10.1016/j.jsis.2016.06.002>).
- Lu, B. Y., Qian, D., Fu, H., and Chen, W. 2018. “Will Supercomputers Be Super-Data and Super-AI Machines?,” *Communications of the ACM* (61:11), pp. 82–87. (<https://doi.org/10.1145/3239556>).
- Lumley, T., Diehr, P., Emerson, S., and Chen, L. 2002. “The Importance of the Normality Assumption in Large Public Health Data Sets,” *Annual Review of Public Health* (23), pp. 151–169. (<https://doi.org/10.1146/annurev.publhealth.23.100901.140546>).
- Lundberg, S. M., and Lee, S. I. 2017. “A Unified Approach to Interpreting Model Predictions,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Long Beach, California, United States, pp. 4768–4777.
- Lynch, S. 2017. “Andrew Ng: Why AI Is the New Electricity,” *Stanford Business*. (<https://www.gsb.stanford.edu/insights/andrew-ng-why-ai-new-electricity>, accessed November 17, 2021).
- Madsen, A. 2019. “Visualizing Memorization in RNNs,” *Distill* (4:3).

- (<https://doi.org/10.23915/distill.00016>).
- Malhotra, N. K., Kim, S. S., and Agarwal, J. 2004. "Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model," *Information System Research* (15:4), pp. 336–355. (<https://doi.org/10.1287/isre.1040.0032>).
- Mao, J., Gan, C., Kohli, P., Tenenbaum, J. B., and Wu, J. 2019. "The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences From Natural Supervision," in *7th International Conference on Learning Representations*, New Orleans, Louisiana, United States.
- Marabelli, M., Newell, S., and Handunge, V. 2021. "The Lifecycle of Algorithmic Decision-Making Systems: Organizational Choices and Ethical Challenges," *The Journal of Strategic Information Systems* (30:3). (<https://doi.org/10.1016/j.jsis.2021.101683>).
- March, J. G. 1991. "Exploration and Exploitation in Organizational Learning," *Organization Science* (2:1), pp. 71–87. (<https://doi.org/10.1287/orsc.2.1.71>).
- Marcinowicz, L., Pawlikowska, T., and Oleszczyk, M. 2014. "What Do Older People Value When They Visit Their General Practitioner? A Qualitative Study," *European Journal of Ageing* (11:4), pp. 361–367. (<https://doi.org/10.1007/s10433-014-0313-0>).
- Martens, D., Baesens, B., Van Gestel, T., and Vanthienen, J. 2007. "Comprehensible Credit Scoring Models Using Rule Extraction from Support Vector Machines," *European Journal of Operational Research* (183:3), pp. 1466–1476. (<https://doi.org/10.1016/j.ejor.2006.04.051>).
- Martens, D., and Provost, F. 2014. "Explaining Data-Driven Document Classifications," *MIS Quarterly* (38:1), pp. 73–99. (<https://doi.org/10.1088/1751-8113/44/8/085201>).
- McCarthy, J. 2007. "What Is Artificial Intelligence?," Stanford. (<http://jmc.stanford.edu/articles/whatisai/whatisai.pdf>).
- McKnight, D. H., Choudhury, V., and Kacmar, C. 2002. "Developing and Validating Trust Measures for E-Commerce: An Integrative Typology," *Information Systems Research* (13:3), pp. 334–359. (<https://doi.org/10.1287/isre.13.3.334.81>).
- Meskó, B., and Görög, M. 2020. "A Short Guide for Medical Professionals in the Era of Artificial Intelligence," *Npj Digital Medicine* (3:126). (<https://doi.org/10.1038/s41746-020-00333-z>).
- Metcalfe, J., and Mischel, W. 1999. "A Hot/Cool-System Analysis of Delay of Gratification:

- Dynamics of Willpower,” *Psychological Review* (106:1), pp. 3–19.
(<https://doi.org/10.1037/0033-295X.106.1.3>).
- Mettler, T., and Rohner, P. 2009. “Situational Maturity Models as Instrumental Artifacts for Organizational Design,” in *Proceedings of the 4th International Conference on Design Science Research in Information Systems and Technology*, Philadelphia, Pennsylvania, United States, pp. 1–9. (<https://doi.org/10.1145/1555619.1555649>).
- Metzger, M. J. 2004. “Privacy, Trust, and Disclosure: Exploring Barriers to Electronic Commerce,” *Journal of Computer-Mediated Communication* (9:4).
(<https://doi.org/10.1111/j.1083-6101.2004.tb00292.x>).
- Migliore, A. 2017. “On the New Regulation of Medical Devices in Europe,” *Expert Review of Medical Devices* (14:12), pp. 921–923.
(<https://doi.org/10.1080/17434440.2017.1407648>).
- Miles, M. B., and Huberman, A. M. 1994. *Qualitative Data Analysis: An Expanded Sourcebook*, (2nd ed.), Beverly Hills, California, United States: SAGE Publications.
- Miller, T. 2019. “Explanation in Artificial Intelligence: Insights from the Social Sciences,” *Artificial Intelligence* (267), pp. 1–38. (<https://doi.org/10.1016/J.ARTINT.2018.07.007>).
- Minutolo, A., Esposito, M., and De Pietro, G. 2017. “A Conversational Chatbot Based on Knowledge-Graphs for Factoid Medical Questions,” *Frontiers in Artificial Intelligence and Applications* (297), pp. 139–152. (<https://doi.org/10.3233/978-1-61499-800-6-139>).
- Mitchell, T. M. 1997. *Machine Learning*, New York, New York, United States: McGraw-Hill.
- Möllering, G. 2005. “Rational, Institutional and Active Trust: Just Do It!?” in *Trust under Pressure: Empirical Investigations of Trust and Trust Building in Uncertain Circumstances*, K. Bijlsma-Frankema and R. Klein Woolthuis (eds.), Edward Elgar Pub, pp. 17–36. (<https://doi.org/10.4337/9781845427962.00006>).
- Moncrief, W. C. 2017. “Are Sales as We Know It Dying . . . Or Merely Transforming?,” *Journal of Personal Selling & Sales Management* (37:4), pp. 271–279.
(<https://doi.org/10.1080/08853134.2017.1386110>).
- Monroe, D. 2018. “Chips for Artificial Intelligence,” *Communications of the ACM* (61:4), pp. 15–17. (<https://doi.org/10.1145/3185523>).
- Morgan, S. E., Harrison, T. R., Long, S. D., Afifi, W. A., Stephenson, M. S., and Reichert, T.

