Q

# I. Introduction

The introduction to this dissertation comprises the motivation for my research, the description of the research gaps, the structure of the dissertation, and outlines of the contexts and designs of the studies carried out in Part B. In Section A.I.5, I conclude this chapter by describing the anticipated contributions for research and practice.

# I.1 Motivation

Purposeful decision making in organizational contexts is an essential competence that determines the path of an organization and its positioning within an industry. The amount of consideration put into making decisions can make the difference between success and failure of organizations. Top-performing organizations "make decisions based on rigorous analysis at more than double the rate of lower performers" (LaValle et al. 2011, p. 22) and are significantly less likely to rely on intuition (LaValle et al. 2011).

For organizations, the transformational value of information technology (IT) has been long discovered (Gregor et al. 2006). IT-enabled organizational transformation may vield "productivity increases by reducing costs and, more importantly, by allowing organizations to increase output quality, along with offering new products and improved customer service" (Gregor et al. 2006, p. 250). More recently, algorithmic decision making (Newell and Marabelli 2015), as a consequence of IT-related progresses, has had successful applications. As a specific instance of algorithmic analyses, the transformational impact of business intelligence and data-driven analyses on organizational decision making processes has been recognized (Sharma et al. 2014). Researchers highlight their great potential for obtaining "insights ... from the highly detailed, contextualized, and rich contents of relevance to any business or organization" (Chen et al. 2012, p. 1168), especially due to the progressing digitalization across all areas of life and the datafication (Galliers et al. 2017) that steadily increase the availability of data as well as the means of collecting them (Berente and Seidel 2014). However, while decision making is one of the main areas affected by business intelligence, it is also clear that data-driven decision making does not automatically guarantee increased value for an organization (Sharma et al. 2014). For an IT- or business-intelligence-enabled transformation, it is also necessary to apply changes to the processes and the structure of an organization to accommodate new technology and capture its potential benefits (Sharma et al. 2014).

With the emergence of big data, researchers and practitioners witness a hitherto unseen opportunity of gaining valuable insights from rich data. Given the increasing availability of individual-level data – for example, through computing in everyday life (Yoo 2010) –, analysts can unveil complex and fine-grained relationships between entities of interest. Big data is associated with two properties that are deemed to carry the greatest potential benefits for information systems (IS) research and practice, namely volume and variety (Abbasi et al. 2016).

The opportunities arising from the new kinds of data call for adequate methods of ML-based analysis tailored to big data (Chen et al. 2012). With the emergence of big data tools and the growing number of sophisticated ML models that are able to learn highly complex relationships from data, both research and practice may benefit from increasingly intelligent methods of generating insights from data (Shmueli and Koppius 2011). We have already witnessed researchers developing automated data analysis methods and tools to support decision making processes on both operational and strategic levels in diverse industries (Kruse et al. 2016: Lucas Jr et al. 2013: Newell and Marabelli 2015: Woerner and Wixom 2015). In practice. however, due to managerial barriers based on a lack of understanding of how to effectively use analytics, the adoption of data-driven decision making methods progresses slowly, despite many organizations having large amounts of rich and high-guality data at their disposal (LaValle et al. 2011). This lack of understanding is partially owed to the fact that complex ML models often constitute black boxes, in the sense that the process through which they generate their outputs are incomprehensible to outsiders (Abbasi et al. 2016). Regarding the design of ML models, their explanatory power and predictive power are often viewed as mutual antagonists (Shmueli 2010; Shmueli and Koppius 2011): By increasing a model's complexity to achieve higher predictive accuracies, the model forfeits part of its explanatory power.

Given the knowledge base and the research methods of the field, Goes (2014) views IS researchers as being excellently positioned to contribute to big data research. He argues that big data constitutes an interesting field of research for a community traditionally engaged in interdisciplinary research. Similarly, Agarwal et al. (2014) state that the IS community possesses knowledge of data management as well as value creation through data. They suggest that the IS community, therefore, holds a comparative advantage for conducting research on big data. Thus, this dissertation draws on this body of knowledge and on contributions from the ML and artificial intelligence (AI) communities to address the challenges associated with big data analytics and data-driven decision making for complex problems. More specifically a nascent design theory (Gregor and Hevner 2013) is developed based on existing ML and AI theory that guides the design of ML systems towards being able to cope with large volumes of data that flow through the system with a high velocity, as well as exploit these properties along with a potentially high data variety to uncover complex relationships between entities of interest. At the same time, the principles proposed in the design theory enable a better understanding of the results produced by ML systems without reducing their overall complexity, which enables in-depth analyses of given problems. At the end of this dissertation, I review the contributions of the individual studies, describe their combined overall contribution, and discuss limitations and implications for research and practice.

