

SCHLAPPA, MARTIN

*OPTIMIZING PRODUCTION PROCESSES VIA IMPROVED RESILIENCE
AND STATE-OF-THE-ART AI TECHNOLOGIES*

Dissertation

for obtaining the degree of Doctor of Business and Economics

(Doctor rerum politicarum – Dr. rer. pol.)

at WHU – Otto Beisheim School of Management

June 19, 2023

First Advisor: Professor Stefan Spinler

Second Advisor: Professor Liji Shen

Martin Schlappa: *Optimizing Production Processes via Improved Resilience and State-of-the-art AI Technologies*

© June 19, 2023

To my wife, my parents, and my sister.

ABSTRACT

Disruptions regularly hit economies. Scholars and industry experts suggest many strategies to avoid disruptions or handle them effectively. Two things are repeatedly mentioned: Increasing resilience and deploying artificial intelligence (AI) technologies. In this dissertation, we¹ look at both aspects and focus our efforts on production processes. To get a well-rounded view, we apply various research methods, i.e., surveys, case studies, and systematic literature reviews (SLRs).

In our first paper (Chapter 2), we investigate the perceived organizational resilience of companies in the German manufacturing industry. We perform an SLR to analyze existing research on organizational resilience measures. We see that existing (qualitative) resilience measures are complex, challenging to interpret, and therefore, hard to scale and apply across multiple industries. Based on this, we develop a novel, low-threshold resilience measure consisting of six resilience items about the past perceived internal/external resilience, current perceived internal/external resilience, and anticipated need for internal/external resilience, called the Enterprise Resilience Index (ERI). Finally, we conduct an empirical study with ~200 German experts across various industries. Our survey shows that the German manufacturing industry perceives itself as relatively resilient, with significant differences between industries and company sizes. We also see that they anticipate a high need for external resilience across industries in the future. Most strikingly, the Machinery industry shows the lowest ERI levels while it anticipates a relatively high need for resilience, showing the development need for this industry in terms of resilience.

To explore the aspect of AI, we focus on waste incineration plants (WIP) in Chapter 3. WIPs have various levels of automation, but they still rely on manual operations by human operators. Consequently, the combustion process is managed rather inefficiently, and steam outputs and emission levels are not optimal. Thus, we investigate how reinforcement learning (RL) can help enhance process automation and thus optimize the combustion process, e.g., by making more frequent and diverse interventions. An RL agent is trained via trial and error with a reward function that includes the optimization criteria. Since the actual equipment, i.e., the real WIP, cannot be used as the training environment, a digital twin is built using original plant data and a neural network. The RL agent is then trained in this offline environment with the deep Q-network algorithm (DQN). Our work demonstrates that a digital twin of a WIP can be built in a data-driven way. We show that the RL agent outperforms the human operator, increasing the steam output by 7.4% and reducing the oxygen level by 3.6%. Thus, applying RL might benefit the plant operator financially due to increased output and the environment in terms of reduced emission levels.

Finally, we look at a practical aspect of AI: AI readiness and adoption (Chapter 4). Many companies across various sectors have adopted AI technologies. However, the supposedly high adoption rates are misleading since many applications are rather experimental and not applied in key business areas. We believe that this limited AI adoption arises from a lack of AI

¹ The term “we” refers to the authors of the respective chapters, as noted at the beginning of each chapter.

readiness. We conduct a case study in the waste incineration industry with over 160 clients and investigate which strategies facilitate AI adoption in not-AI-ready industries. Based on these interactions, we distill five strategies that counter typical AI readiness barriers, thus increasing AI readiness: education, trust, customer centricity, focus, and collaboration. These strategies focus on transforming businesses just as much as necessary to prepare them for the AI technology that is supposed to be implemented. With increased AI readiness, chances for AI adoption rise. We are convinced that these strategies can be applied in various environments.

In summary, this dissertation gives empirical evidence and expands the literature on organizational resilience and benchmarking, reinforcement learning and digital twins, and AI readiness and adoption.

ACKNOWLEDGMENTS

This dissertation was prepared at the Chair of Logistics Management at WHU – Otto Beisheim School of Management between 2019 and 2022.