2005. "Family Discussions about Organ Donation: How the Media Influences Opinions about Donation Decisions," *Clinical Transplantation* (19:5), pp. 674–682. (<https://doi.org/10.1111/j.1399-0012.2005.00407.x>).
- MSV, J. 2018. "Here Are Three Factors That Accelerate The Rise Of Artificial Intelligence," *Forbes*. (<https://www.forbes.com/sites/janakirammsv/2018/05/27/here-are-three-factors-that-accelerate-the-rise-of-artificial-intelligence/#639dff18add9>, accessed November 19, 2018).
- Murphy, M. D., Pinheiro, D., Iyengar, R., Lim, G., Menezes, R., and Cadeiras, M. 2020. "A Data-Driven Social Network Intervention for Improving Organ Donation Awareness among Minorities: Analysis and Optimization of a Cross-Sectional Study," *Journal of Medical Internet Research* (22:1), pp. 1–15. (<https://doi.org/10.2196/14605>).
- Myers, M. D. 2013. *Qualitative Research in Business and Management*, (2nd ed.), Los Angeles, California, United States: SAGE Publications.
- Myers, M. D., and Newman, M. 2007. "The Qualitative Interview in IS Research: Examining the Craft," *Information and Organization* (17:1), pp. 2–26. (<https://doi.org/10.1016/j.infoandorg.2006.11.001>).
- Naidoo, R. 2020. "A Multi-Level Influence Model of COVID-19 Themed Cybercrime," *European Journal of Information Systems* (29:3), pp. 306–321. (<https://doi.org/10.1080/0960085X.2020.1771222>).
- Nan, N., and Tanriverdi, H. 2017. "Unifying the Role of IT in Hyperturbulence and Competitive Advantage Via a Multilevel Perspective of IS Strategy," *MIS Quarterly* (41:3), pp. 937–958. (<https://doi.org/10.25300/misq/2017/41.3.12>).
- Nascimento, A. M., da Cunha, M. A. V. C., de Souza Meirelles, F., Scornavacca, E., and de Melo, V. V. 2018. "A Literature Analysis of Research on Artificial Intelligence in Management Information System (MIS)," in *Proceedings of the 24th Americas Conference on Information Systems*, New Orleans, Louisiana, United States.
- Nguyen, T. T. H., Tapanainen, T., and Obi, T. 2014. "A Review of Information and Communication Technology (ICT) Training for Elderly People-Toward Recommendations for Developing Countries," in *Proceedings of the 18th Pacific Asia Conference on Information Systems*, Chengdu, China.
- Norberg, P. A., and Horne, D. R. 2007. "Privacy Attitudes and Privacy-Related Behavior,"

- Psychology & Marketing* (24:10), pp. 829–847. (<https://doi.org/10.1002/mar.20186>).
- Norberg, P. A., Horne, D. R., and Horne, D. A. 2007. “The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors,” *Journal of Consumer Affairs* (41:1), pp. 100–126. (<https://doi.org/10.1111/j.1745-6606.2006.00070.x>).
- O’Guinn, T. C., and Faber, R. J. 1989. “Compulsive Buying: A Phenomenological Exploration,” *Journal of Consumer Research* (16:2), p. 147. (<https://doi.org/10.1086/209204>).
- Olah, C., Satyanarayan, A., Johnson, I., Carter, S., Schubert, L., Ye, K., and Mordvintsev, A. 2018. “The Building Blocks of Interpretability,” *Distill* (3:3). (<https://doi.org/10.23915/distill.00010>).
- Pachur, T., Hertwig, R., and Steinmann, F. 2012. “How Do People Judge Risks: Availability Heuristic, Affect Heuristic, or Both?,” *Journal of Experimental Psychology: Applied* (18:3), pp. 314–330. (<https://doi.org/10.1037/a0028279>).
- Panch, T., Mattie, H., and Celi, L. A. 2019. “The ‘Inconvenient Truth’ about AI in Healthcare,” *Npj Digital Medicine* (2:1), pp. 4–6. (<https://doi.org/10.1038/s41746-019-0155-4>).
- Panetta, K. 2018. “Gartner Top Strategic Predictions for 2019 and Beyond,” *Gartner*. (<https://www.gartner.com/smarterwithgartner/gartner-top-strategic-predictions-for-2019-and-beyond/>, accessed May 29, 2019).
- Paton, C., and Kobayashi, S. 2019. “An Open Science Approach to Artificial Intelligence in Healthcare,” *Yearbook of Medical Informatics* (28:1), pp. 47–51. (<https://doi.org/10.1055/s-0039-1677898>).
- Paulk, M. C., Curtis, B., Chrissis, M. B., and Weber, C. V. 1993. “Capability Maturity Model, Version 1.1,” *IEEE Software* (10:4), pp. 18–27. (<https://doi.org/10.1201/9781420052862.axf>).
- Pfeuffer, N., Benlian, A., Gimpel, H., and Hinz, O. 2019. “Anthropomorphic Information Systems,” *Business and Information Systems Engineering* (61:4), Springer Fachmedien Wiesbaden, pp. 523–533. (<https://doi.org/10.1007/s12599-019-00599-y>).
- Ply, J. K., Moore, E., Williams, C. K., and Thatcher, J. B. 2012. “IS Employee Attitudes and Perceptions at Varying Levels of Software Process Maturity,” *MIS Quarterly* (36:2), pp. 601–624. (<https://doi.org/10.2307/41703469>).