## I.2 Research Agenda

The development of the nascent design theory and its exemplary application to three different research problems are divided into two parts. The first part deals with the development of the theory itself, whereas the second part applies the theory in different contexts in which ML can be used to drive decision making. Figure A:1 sketches the research studies carried out in this dissertation, which I briefly describe in the following paragraphs, along with their roles in the

Q

R

overall research design and their relationships among each other. Chapter A.II details the theoretical background to the relevant topics covered in the remainder of this dissertation.



#### Figure A:1. Overview of the research agenda.

The focus of this research lies on analysis tasks that aim to model non-shallow relationships among data. Existing research has extensively dealt with creating ML models and fine-tuning them to different tasks (e.g., decision trees (Breiman et al. 1984), support vector machines (Cortes and Vapnik 1995), and neural networks (Cybenko 1989)). However, employing ML in practice typically goes beyond applying a single ML model to a given task. Instead, an ML system often constitutes a composition of multiple ML models and data processing modules. While we increasingly witness research contributions in the form of level one design science artifacts (Gregor and Hevner 2013) in the IS literature that provide ML systems for particular practical problems, the body of literature on ML lacks formal guidelines regarding the design of such systems.

In the previous section, I described the benefits of ML systems for analytical decision making and the typical challenges they entail. Given the ubiquitousness of IS (Vodanovich et al. 2010), the, thereby, increasing availability of data (Berente and Seidel 2014), and the cruciality of delivering information in a timely manner (Kiron et al. 2012), ML systems carry great potential for supporting decision making processes. However, in order to so, they have to cope with and

take advantage of a high volume, velocity, and variety of data. The given problem leads me to the main research inquiry (RI) of this dissertation, which seeks to answer the following question:

RI 1. How can ML systems be designed to address complex decision making problems while taking advantage of big data and adequately handling its technical and managerial adoption barriers?

In answering this research question, a nascent, level two design theory (Gregor and Hevner 2013) is developed, which I name the Division-of-Labor (DoL) framework. This framework provides design principles for supervised ML problems that deal with data sets of intermediate to large sizes. This design theory is particularly suited for data analysis problems revolving around big data sets with high volumes and varieties (Abbasi et al. 2016). Furthermore, it addresses issues regarding data distributions of complex shapes, high-velocity data, and the interpretability of ML systems. It is applicable in different use cases such as any kinds of forecasts that require consideration of a large number of factors (e.g., user behavior predictions), analyses of large company data that flows into a system in a rapid manner, analyses that require a high level of anonymity and protection of consumer data (e.g., medical analyses), self-organizing systems that not only analyze data but also perform actions autonomously (e.g., autonomous vehicles), and complex, descriptive analyses in research. In order to showcase the design theory's potential, three expository instantiations (Gregor and Jones 2007) are provided. First, a decision support system is designed for making treatment decisions for head and neck cancer patients. This area of application is characterized by the extraordinary cruciality of decisions, which can spell the difference between life and death (Aron et al. 2011). Second, a recommendation system is designed for consumers on two-sided digital platforms that typically involve an intermediate or high number of interactions. Knowing the consumers' preferences to match them with providers is of the highest importance for digitally enabled two-sided platforms (Parker et al. 2016). However, being able to predict consumers' choices is a highly complex problem with platforms on which diverse consumers and products or services are brought together. Third and finally, the efficiency-comfort tradeoff in autonomous vehicles (i.e., the trade-off between optimizing for efficiency and the passenger's individual driving preferences) is formulated and a real-time ML system is designed that addresses this trade-off.