First, I want to thank Prof. Dr. Stefan Spinler for his support of this dissertation. He connected me with my research partners and guided me through the process, from the initial research proposal to the final write-up. I am glad that he gave me the freedom to try something new that goes beyond the typical research projects at his chair. I highly appreciate his detail orientation, encouragement to managerial implications, and trust in my abilities.

I also want to thank Prof. Dr. Liji Shen for supporting my dissertation as a second advisor and offering valuable feedback during my research.

Moreover, I would like to my two research partners:

The SPAICER project (“Scalable adaptive production systems through AI-based resilience optimization”) funded by the Federal Ministry for Economic Affairs and Climate Action (BMWK) (funding reference 01MK20015A). I thank all SPAICER partners for the insightful and valuable discussions around AI, resilience, supply chains, and the “Mittelstand”. I want to highlight the support of my co-authors Jan Sporkmann, Martin Unterberg, and Sebastian Bouschery.

Uniper Technologies, especially Jonas Hegemann and Tobias Mathur. Thank you for giving me access to valuable data, enabling my empirical work, and supporting me as co-authors in my research endeavors.

Last but not least, I want to thank the numerous people who supported me, challenged and read my work, listened to my ideas, worked through my code, and so on.

Thank you.

CONTENTS

- 1 Introduction..... 1
 - 1.1 Disruptions and resilience 1
 - 1.2 Artificial intelligence and reinforcement learning 1
 - 1.3 Contributions of this dissertation 2
- 2 Benchmarking Organizational Resilience in the German Manufacturing Industry 4
 - 2.1 Introduction 4
 - 2.2 Literature review 5
 - 2.2.1 Question formulation 6
 - 2.2.2 Locating studies 6
 - 2.2.3 Study selection and evaluation 6
 - 2.2.4 Analysis and synthesis 8
 - 2.2.5 Reporting and using the results..... 13
 - 2.3 Methodology 13
 - 2.3.1 Survey design..... 14
 - 2.3.2 Enterprise Resilience Index 15
 - 2.3.3 Data collection 16
 - 2.4 Results 17
 - 2.4.1 Data pre-processing 17
 - 2.4.2 Descriptive statistics 17
 - 2.4.3 Data analysis 20
 - 2.5 Discussion 25
 - 2.6 Conclusion..... 26
- 3 Optimizing Control of Waste Incineration Plants Using Reinforcement Learning and Digital Twins..... 28
 - 3.1 Introduction 28
 - 3.2 Literature review 29
 - 3.2.1 Reinforcement learning in thermal power plants..... 29
 - 3.2.2 Reinforcement learning in waste incineration plants..... 31
 - 3.3 Methodology 31
 - 3.3.1 Data 32
 - 3.3.2 Digital twin of the WIP..... 35
 - 3.3.3 Implementation of DQN 38
 - 3.4 Results and discussion..... 40

3.5	Conclusion.....	43
4	Overcoming Barriers to AI Adoption: A Case Study from the Waste Incineration Industry	44
4.1	Introduction	44
4.2	Literature review	45
4.2.1	Organizational AI readiness and adoption.....	45
4.2.2	Case studies on AI readiness and adoption.....	47
4.3	Methodology	51
4.3.1	Getting started.....	51
4.3.2	Selecting cases	51
4.3.3	Crafting instruments and protocols.....	52
4.3.4	Entering the field & Analyzing data.....	53
4.3.5	Shaping hypotheses & Enfolding literature.....	54
4.3.6	Reaching closure.....	54
4.4	Results	54
4.5	Discussion	58
4.6	Conclusion.....	59
5	Conclusion	61
5.1	Summary	61
5.2	Outlook.....	63
A	Appendix to Chapter 2	65
B	Appendix to Chapter 3.....	67
	References	69