- Poba-Nzaou, P., Uwizeyemungu, S., Raymond, L., and Paré, G. 2014. "Motivations Underlying the Adoption of ERP Systems in Healthcare Organizations: Insights from Online Stories," *Information Systems Frontiers* (16:4), pp. 591–605. (<https://doi.org/10.1007/s10796-012-9361-1>).
- Podsakoff, P. M., MacKenzie, S. B., Lee, J. Y., and Podsakoff, N. P. 2003. "Common Method Biases in Behavioral Research: A Critical Review of the Literature and Recommended Remedies," *Journal of Applied Psychology* (88:5), pp. 879–903. (<https://doi.org/10.1037/0021-9010.88.5.879>).
- Poepplbuss, J., Niehaves, B., Simons, A., and Becker, J. 2011. "Maturity Models in Information Systems Research: Literature Search and Analysis," *Communications of the Association for Information Systems* (29), pp. 505–532. (<https://doi.org/10.17705/1cais.02927>).
- Porter, M. E. 1985. "Technology and Competitive Advantage," *The Journal of Business Strategy* (5:3), pp. 60–78. (<https://doi.org/10.1108/eb039075>).
- Prasser, F., Kohlmayer, F., Lautenschläger, R., and Kuhn, K. A. 2014. "ARX - A Comprehensive Tool for Anonymizing Biomedical Data," in *Proceedings of Annual Symposium American Medical Informatics Association*, Washington, D.C., United States, pp. 984–993.
- Pumplun, L., Fecho, M., Wahl, N., Peters, F., and Buxmann, P. 2021. "Maturity Model - Determine Your Clinic's Readiness for ML-Supported Diagnostics." (<https://ml-maturity-clinics.netlify.app/>, accessed September 23, 2021).
- Pumplun, L., Tauchert, C., and Heidt, M. 2019. "A New Organizational Chassis for Artificial Intelligence - Exploring Organizational Readiness Factors," in *Proceedings of the 27th European Conference on Information Systems*, Stockholm & Uppsala, Sweden.
- Qiu, L., and Benbasat, I. 2009. "Evaluating Anthropomorphic Product Recommendation Agents: A Social Relationship Perspective to Designing Information Systems," *Journal of Management Information Systems* (25:4), pp. 145–182. (<https://doi.org/10.2753/MIS0742-1222250405>).
- Qureshi, I., Hom, H., Kong, H., and Compeau, D. 2009. "Assessing Between-Group Differences in Information Systems Research: A Comparison of Covariance- and Component-Based SEM," *MIS Quarterly* (33:1), pp. 197–214. (<https://doi.org/10.2307/20650285>).

- Raber, D., Winter, R., and Wortmann, F. 2012. "Using Quantitative Analyses to Construct a Capability Maturity Model for Business Intelligence," in *Proceedings of the 45th Hawaii International Conference on System Sciences*, Maui, Hawaii, United States, pp. 4219–4228.
- Raber, D., Wortmann, F., and Winter, R. 2013. "Towards the Measurement of Business Intelligence Maturity," in *Proceedings of the 21st European Conference on Information Systems*, Utrecht, Netherlands.
- Raghupathi, W., and Raghupathi, V. 2014. "Big Data Analytics in Healthcare: Promise and Potential," *Health Information Science and Systems* (2:1). (<https://doi.org/10.1186/2047-2501-2-3>).
- Rai, A., Constantinides, P., and Sarker, S. 2019. "Editor's Comments: Next-Generation Digital Platforms: Toward Human-AI Hybrids," *MIS Quarterly* (43:1), iii–ix.
- Raisch, S., and Krakowski, S. 2021. "Artificial Intelligence and Management: The Automation-Augmentation Paradox," *Academy of Management Review* (46:1), pp. 192–210. (<https://doi.org/10.5465/2018.0072>).
- Rana, N. P., Chatterjee, S., Dwivedi, Y. K., and Akter, S. 2021. "Understanding Dark Side of Artificial Intelligence (AI) Integrated Business Analytics: Assessing Firm's Operational Inefficiency and Competitiveness," *European Journal of Information Systems* (00:00), Taylor & Francis, pp. 1–24. (<https://doi.org/10.1080/0960085X.2021.1955628>).
- Rana, R., Staron, M., Hansson, J., Nilsson, M., and Meding, W. 2014. "A Framework for Adoption of Machine Learning in Industry for Software Defect Prediction," in *Proceedings of the 9th International Conference on Software Engineering and Applications*, Vienna, Austria, pp. 383–392. (<https://doi.org/10.5220/0005099303830392>).
- Ransbotham, S., Khodabandeh, S., Kiron, D., Candelon, F., Chu, M., and Lafountain, B. 2020. "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., and Reeves, M. 2017. "Reshaping Business With Artificial Intelligence," *MIT Sloan Management Review*. (<https://sloanreview.mit.edu/projects/reshaping-business-with-artificial-intelligence/>).