The development of the expository instantiations based on the nascent design theory are, thus, subdivided into the following three research inquiries:

- RI 2. Defining the requirements for designing an ML-based decision support system for head and neck cancer treatments; designing the system according to the DoL framework to enhance the decision making process.
- RI 3. Determining the specific characteristics of consumer data generated on two-sided platforms; designing an ML system according to the DoL framework that maintains highly individual user profiles to models their preferences in detail.
- RI 4. Conceptualizing an ML system based on the DoL framework that models driving preferences of users, learns efficient driving patterns, and resolves the trade-off

Q

between the preferred and an efficient driving behavior; designing the system as a realtime system that can be used in dynamic situations that occur with autonomous vehicles; evaluating the system in a simulation environment based on real carsharing data.

In order to evaluate the proposed artifacts, they are implemented as computer programs using the programming languages *Java*, *Python 2.7*, and *Python 3.6* in combination with data processing and ML libraries such as *NumPy*, *SciPy*, *Scikit-learn*, and *TensorFlow*. The implementations comprise the constructs of the design theory and abstract definitions of their functions as well as concrete formulations of the processes involved in the expository instantiations. The implementations correspond to the formulae provided in this dissertation, which allow for a precise replication of the findings.

## I.3 Structure of the Dissertation

This cumulative dissertation is comprised of three parts. In Part A, I describe the motivation for my research (A.I.1), present the research agenda (AI.2), outline the structure of the dissertation (AI.3), define the research context and design (AI.4), and state the contributions of the dissertation (AI.5). Chapter A.II is concerned with the theoretical background of this dissertation comprising (artificial) intelligence in general, ML and big data analytics, and computational data analytics in IS research.

Part B constitutes the main part of this dissertation. In that part, four studies are carried out that address the design of ML systems for complex decision making problems (see Table A-1). Chapter B.I comprises one research article, in which a nascent design theory for ML systems is derived. In the three articles comprised by Chapter B.II, the design theory is applied to design applications for specific problems in different contexts.

| No | Outlet   | Status    | Ranking<br>(VHB) | Section | RI   | Main contribution   |
|----|--|-----------|------------------|---------|------|---|
| 1  | Proceedings of the<br>Twelfth International<br>Conference on<br>Design Science<br>Research in<br>Information Systems<br>and Technology | Published | С                | B.I     | 1    | Nascent design theory (DoL framework) for machine<br>learning systems dealing with complex analysis<br>problems. The theory addresses real-time (incremental)<br>machine learning and the benefits of computing<br>machine learning models in a distributed manner.   |
| 2  | Proceedings of the<br>Thirty-Eighth<br>International<br>Conference on<br>Information Systems   | Published | A                | B.II    | 1, 2 | Medical decision support system that uses machine<br>learning to provide physicians with treatment<br>recommendations for head and neck cancer patients.<br>The system adheres to the principles of the DoL<br>framework and lays the foundation for deriving design<br>principles for machine-learning-based decision support<br>systems for chronic diseases in general.  |
| 3  | Proceedings of the<br>Fortieth International<br>Conference on<br>Information Systems   | Accepted  | A                | B.II    | 1, 3 | Recommendation system for consumers on two-sided<br>platforms with the purpose of recommending the most<br>desirable content/products for each user. It adheres to<br>the DoL framework and employs a novel technique<br>based on time-series embedding, using a combination<br>of feedforward and recurrent neural networks to create<br>content-based, context-sensitive, and user-specific<br>recommendations. |

Table A-1. Overview of studies included in the thesis.

Electronic and a state of the second s

|  | Journal of Submitted 1,4 arcadually learning a trade-off between resource efficiency, and behavioral preferences. It adheres to the principles of the DoL framework to allow for a novel way of Learning manaer.   4 Management (second A B.II 1,4 gradually learning a trade-off between resource efficiency, and behavioral preferences. It adheres to the principles of the DoL framework to allow for a novel way of learning manaer. The artifact was tested in a simulation environment with a minimal data set obtained from a carsharing provider. |
|--|--|
|--|--|

In Part C, I summarize and synthesize the findings of Part B, discuss limitations and implications for theory and practice, and finally conclude the dissertation and briefly state future research opportunities. The structure of the dissertation is illustrated in Figure A:2.