LIST OF FIGURES

Fig. 2.1:	Five-step systematic literature review	6
Fig. 2.2:	Inclusion/exclusion funnel	7
Fig. 2.3:	Samples sizes of all 18 relevant papers	12
Fig. 2.4:	Dimensions of resilience	14
Fig. 2.5:	Boxplot visualization of all responses	19
Fig. 2.6:	Distribution of responses to items Q1, Q2, Q3, and Q4	20
Fig. 2.7:	Distribution of responses to items Q5 and Q6	20
Fig. 2.8:	Visualization of the ERI per industry	24
Fig. 3.1:	Operators and environments	31
Fig. 3.2:	Overarching design of the implementation approach	32
Fig. 3.3:	Typical behavior of the human operator	34
Fig. 3.4:	Analysis of performance variables: (a) steam output and (b) oxygen level.....	35
Fig. 3.5:	Inputs for and outputs of the digital twin	36
Fig. 3.6:	Training and test loss for the best performing settings	37
Fig. 3.7:	Comparison of real and simulated oxygen and steam levels	37
Fig. 3.8:	Training process using DQN and the digital twin.....	39
Fig. 3.9:	Comparison of the original behavior of the human operator and the suggested behavior by the RL agent for 1 episode and 1 control lever	42
Fig. 4.1:	Main challenges for AI readiness and adoption.....	50
Fig. 4.2:	Process for case study research	51
Fig. 4.3:	Strategies to increase organizational AI readiness factors.....	56
Fig. 4.4:	Staged implementation of AI solution	57

LIST OF TABLES

Table 2.1: Categorization of relevant articles	10
Table 2.2: Overview of purpose and research method.....	12
Table 2.3: Survey participants by industry cluster, revenue, and size	18
Table 2.4: Descriptive analysis of all responses	19
Table 2.5: Spearman’s rank correlation matrix	21
Table 2.6: Mann-Whitney-U test for various hypotheses	22
Table 2.7: Means of ERI and NFR per industry cluster.....	23
Table 2.8: Item-wise results for the Mann-Whitney-U tests for SMEs vs. LEs, items Q1, Q2, Q3, and Q4	24
Table 3.1: Analysis of control variables.....	33
Table 3.2: Mean episode results after training a selected number of episodes	40
Table 3.3: Mean episode rewards.....	40
Table 3.4: Mean oxygen and steam levels	41
Table 4.1: Overview of articles assessing organizational AI readiness and adoption empirically.....	48
Table 4.2: Split of customers by plant type and country.....	53
Table 4.3: Assessment of AI readiness in WIPs using organizational AI readiness factors....	55

LIST OF ABBREVIATIONS AND ACRONYMS

A3C	Asynchronous advantage actor-critic
ADHDP	Action dependent heuristic dynamic programming
AI	Artificial intelligence
B2B	Business-to-business
BCI	Business climate index
BMWK	Federal Ministry for Economic Affairs and Climate Action
CFB	Circulating fluidized bed
DCS	Distributed control system
DDPG	Deep deterministic policy gradient
DOI	Diffusion of innovation theory
DQN	Deep Q-network
ERI	Enterprise Resilience Index
ERP	Enterprise resource planning
GDP	Gross domestic product
LE	Large enterprise
LSTM	Long short-term memory
MORE	Model-based offline RL with restrictive exploration
MPC	Model predictive control
MSE	Mean-squared error
MSME	Micro, small, and medium enterprise
NFQ	Neural fitted Q-iteration
NFR	Anticipated need for resilience
OEM	Original equipment manufacturer
PCA	Principal component analysis
PID	Proportional–integral–derivative
RL	Reinforcement learning
RNN	Recurrent neural networks
SD	Standard deviation
SLR	Systematic literature review
SME	Small and medium-sized enterprise
SPAICER	Scalable adaptive production systems through AI-based resilience optimization
SUMO	Simulation of urban mobility
TAM	Technology acceptance model
TOE	Technology-organization-environment framework
UTAUT	Unified theory of acceptance and use of technology
VAM	Vinyl acetate monomer
VISTA	Virtual image synthesis and transformation for autonomy
WIP	Waste incineration plant

1 INTRODUCTION

Disruptions regularly hit economies. Some are known and happen frequently; others are unknown and hit seemingly out of nowhere. A very recent and devastating example is the outbreak of the war in Ukraine, which has created massive challenges across Europe. For instance, the energy market in Germany has come under tremendous pressure. Gasoline prices are at all-time highs, and previously unquestioned and stable gas flows are suddenly uncertain.