- Reese, H. 2016. "Top 10 AI Failures of 2016." (<https://www.techrepublic.com/article/top-10-ai-failures-of-2016/>, accessed November 20, 2018).
- Renaud, K., and Biljon, J. Van. 2008. "Senior Technology Acceptance & Adoption Model S-TAM," in *Proceedings of the 2008 Annual Research Conference of the South African Institute of Computer Scientists and Information Technologists on IT Research in Developing Countries*, Wilderness, South Africa, pp. 210–219.
- Reynolds, T. J., and Olson, J. C. 2001. *Understanding Consumer Decision Making: The Means-End Approach to Marketing and Advertising Strategy*, Mahwah, New Jersey, United States: Psychology Press.
- Ribeiro, M. T., Singh, S., and Guestrin, C. 2016. "Why Should I Trust You?: Explaining the Predictions of Any Classifier," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, California, United States, February 16, pp. 1135–1144. (<https://doi.org/10.1145/2939672.2939778>).
- RKI. 2020. "Corona Data Donation Blog." (<https://corona-datenspende.de/science/en/>, accessed September 29, 2020).
- Rogers, E. M. 1995. *Diffusion of Innovations*, (4th ed.), New York, New York, United States: Free Press.
- Roski, J., Chapman, W., Heffner, J., Hutchinson, F., Trivedi, R., Del Fiore, G., Kukafka, R., Bleicher, P., Estiri, H., Klann, J., and Pierce, J. 2019. "How Artificial Intelligence Is Changing Health and Health Care," in *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril* (1st ed.), M. Matheny, S. Israni, M. Ahmed, and D. Whicker (eds.), Washington, D.C., United States: National Academy of Medicine, pp. 59–88.
- Rowe, F., Ngwenyama, O., and Richet, J.-L. 2020. "Contact-Tracing Apps and Alienation in the Age of COVID-19," *European Journal of Information Systems* (29:5), pp. 545–562. (<https://doi.org/10.1080/0960085X.2020.1803155>).
- Rudin, C. 2019. "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," *Nature Machine Intelligence* (1:5), pp. 206–215. (<https://doi.org/10.1038/s42256-019-0048-x>).
- Russell, S. J., and Norvig, P. 2016. *Artificial Intelligence: A Modern Approach*, (3rd ed.), Boston, Massachusetts, United States: Addison Wesley.

- Rzepka, C., and Berger, B. 2018. "User Interaction with AI-Enabled Systems: A Systematic Review of IS Research," in *Proceedings of 39th International Conference on Information Systems*, San Francisco, California, United States.
- Saldaña, J. 2009. *The Coding Manual for Qualitative Researchers*, (1st ed.), London, United Kingdom: SAGE Publications.
- Sandhu, S., Lin, A. L., Brajer, N., Sperling, J., Ratliff, W., Bedoya, A. D., Balu, S., O'Brien, C., and Sendak, M. P. 2020. "Integrating a Machine Learning System into Clinical Workflows: Qualitative Study," *Journal of Medical Internet Research* (22:11). (<https://doi.org/10.2196/22421>).
- Sanyal, S. 2018. "How Is AI Revolutionizing Elderly Care," *Forbes*. (<https://www.forbes.com/sites/shourjyasanyal/2018/10/31/how-is-ai-revolutionizing-elderly-care/#44775f8e07d6>, accessed May 29, 2019).
- Sarker, S., Xiao, X., and Beaulieu, T. 2013. "Qualitative Studies in Information Systems: A Critical Review and Some Guiding Principles," *MIS Quarterly* (37:4), iii–xviii.
- Schneider, T., and Tkachenko, O. 2018. "Towards Efficient Privacy-Preserving Similar Sequence Queries on Outsourced Genomic Databases," in *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*, Toronto, Ontario, Canada, pp. 71–75.
- Schreiner, M., and Hess, T. 2015. "Why Are Consumers Willing to Pay for Privacy? An Application of the Privacy-Freemium Model to Media Companies," in *Proceedings of the 23rd European Conference on Information Systems*, Münster, Germany.
- Seeger, A.-M., Pfeiffer, J., and Heinzl, A. 2018. "Designing Anthropomorphic Conversational Agents: Development and Empirical Evaluation of a Design Framework," in *Proceedings of the 39th International Conference on Information Systems*, San Francisco, California, United States.
- Seeger, A. M., Pfeiffer, J., and Heinzl, A. 2021. "Texting with Humanlike Conversational Agents: Designing for Anthropomorphism," *Journal of the Association for Information Systems* (22:4), pp. 931–967. (<https://doi.org/10.17705/1jais.00685>).
- Shahid, N., Rappon, T., and Berta, W. 2019. "Applications of Artificial Neural Networks in Health Care Organizational Decision-Making: A Scoping Review," *PLoS ONE* (14:2), pp. 1–22. (<https://doi.org/10.1371/journal.pone.0212356>).
- Shaw, J., Rudzicz, F., Jamieson, T., and Goldfarb, A. 2019. "Artificial Intelligence and the