| A. Foundation   |   |  |   |   |  |  |  |  |
|---|---|--|---|---|--|--|--|--|
| A.I Introduction  |   |  | A.II Theoretical Background   |   |  |  |  |  |
| A.I.1 Motivation  | A.I.2 Research Agenda   |  | A.II.1 Intelligence and<br>Intelligent Agents   | A.II.2 Business<br>Intelligence and Data-<br>Driven Decision Making |  |  |  |  |
| Dissertation  | Dissertation A.I.4 Research Context and Design                    |  |   | d Big Data Analytics  |  |  |  |  |
| A.I.5 Contributions   |   |  |   |   |  |  |  |  |
| B. Developing and Validating a Design Theory for Complex Machine Learning Systems |   |  |   |   |  |  |  |  |
| B.I Distributed Cognitiv<br>Deriving Design Princip<br>Theory                     | e Expert Systems:<br>les from Existing                            |  | B.II Distributed Cognitive Expert Systems in<br>Action: Developing Specific Instantiations of<br>Machine Learning Systems |   |  |  |  |  |
| Study 1: Towards Distributed Cognitive Expert<br>Systems                          |   |  | Study 2: Distributed<br>Cognitive Expert<br>Systems in Cancer Data<br>Analytics   | Study 3: Advancing<br>Recommendations on<br>Two-Sided Platforms     |  |  |  |  |
|   |   |  | Study 4: A Machine Learning Approach to the<br>Efficiency-Comfort Trade-Offin Everyday-Life<br>Automation                 |   |  |  |  |  |
| C. Contributions  |   |  |   |   |  |  |  |  |
| C.I Findings and Results  |   |  | C.II Limitations and Implications   |   |  |  |  |  |
| C.I.1-3 Reflections on the<br>Development of the<br>Individual Artifacts          | C.I.4 Lessons learned<br>and Revisitation of the<br>Design Theory |  | C.II.1 Limitations  | C.II.2 Implications for<br>Practice                                 |  |  |  |  |
|   |   |  | C.II.3 Implications for Research  |   |  |  |  |  |
| C.III Conclusion and Outlook  |   |  |   |   |  |  |  |  |

Figure A:2. Structure of the dissertation.

# I.4 Research Context and Design

As stated in the Research Agenda Section, the studies are carried out in different settings. Study 1 is a purely theoretical approach to designing ML systems. Drawing on existing theory on AI, mathematical optimization, and basic economics, a nascent design theory is derived. In order to provide the theory's practicability (Gregor and Hevner 2013), expository instantiations (Gregor and Jones 2007) are provided in three subsequent studies. Study 2 considers data on cancer patients collected by the Department for Oral and Maxillofacial Surgery of the University Medical Center Göttingen. These data and data collected in medical environments in general are characterized by a large number of missing values (Wells et al. 2013). Treatment decisions embody great impacts for both patients and physicians and the large number of factors that determine the outcome of a treatment makes it exceptionally difficult for physicians to make the right decision. Furthermore, they are hesitant to adopt IS partially due to the additional administrative efforts that would arise thereafter (Romanow et al. 2012). An artifact is designed that considers these circumstances to make a step towards practicable ML for medical decision support systems. Study 3 addresses shortcomings of contemporary recommendation systems, which are well known because they have been studied in regard to a large variety of recommendation techniques (Cano and Morisio 2017). A novel user-specific, context-aware, content-based, and collaborative recommendation system is developed that may effectively process static and time-varying variables of users and take into account their history and similarities to other users in order to make more precise predictions regarding their future actions on two-sided, digital platforms. The proposed system is evaluated using data gathered on Expedia's hotel booking platform. Finally, Study 4 designs an artifact for a real-time environment, namely autonomous vehicles. More specifically, an ML system is developed composed of three ML modules: a module for modeling users' preferred driving behaviors, a module for modeling the relationship between driving behavior and efficiency, and a module that builds upon these two modules to resolve a postulated trade-off between efficient and comfortable driving. The proposed design may be adapted for other situations in which individuals' goals, especially regarding processes that play a role in people's everyday lives (Yoo 2010), may run against system-level goals.



Figure A:3. Research design.

Studies 2, 3, and 4 not only provide new ways of addressing ML problems in different application areas. Combined with Study 1, they also become part of a nascent design theory that is put to work in different contexts to showcase its effectiveness in addressing different

challenges when employing ML and big data analytics. The overall design of the research is depicted in Figure A:3.