Examples like this show that it is vital to investigate strategies to avoid disruptions or handle them effectively. Literature and experts suggest many strategies, but two ideas are repeatedly mentioned: Increasing *resilience* and deploying *artificial intelligence (AI) technologies*.

1.1 Disruptions and resilience

“Resilience” is one of those words that is frequently used when bad things happen. People, companies, supply chains, environments, and many others can be resilient (or not). Resilience is a “multidisciplinary concept” (Ponis & Koronis, 2012, p. 921), existing in many realms, including ecology, environmental studies, psychology, engineering, and management (Hosseini et al., 2016; Kamalahmadi & Parast, 2016). In the management world, many definitions exist with slight nuances that distinguish them. For instance, Tukamuhabwa et al. describe resilience as “the adaptive capability [...] to prepare for and/or respond to disruptions, to make a timely and cost effective recovery, and therefore progress to a post-disruption state of operations – ideally, a better state than prior to the disruption” (2015, p. 5599).

The definition highlights the essential aspect of resilience: disruptions. Without disruptions, resilience cannot be experienced. Fundamentally, resilience is grounded on the idea that some disruptions are inevitable, whether predictable or unpredictable (Fiksel et al., 2015; Hohenstein et al., 2015; Pires Ribeiro & Barbosa-Povoa, 2018). The causes of such disruptions can be diverse: economic crises, humanitarian crises, medical emergencies, natural disasters, political conflicts, or wars (Henry & Ramirez-Marquez, 2016; Hohenstein et al., 2015; Jüttner & Maklan, 2011). The ability to manage disruptions can be an essential success factor for companies (Christopher & Lee, 2004).

Following a “typical” disruption profile (Sheffi & Rice, 2005), resilience management involves four phases: *preparation*, *response*, *recovery*, and *growth*. Within each phase, there are various strategies to increase a company’s resilience, for instance, flexibility, collaboration, redundancies, visibility, or velocity (Jüttner & Maklan, 2011; Kamalahmadi & Parast, 2016; Pettit et al., 2013; Ponomarov & Holcomb, 2009; Singh et al., 2019). Some of these strategies are facilitated by AI technologies (Bauer et al., 2021; Hosseini & Ivanov, 2020; Öksüz et al., 2021; Spieske & Birkel, 2021).

1.2 Artificial intelligence and reinforcement learning

Similar to resilience, “AI” is another “hot topic”. “97.2% of executives said their firms were investing in Big Data and AI initiatives” (NewVantage Partners, 2019, p. 4). AI technologies can be used to optimize operations, enhance customer satisfaction, enable

personalization or increase productivity (Chui et al., 2018; Grover et al., 2022). More data, more processing power, and better algorithms enable AI technologies (Bughin et al., 2017).

Many technologies can be described as AI, and AI itself can be distinguished in various ways (Benbya et al., 2020; Davenport, 2021). Machine learning is a key area of AI, which can be split into supervised, unsupervised, and reinforcement learning (RL). Whereas supervised and unsupervised learning have been widely applied, applications of RL are relatively rare. The idea behind RL is to train a machine the way humans learn: by trial and error (Sutton & Barto, 2018). Instead of complicated if-then-relations, in RL, a reward function that rewards “good behavior” is defined. An agent can execute one of the pre-defined actions in a digital environment and learn from the feedback (i.e., rewards) it receives. In the last decade, breakthroughs in RL occurred, such as AlphaGo’s win in the board game Go against the best human player in 2016 (Silver et al., 2016). Given the complexity of this approach and the need for an environment, most use cases are implemented in games, primarily by big tech firms.