- Implementation Challenge,” *Journal of Medical Internet Research* (21:7). (<https://doi.org/10.2196/13659>).
- Shawar, B. A., Atwell, E., and Roberts, A. 2005. “FAQchat as an Information Retrieval System,” in *Proceedings of the 2nd Language and Technology Conference*, Poznań, Poland, pp. 274–278.
- Shook, J., Smith, R., and Antonio, A. 2018. “Transparency and Fairness in Machine Learning Applications,” *Texas A&M Journal of Property Law*, pp. 443–462. (<https://doi.org/10.3868/s050-004-015-0003-8>).
- Shrestha, L. B. 2000. “Population Aging in Developing Countries,” *Health Affairs* (19:3), pp. 204–212. (<https://doi.org/10.1377/hlthaff.19.3.204>).
- Shyu, Y. I. L., Tang, W. R., Tsai, W. C., Liang, J., and Chen, M. C. 2006. “Emotional Support Levels Can Predict Physical Functioning and Health Related Quality of Life Among Elderly Taiwanese With Hip Fractures,” *Osteoporosis International* (17:4), Springer, pp. 501–506. (<https://doi.org/10.1007/s00198-005-0020-y>).
- Sidorova, A., and Rafiee, D. 2019. “AI Agency Risks and Their Mitigation through Business Process Management: A Conceptual Framework,” in *Proceedings of the Hawaii International Conference on System Sciences*, Wailea, Hawaii, United States, pp. 5837–5845.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., Van Den Driessche, G., Graepel, T., and Hassabis, D. 2017. “Mastering the Game of Go without Human Knowledge,” *Nature* (550:7676), pp. 354–359. (<https://doi.org/10.1038/nature24270>).
- Siminoff, L., Mercer, M. B., Graham, G., and Burant, C. 2007. “The Reasons Families Donate Organs for Transplantation: Implications for Policy and Practice,” *Journal of Trauma* (62:4), pp. 969–978. (<https://doi.org/10.1097/01.ta.0000205220.24003.51>).
- Skatova, A., and Goulding, J. 2019. “Psychology of Personal Data Donation,” *PLoS ONE* (14:11). (<https://doi.org/10.1371/journal.pone.0224240>).
- Skatova, A., Ng, E., and Goulding, J. 2014. “Data Donation: Sharing Personal Data for Public Good?,” in *Proceedings of 5th Digital Economy All Hands Meeting*, London, United Kingdom. (<https://doi.org/10.13140/2.1.2567.8405>).

- Slovic, P., Finucane, M. L., Peters, E., and MacGregor, D. G. 2007. "The Affect Heuristic," *European Journal of Operational Research* (177:3), pp. 1333–1352. (<https://doi.org/10.1016/j.ejor.2005.04.006>).
- Slovic, P., Finucane, M., Peters, E., and MacGregor, D. G. 2002. "Rational Actors or Rational Fools: Implications of the Affect Heuristic for Behavioral Economics," *The Journal of Socio-Economics* (31:4), pp. 329–342. ([https://doi.org/10.1016/S1053-5357\(02\)00174-9](https://doi.org/10.1016/S1053-5357(02)00174-9)).
- Smith, J. H., Dinev, T., and Xu, H. 2011. "Information Privacy Research: An Interdisciplinary Review," *MIS Quarterly* (35:4), pp. 989–1015. (<https://doi.org/10.2307/41409970>).
- Smith, S. P., Johnston, R. B., and Howard, S. 2011. "Putting Yourself in the Picture: An Evaluation of Virtual Model Technology as an Online Shopping Tool," *Information Systems Research* (22:3), pp. 640–659. (<https://doi.org/10.1287/isre.1090.0279>).
- Solís, W. V., Cedillo, P., Parra, J., Guevara, A., and Ortiz, J. 2017. "Intelligent Pillbox: Evaluating the User Perceptions of Elderly People," in *Proceedings of the 26th International Conference on Information Systems Development*, Larnaca, Cyprus.
- Spann, M., Skiera, B., and Schäfers, B. 2004. "Measuring Individual Frictional Costs and Willingness-to-Pay via Name-Your-Own-Price Mechanisms," *Journal of Interactive Marketing* (18:4), pp. 22–36. (<https://doi.org/10.1002/dir.20022>).
- Steele, R., Lo, A., Secombe, C., and Wong, Y. K. 2009. "Elderly Persons' Perception and Acceptance of Using Wireless Sensor Networks to Assist Healthcare," *International Journal of Medical Informatics* (78), pp. 788–801. (<https://doi.org/10.1016/j.ijmedinf.2009.08.001>).
- Stogdill, R. M. 1950. "Leadership, Membership, and Organization," *Psychological Bulletin* (47:1), pp. 1–14. (<http://www.ncbi.nlm.nih.gov/pubmed/19586159>).
- Strich, F., Mayer, A. S., and Fiedler, M. 2021. "What Do i Do in a World of Artificial Intelligence? Investigating the Impact of Substitutive Decision-Making Ai Systems on Employees' Professional Role Identity," *Journal of the Association for Information Systems* (22:2), pp. 304–324. (<https://doi.org/10.17705/1jais.00663>).
- Suh, B., and Han, I. 2002. "Effect of Trust on Customer Acceptance of Internet Banking," *Electronic Commerce Research and Applications* (1:3–4), pp. 247–263. ([https://doi.org/10.1016/S1567-4223\(02\)00017-0](https://doi.org/10.1016/S1567-4223(02)00017-0)).