In the field of IS, we may distinguish five general streams of research, which I describe in the following according to Banker and Kauffman (2004). The research stream value of information is concerned with understanding the value that information may contribute. It determines the worth of certain information for individual decision makers and firms in a market context. The stream of *human-computer systems design* studies cognitive processes based on behavioral decision theory to model individual and group dynamics when interacting with IS. IS organization and strategy is a research stream that considers the value of IS investments on levels of analysis that go beyond the individual level, for example, the system level, business process level, strategic level, and organizational level. The research stream economics of IS and IT deals with the value of IT, the adoption of network technologies, and IT-based coordination for organizations, markets, and industries. Finally, the decision support and design science research stream - a science of the artificial (Gregor 2009; Simon 1996) - deals with designing and implementing IS. This dissertation, taking action with respect to complex decision making problems by designing ML systems, contributes to the latter research stream. In contrast to behavioral science - a traditional paradigm in IS research -, which is rooted in natural science and aims to explain and predict social phenomena (Hevner et al. 2004: Shmueli and Koppius 2011), design science research is concerned "not with how things are but with how they might be" (Simon 1996, p. xii). Research conducted based on the design science paradigm, thus, produces prescriptive knowledge, which describes the steps to achieve a certain goal (Gregor and Jones 2007). In IS research, design science contributions build IS. which gather, process, and provide information and, thus, often take on the shape of decision support systems (Arnott and Pervan 2012), which are also a point of focus of this dissertation in dealing with decision support and autonomous decision making systems. Here, the design of these systems follows a top-down approach by first developing a design theory, which represents a meta artifact with the purpose to address a class of problems (livari 2015), and then applying this design theory to specific instances of the problem class.

From a philosophical perspective, inquiries in IS research – and social science research in general – each follow one of two main epistemologies, namely positivism and interpretivism (Gregor 2006). "Positivism is largely concerned with the testing, confirmation and falsification, and predictive ability of generalizable theories" (Wynn and Williams 2012, p. 788) and assumes the existence of an objective reality, whereas interpretivism considers reality to be a subjective view of each individual on a given phenomenon in a specific context (Wynn and Williams 2012). According to Gregor (2006), design theories require some kind of realist ontology. In social science research, we witness the increasingly important role of a realist epistemological view, namely critical realism (Bhaskar 1998), which states that there exists an independent reality, albeit not fully understandable or observable by humans, and requires research to examine the shape of components and interactions in this reality based on given observations (Wynn and Williams 2012). This dissertation is concerned with developing methods to automatically derive insights from data. For this purpose, I follow the ontology of critical realism, assuming that given observations are indicators that an observer may use to model relationships that

9

Q

they believe to exist in a hidden reality. Furthermore, with human knowledge of reality being fallible (Wynn and Williams 2012) and assuming that this statement holds true for any observer that obeys the laws of this reality (i.e., both biological and artificial agents (Pfeifer and Bongard 2006)), the proposed design theory advocates for adaptive systems that not only create representations of their perceived realities but also adapt these representations when new observations are made.

Table A-2 summarizes the research design of this work.

| No. | RI   | Epistemology     | Paradigm          | Methodology<br>(Seminal work)  | Data collection  | Data analysis   |
|-----|------|------------------|-------------------|--|--|---|
| 1   | 1    | Critical Realism | Design<br>Science | Top-down development of a<br>design theory (Gregor and<br>Hevner 2013)         | -/-  | -/-   |
| 2   | 1, 2 |                  |                   | Design according to the DoL<br>framework (Tofangchi,<br>Hanelt and Kolbe 2017) | Collection of structured<br>and unstructured patient<br>data by the medical staff<br>of the University Medical<br>Center Göttingen | Machine learning<br>(clustering, different<br>kinds of classification<br>models)<br>Machine learning<br>(dimensionality<br>reduction, different<br>kinds of classification<br>models) |
| 3   | 1, 3 |                  |                   | Design according to the DoL<br>framework (Tofangchi,<br>Hanelt and Kolbe 2017) | Quantitative consumer<br>data automatically<br>collected through the<br>Expedia platform   |   |
| 4   | 1, 4 |                  |                   | Design according to the DoL<br>framework (Tofangchi,<br>Hanelt and Kolbe 2017) | Real-time usage data<br>automatically collected<br>in carsharing vehicles  | Machine learning<br>(different kinds of<br>regression models,<br>density estimation)  |

Table A-2. Overview of research design and core research questions.