However, this approach is also well suited for many control problems, which tend to occur in more traditional, engineering-driven companies. They typically have control problems that are hard to automate using conventional methods, thus requiring human assistance. Humans, however, are limited in processing large amounts of data and performing at continuously high levels, thus causing disruptions. AI can help overcome the issues of conventional automation methods as well as human operations. With the support of AI, processes can become not only less prone to disruptions but also optimized in terms of various performance characteristics, for instance, by identifying patterns or anomalies in processes. This way, processes can become more resilient. At the same time, human capacity can be freed for more meaningful tasks such as regular management activities, preparation for disruptions, or creative work unsuitable for AI (Grover et al., 2022).

Even though the advantages of AI are evident, full-scale applications are rarely seen in the real world. It seems that AI is adopted a lot, but it is usually only experimental, and use cases fail to scale (Bughin et al., 2017; Holmström, 2021). Many AI projects fail altogether (Brock & von Wangenheim, 2019). Companies try to adopt AI even though they are not ready for AI (yet). Unfortunately, organizations need to be ready for AI before adopting AI (AlSheibani et al., 2018; Issa et al., 2021; Jöhnk et al., 2021; Pumplun et al., 2019).

1.3 Contributions of this dissertation

This dissertation is cumulative and based on three research projects. In the following three chapters, we² present each project separately:

In Chapter 2, we study the perceived organizational resilience of the German manufacturing industry. We survey ~200 professionals to understand better their company’s past, current, and future organizational resilience. Moreover, we develop the Enterprise Resilience Index (ERI), a novel, low-threshold, survey-based measure for perceived organizational resilience. The survey results suggest a medium resilience level across the manufacturing industry, with differences between sub-industries: For instance, participants from the Chemicals industry perceive their companies as relatively more resilient whereas

² The term “we” refers to the authors of the respective chapters, as noted at the beginning of each chapter.

participants from the Machinery industry think the opposite. We also see that current resilience levels are below past (i.e., last year's) levels, confirming the value-add of continuously measuring resilience as a dynamic capability. We also gather a strong anticipated need for external resilience, arguably because the German manufacturing industry expects disruption from external sources. This need is especially pronounced in large enterprises (LEs) with significantly higher levels than small and medium-sized enterprises (SMEs). In the future, our survey can regularly benchmark various industries in terms of perceived organizational resilience.

In Chapter 3, we use RL to control and optimize the combustion process of a waste incineration plant (WIP). Since the actual equipment, i.e., the real WIP, cannot be used as the training environment, a digital twin is built using original plant data and a neural network. The RL agent is then trained in this offline environment with the deep Q-network (DQN) algorithm via trial and error with a reward function that includes multiple optimization criteria. We show that a digital twin of a WIP can be built in a data-driven way with limited expert knowledge. The RL agent outperforms the human operator, increasing the steam output by 7.4% (converted to electricity and heat) and reducing the oxygen level by 3.6%. Thus, the application of RL benefits the plant operator in financial terms due to increased output and the environment in terms of reduced emission levels. This work helps bridge the gap between theoretical research on control problems and real-life industrial issues. In the future, this research can be expanded to other optimization targets, such as total costs or the pollution of the boiler. In a broader sense, this application can be seen as a general-purpose framework, showcasing the power of RL to control complex processes without the need for very specialized models of the process or expert knowledge.

In Chapter 4, we look at AI readiness and adoption. Using the case study method, we investigate which strategies facilitate AI adoption in traditional, i.e., not AI-ready industries. We provide empirical evidence to a topic that is usually studied conceptually. More specifically, we examine Uniper's efforts in implementing their AI solution "Operait" across the waste incineration industry and assess interactions and discussions with over 160 clients over two years. Based on these interactions, we distill five strategies that counter typical AI readiness barriers, thus increasing AI readiness: education, trust, customer centricity, focus, and collaboration. These strategies focus on transforming the businesses just as much as necessary to prepare them for the AI technology that is supposed to be implemented. With increased AI readiness, chances for AI adoption rise. We are convinced that these strategies can be applied in various environments.

In the final chapter, we synthesize our three research projects, highlight the managerial implications of our work, and lay out opportunities for future research.