- Sulaiman, H., and Wickramasinghe, N. 2014. "Assimilating Healthcare Information Systems in a Malaysian Hospital," *Communications of the Association for Information Systems* (34:1), pp. 1291–1318. (<https://doi.org/10.17705/1cais.03477>).
- Sun, J., and Qu, Z. 2015. "Understanding Health Information Technology Adoption: A Synthesis of Literature from an Activity Perspective," *Information Systems Frontiers* (17:5), pp. 1177–1190. (<https://doi.org/10.1007/s10796-014-9497-2>).
- Sun, T. Q., and Medaglia, R. 2019. "Mapping the Challenges of Artificial Intelligence in the Public Sector: Evidence from Public Healthcare," *Government Information Quarterly* (36:2), Elsevier, pp. 368–383.
- Taylor, K. 2015. "Connected Health: How Digital Technology Is Transforming Health and Social Care," London, United Kingdom.
(https://www2.deloitte.com/content/dam/Deloitte/pl/Documents/Reports/pl_deloitte-uk-connected-health.pdf).
- Taylor, S., and Todd, P. A. 1995. "Understanding Information Technology Usage: A Test of Competing Models," *Information Systems Research* (6:2), pp. 144–176.
(<https://doi.org/10.1287/isre.6.2.144>).
- Templeton, G. F. 2011. "A Two-Step Approach for Transforming Continuous Variables to Normal: Implications and Recommendations for IS Research," *Communications of the Association for Information Systems* (28:1), pp. 41–58.
(<https://doi.org/10.17705/1cais.02804>).
- The Economist. 2017. "Regulating the Internet Giants: The World's Most Valuable Resource Is No Longer Oil, but Data." (<https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data>, accessed November 19, 2018).
- The Economist. 2018. "For Artificial Intelligence to Thrive, It Must Explain Itself." (<https://www.economist.com/science-and-technology/2018/02/15/for-artificial-intelligence-to-thrive-it-must-explain-itself>, accessed March 20, 2019).
- Thompson, J. M., Buchbinder, S. B., and Shanks, N. H. 2017. "An Overview of Healthcare Management," in *Introduction to Health Care Management* (3rd ed.), S. B. Buchbinder and N. H. Shanks (eds.), pp. 2–16.
- Thrall, J. H., Li, X., Li, Q., Cruz, C., Do, S., Dreyer, K., and Brink, J. 2018. "Artificial Intelligence and Machine Learning in Radiology: Opportunities, Challenges, Pitfalls, and

- Criteria for Success.,” *Journal of the American College of Radiology* (15:3), pp. 504–508. (<https://doi.org/10.1016/j.jacr.2017.12.026>).
- Topol, E. J. 2019. “High-Performance Medicine: The Convergence of Human and Artificial Intelligence,” *Nature Medicine* (25:1), pp. 44–56. (<https://doi.org/10.1038/s41591-018-0300-7>).
- Tresp, V., Marc Overhage, J., Bundschuh, M., Rabizadeh, S., Fasching, P. A., and Yu, S. 2016. “Going Digital: A Survey on Digitalization and Large-Scale Data Analytics in Healthcare,” *Proceedings of the IEEE* (104:11), pp. 2180–2206. (<https://doi.org/10.1109/JPROC.2016.2615052>).
- Tsouna-Hadjis, E., Vemmos, K. N., Zakopoulos, N., and Stamatelopoulos, S. 2000. “First-Stroke Recovery Process: The Role of Family Social Support,” *Archives of Physical Medicine and Rehabilitation* (81:7), pp. 881–887. (<https://doi.org/10.1053/apmr.2000.4435>).
- United Nations. 2008. “International Standard Industrial Classification of All Economic Activities (ISIC),” New York, New York, United States. (https://unstats.un.org/unsd/publication/seriesm/seriesm_4rev4e.pdf).
- United Nations. 2017. “World Population Ageing,” New York, New York, United States. (https://www.un.org/en/development/desa/population/publications/pdf/ageing/WPA2017_Highlights.pdf).
- Urquhart, C., Lehmann, H., and Myers, M. D. 2010. “Putting the ‘Theory’ Back into Grounded Theory: Guidelines for Grounded Theory Studies in Information Systems,” *Information Systems Journal* (20:4), pp. 357–381. (<https://doi.org/10.1111/j.1365-2575.2009.00328.x>).
- Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., Li, B., Madabhushi, A., Shah, P., Spitzer, M., and Zhao, S. 2019. “Applications of Machine Learning in Drug Discovery and Development,” *Nature Reviews Drug Discovery* (18:6), pp. 463–477. (<https://doi.org/10.1038/s41573-019-0024-5>).
- Varghese, J. 2020. “Artificial Intelligence in Medicine: Chances and Challenges for Wide Clinical Adoption,” *Visceral Medicine*, pp. 443–449. (<https://doi.org/10.1159/000511930>).
- Veale, M., Van Kleek, M., and Binns, R. 2018. “Fairness and Accountability Design Needs