# I.5 Contributions

The major contribution of this dissertation is a nascent design theory (Gregor and Hevner 2013) that builds upon existing theory (i.e., "kernel theories" (Hevner et al. 2004, p. 76)) in artificial intelligence, mathematical optimization, and economics to provide design principles for complex ML systems. It can be classified as a theory for design and action (Type V according to Gregor (2006)) and comprises prescriptive knowledge for a given class of problems. As a theory involving constructs, methods, models, and design principles to address a certain class of problems, it constitutes a level 2 (nascent) design theory (Gregor and Hevner 2013). The design theory deals with how to address complex decision making problems (diverse data sets and complex relationships to uncover) through ML based on a composition of models while maintaining the interpretability of the designed systems, the durability of employed models (i.e., efficient incorporation of new data), and their adaptability to environmental changes that may alter previously identified relationships. Although ML research has contributed individual methods for addressing some of the above-mentioned challenges (e.g., ML based on composing multiple models through bagging (Breiman 2001) or boosting (Schapire 1990)), there exists, to the best of my knowledge, no extensive theory for the design of ML systems that addresses the technical challenges that occur in business contexts that involve big data. From the perspective of AI research, this dissertation contributes to the research stream of collective intelligence (al-Rifaie et al. 2012; Beni and Wang 1993; Lévy 1997; Reynolds 1987; Rosenberg and Pescetelli 2017): It shows how a complex autonomous system can be constructed by using simple ML models as buildings blocks that compose a recursive collective

10

system. The system may exhibit diverse behavior that is determined by the behavior of its building blocks, which are agents that act autonomously towards a common goal. With respect to big data analytics in IS research, design principles are contributed that not only allow addressing the challenges of big data but also leverage the high volume and variety of big data to build systems that process tasks of a high complexity. For these reasons, I classify this contribution as an "invention" (Gregor and Hevner 2013, p. 345). That is, the design theory deals with an application domain with a low maturity (ML and big data research in general is still at an early stage, with many uncertainties regarding the suitability of different ML models for different kinds of problems) and a low solution maturity (the given problem has not been adequately addressed yet).

Aside from the main contribution, this work includes several minor contributions: The provided expository instantiations (Gregor 2007) for three different application contexts, which require building models to get an understanding of the studied problems, contribute to theory for explaining (Type II) and theory for predicting (Type III). More specifically, this work develops (1) data science methods for head and neck cancer treatments that identify patient attributes that are most crucial to making treatment decisions and predict optimal treatments, (2) a recommendation system for the Expedia platform that makes hotel recommendations based on users' historical behavior, their similarity to other users, and uses a novel approach to model user contexts, and (3) a real-time system for autonomous vehicles that uses data-driven techniques to resolve a user-specific trade-off between efficient and comfortable driving. These situated artifacts - brought into material existence by applying the nascent design theory - can be, in themselves, understood as contributions to design science research (Gregor and Hevner 2013). Each of these artifacts solves a particular problem and may be the first step in the emergence of new theory that addresses a broader spectrum of related phenomena (Gregor and Hevner 2013). Moreover, the research framework surrounding these artifacts offers more abstract method descriptions and design principles that can be operationalized in further contexts. These contextual studies (1) describe how, in typical medical environments in which physicians decide upon treatments and log their decisions, patient attributes crucial to treatment decisions can be discovered in a data-driven manner and which design principles for treatment decision support systems can be generalized to chronic diseases, (2) show that a recommendation system for two-sided platforms, in general, learns more effectively by taking into account collaborative, content-based, contextual, and user-specific variables at the same time and how to integrate static and time-varying data, and (3) provide design principles for a trade-off module for autonomous agents whose interactions with humans is characterized by conflicting system-level goals (e.g., efficiency) and individual-level goals (e.g., comfort of users). On a more abstract level, these artifacts are expository instantiations (Gregor and Jones 2007) of the design theory and therein contribute to our understanding of it. Being in a nascent state, this theory requires practice-oriented research to test its propositions and drive its maturity.

Finally, the proposed design theory also constitutes a contribution to business practice, dealing with increasingly significant issues related to employing big data systems. While big data enables more sophisticated analyses of business problems, their accommodation requires

Q