- for Algorithmic Support in High-Stakes Public Sector Decision-Making,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montréal, Canada, p. 440. (<https://doi.org/10.1145/3173574.3174014>).
- Venkatesh, V., Davis, F. D., and Morris, M. G. 2007. “Dead or Alive? The Development, Trajectory and Future of Technology Adoption Research,” *Journal of the Association for Information Systems* (8:4), pp. 267–286. (<https://doi.org/10.17705/1jais.00120>).
- Venkatesh, V., Morris, M., Gordon, B., and Davis, F. 2003. “User Acceptance of Information Technology: Toward a Unified View,” *MIS Quarterly* (27:3), pp. 425–478. (<https://doi.org/10.2307/30036540>).
- Venkatesh, V., Thong, J. Y., and Xu, X. 2012. “Consumer Acceptance and Use of Information Technology: Extending the Unified Theory of Acceptance and Use of Technology,” *MIS Quarterly* (36:1), pp. 157–178. (<https://doi.org/10.2307/41410412>).
- Vichitvanichphong, S., Talaei-Khoei, A., and Kerr, D. 2017. “Elderly’s Perception about The Value of Assistive Technologies for Their Daily Living: Impacting Factors and Theoretical Support,” in *Proceedings of the 50th Hawaii International Conference on System Sciences*, Waikoloa Village, Hawaii, United States, pp. 3678–3685.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., Vezhnevets, A. S., Leblond, R., Pohlen, T., Dalibard, V., Budden, D., Sulsky, Y., Molloy, J., Paine, T. L., Gulcehre, C., Wang, Z., Pfaff, T., Wu, Y., Ring, R., Yogatama, D., Wünsch, D., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Kavukcuoglu, K., Hassabis, D., Apps, C., and Silver, D. 2019. “Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning,” *Nature* (575:7782), pp. 350–354. (<https://doi.org/10.1038/s41586-019-1724-z>).
- Wagner, A., and Mesbah, N. 2019. “Too Confident to Care: Investigating Overconfidence in Privacy Decision-Making,” in *Proceedings of the 42th European Conference on Information Systems*, Stockholm & Uppsala, Sweden.
- Wagner, A., Wessels, N., Buxmann, P., and Krasnova, H. 2018. “Putting a Price Tag on Personal Information - a Literature Review,” in *Proceedings of the 51st Hawaii International Conference on System Sciences*, Waikoloa Village, Hawaii, United States, pp. 3760–3769.

- Wakefield, R. 2013. "The Influence of User Affect in Online Information Disclosure," *Journal of Strategic Information Systems* (22:2), pp. 157–174. (<https://doi.org/10.1016/j.jsis.2013.01.003>).
- Wang, H., Zhou, Z., Li, Y., Chen, Z., Lu, P., Wang, W., Liu, W., and Yu, L. 2017. "Comparison of Machine Learning Methods for Classifying Mediastinal Lymph Node Metastasis of Non-Small Cell Lung Cancer from 18F-FDG PET/CT Images," *EJNMMI Research* (7:11), EJNMMI Research. (<https://doi.org/10.1186/s13550-017-0260-9>).
- Wang, W., and Benbasat, I. 2016. "Empirical Assessment of Alternative Designs for Enhancing Different Types of Trusting Beliefs in Online Recommendation Agents," *Journal of Management Information Systems* (33:3), pp. 744–775. (<https://doi.org/10.1080/07421222.2016.1243949>).
- Wang, W., Qiu, L., Kim, D., and Benbasat, I. 2016. "Effects of Rational and Social Appeals of Online Recommendation Agents on Cognition- and Affect-Based Trust," *Decision Support Systems* (86), pp. 48–60. (<https://doi.org/10.1016/j.dss.2016.03.007>).
- Wang, W., and Siau, K. 2018. "Trust in Health Chatbots," in *Proceedings of the 39th International Conference on Information Systems*, San Francisco, California, United States.
- Wang, X., Sun, J., Wang, Y., and Liu, Y. 2021. "Deepen Electronic Health Record Diffusion beyond Breadth: Game Changers and Decision Drivers," *Information Systems Frontiers*. (<https://doi.org/10.1007/s10796-020-10093-6>).
- Watson, D., and Clark, L. A. 1999. *The PANAS-X: Manual for the Positive and Negative Affect Schedule*, The University of Iowa, pp. 1–24. (https://doi.org/10.1007/978-94-007-0753-5_2212).
- Weber, R. P. 1990. *Basic Content Analysis*, (2nd ed.), Newbury Park, California, United States: SAGE Publications.
- Weissman, G. E., Crane-Droesch, A., Chivers, C., Luong, T. B., Hanish, A., Levy, M. Z., Lubken, J., Becker, M., Draugelis, M. E., Anesi, G. L., Brennan, P. J., Christie, J. D., Hanson, C. W., Mikkelsen, M. E., and Halpern, S. D. 2020. "Locally Informed Simulation to Predict Hospital Capacity Needs During the COVID-19 Pandemic," *Annals of Internal Medicine* (173:1), pp. 21–28. (<https://doi.org/10.7326/M20-1260>).
- Weitzman, E. R., Kaci, L., and Mandl, K. D. 2010. "Sharing Medical Data for Health

- Research: The Early Personal Health Record Experience,” *Journal of Medical Internet Research* (12:2). (<https://doi.org/10.2196/jmir.1356>).
- Weizenbaum, J. 1966. “ELIZA - A Computer Program For the Study of Natural Language Communication Between Man and Machine,” *Communications of the ACM* (9:1), pp. 36–71. (<https://doi.org/10.1145/365153.365168>).
- Wessels, N., Wagner, A., Sarswat, J. P., and Buxmann, P. 2019. “What Is Your Selfie Worth? A Field Study on Individuals’ Valuation of Personal Data,” in *Proceedings of the 14th International Conference on Wirtschaftsinformatik*, Siegen, Germany, pp. 1322–1336.
- Whittaker, R. 2012. “Issues in MHealth: Findings from Key Informant Interviews,” *Journal of Medical Internet Research* (14:5). (<https://doi.org/10.2196/jmir.1989>).
- Willcocks, L. 2020. “Robo-Apocalypse Cancelled? Reframing the Automation and Future of Work Debate,” *Journal of Information Technology* (35:4), pp. 286–302. (<https://doi.org/10.1177/0268396220925830>).
- Wilson, C. 2018. “Is It Love or Loneliness? Exploring the Impact of Everyday Digital Technology Use on the Well-Being of Older Adults,” *Ageing Society* (38:7), pp. 1307–1331. (<https://doi.org/10.1017/s0144686x16001537>).
- Wischnevsky, J. D., Damanpour, F., and Méndez, F. A. 2011. “Influence of Environmental Factors and Prior Changes on the Organizational Adoption of Changes in Products and in Technological and Administrative Processes,” *British Journal of Management* (22:1), pp. 132–149. (<https://doi.org/10.1111/j.1467-8551.2010.00700.x>).
- World Health Organization. 2016. “Global Strategy on Human Resources for Health: Workforce 2030,” Geneva, Switzerland. (<https://apps.who.int/iris/handle/10665/250368>).
- World Health Organization. 2019. “Risk Factors of Ill Health Among Older People,” World Health Organization. (<http://www.euro.who.int/en/health-topics/Life-stages/healthy-ageing/data-and-statistics/risk-factors-of-ill-health-among-older-people>, accessed June 16, 2019).
- Wu, J., and Liu, D. 2007. “The Effects of Trust and Enjoyment on Intention to Play Online Games,” *Journal of Electronic Commerce Research* (8:2), pp. 128–140.
- Wu, J., Zan, X., Gao, L., Zhao, J., Fan, J., Shi, H., Wan, Y., Yu, E., Li, S., and Xie, X. 2019. “A Machine Learning Method for Identifying Lung Cancer Based on Routine Blood Indices: Qualitative Feasibility Study,” *JMIR Med Inform* (7:3), p. e13476.

- (<https://doi.org/10.2196/13476>).
- Wulf, F., Westner, M., and Strahringer, S. 2021. "Cloud Computing Adoption: A Literature Review on What Is New and What Still Needs to Be Addressed," *Communications of the Association for Information Systems* (48), pp. 523–561. (<https://doi.org/10.17705/1CAIS.04843>).
- Xu, H., Teo, H. H., Tan, B. C. Y., and Agarwal, R. 2009. "The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services," *Journal of Management Information Systems* (26:3), pp. 135–174. (<https://doi.org/10.2753/MIS0742-1222260305>).
- Xu, K., Ba, J. L., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R. S., and Bengio, Y. 2015. "Show, Attend and Tell: Neural Image Caption Generation with Visual Attention," in *Proceedings of the 32nd International Conference on Machine Learning*, Lille, France, pp. 2048–2057.
- Yan, D., Zhou, Q., Wang, J., and Zhang, N. 2016. "Bayesian Regularisation Neural Network Based on Artificial Intelligence Optimisation," *International Journal of Production Research* (55:8), pp. 2266–2287. (<https://doi.org/10.1080/00207543.2016.1237785>).
- Yang, Z., Sun, J., Zhang, Y., and Wang, Y. 2015. "Understanding SaaS Adoption from the Perspective of Organizational Users: A Tripod Readiness Model," *Computers in Human Behavior* (45), pp. 254–264. (<https://doi.org/10.1016/j.chb.2014.12.022>).
- Ye, R. L., and Johnson, P. E. 1995. "The Impact of Explanation Facilities on User Acceptance of Expert Systems Advice," *MIS Quarterly* (19:2), pp. 157–172. (<https://doi.org/10.2307/249686>).
- Your.MD. 2019. "Your.MD - Health Guide and Symptom Checker." (<https://www.your.md/>, accessed May 28, 2019).
- Zhang, B., and Li, Y. 2017. "Wearable Medical Devices Acceptance and Expectance of Senior Population in China," in *Proceedings of The 17th International Conference on Electronic Business*, Dubai, United Arab Emirates, pp. 241–251.
- Zhang, Y., Qiu, M., Tsai, C. W., Hassan, M. M., and Alamri, A. 2017. "Health-CPS: Healthcare Cyber-Physical System Assisted by Cloud and Big Data," *IEEE Systems Journal* (11:1), pp. 88–95. (<https://doi.org/10.1109/JSYST.2015.2460747>).
- Zhao, X., Lynch, J. G., and Chen, Q. 2010. "Reconsidering Baron and Kenny: Myths and

- Truths about Mediation Analysis,” *Journal of Consumer Research* (37:2), pp. 197–206. (<https://doi.org/10.1086/651257>).
- Zhou, J. 2017. “Data Mining for Individual Consumer Credit Default Prediction under E-Commence Context: A Comparative Study,” in *Proceedings of 38th International Conference on Information Systems*, Seoul, South Korea, p. 22.
- Zhu, K., Dong, S., Xu, S. X., and Kraemer, K. L. 2006. “Innovation Diffusion in Global Contexts: Determinants of Post-Adoption Digital Transformation of European Companies,” *European Journal of Information Systems* (15:6), pp. 601–616. (<https://doi.org/10.1057/palgrave.ejis.3000650>).
- Zhu, K., and Kraemer, K. L. 2005. “Post-Adoption Variations in Usage and Value of E-Business by Organizations: Cross-Country Evidence from the Retail Industry,” *Information Systems Research* (16:1), pp. 61–84. (<https://doi.org/10.1287/isre.1050.0045>).
- Zhu, K., Kraemer, K., and Xu, S. 2003. “Electronic Business Adoption by European Firms: A Cross-Country Assessment of the Facilitators and Inhibitors,” *European Journal of Information Systems* (12:4), pp. 251–268. (<https://doi.org/10.1057/palgrave.ejis.3000475>).
- Zoabi, Y., Deri-Rozov, S., and Shomron, N. 2021. “Machine Learning-Based Prediction of COVID-19 Diagnosis Based on Symptoms,” *Npj Digital Medicine* (4:1), Springer US, pp. 1–5. (<https://doi.org/10.1038/s41746-020-00372-6>).