2 BENCHMARKING ORGANIZATIONAL RESILIENCE IN THE GERMAN MANUFACTURING INDUSTRY

This chapter is based on Schlappa, Sporkmann, et al. (2022).³

2.1 Introduction

Companies in the German manufacturing industry are subject to disruptions, which can negatively affect their performance, such as productivity or sustainability (Tang & Nurmaya Musa, 2011). For example, the recent COVID-19 pandemic and the ongoing war in Ukraine have caused significant (unanticipated) disruptions in the industry (Queiroz et al., 2022). The ability to manage disruptions can be defined as organizational resilience – a research area with substantial advancements in recent years. Organizational resilience is a pivotal capability that distinguishes a successful company from others (Christopher & Lee, 2004) and yields competitive advantages (Norrman & Jansson, 2004; Sheffi & Rice, 2005).

In a nutshell, organizational resilience describes the level to which companies can “keep the lights on”. After getting hit by a disruption, a resilient company recovers in a timely and cost-effective fashion (Christopher & Peck, 2004; Ponomarov & Holcomb, 2009). As a very “multidisciplinary concept” (Hosseini, Morshedlou, et al., 2019, p. 125), organizational resilience builds on the idea that some disruptions cannot be avoided ex-ante (Hohenstein et al., 2015; Pires Ribeiro & Barbosa-Povoa, 2018), and it acknowledges that some are even “unpredictable or unknowable before the fact” (Fiksel et al., 2015, p. 81).

Organizational resilience needs to be managed actively in a coordinated management approach, which we call resilience management. Managers should act within the individual “resilience fitness space” (Pettit et al., 2013, p. 47). They cannot implement all possible resilience measures at once, given that (financial) resources are naturally limited. Instead, they should analyze the resilience measures relative to the expected disruptions and associated benefits and implement the most efficient ones. To do this, managers require reliable information on the current resilience level. However, we find no consensus on how organizational resilience can be measured.

Various concepts for measuring organizational resilience have been published, linking organizational resilience to existing performance indicators within a company. In the SPAICER project⁴, we tried deploying such a concept with our industry partners. However, we have learned firsthand that existing resilience measures are complex, difficult to deploy, challenging to interpret, and therefore, hard to scale and apply across multiple industries. If anything, these resilience measures are context-specific and tailored only to the individual company.

³ This unpublished working paper with the title “Benchmarking Organizational Resilience in the German Manufacturing Industry” was written by Martin Schlappa, Jan Sporkmann, Martin Unterberg, and Sebastian Bouschery and has been submitted to “Production Engineering”. It is currently being revised and resubmitted.

⁴ The SPAICER project focuses on AI-based resilience optimization (<https://www.spaicer.de/en/>). Partners from academia and industry participate in this applied research.

Based on our experience, we developed a novel, low-threshold resilience measure, a survey-based index for (perceived) resilience levels, targeted at the German manufacturing industry. We transfer the well-established concept of the ifo Business Climate Index (BCI) to organizational resilience and query companies about (1) their past perceived internal/external resilience, (2) their past perceived internal/external resilience, and (3) their anticipated need for internal/external resilience. Moreover, we develop the Enterprise Resilience Index (ERI), an easy-to-adapt indicator for the organizational resilience level. It is meant to supplement data from company-specific resilience indices and is well-suited for industry-wide benchmarks.

Our paper makes the following contributions: First, we conduct a systematic literature review (SLR) and create transparency about the state-of-the-art qualitative organizational resilience measures. We are the first to provide a broad overview of qualitative organizational resilience measures for future studies to build on. Second, we analyze and benchmark organizational resilience in the German manufacturing industry. Our data analysis is motivated by insights gathered from discussions with our industry partners in the SPAICER project: among others, we analyze statistical correlations between the past and the current organizational resilience, benchmark the organizational resilience of small and medium-sized enterprises (SMEs) and large enterprises (LEs), and investigate industry-related differences. Thus, we increase the understanding of our resilience measure and help practitioners establish a more targeted resilience management in the future. Third, our ERI represents a new resilience measure that is easy to implement at scale, providing practitioners with a tool to benchmark their (perceived) organizational resilience level in cross-industry benchmarks. The ERI is designed to serve in longitudinal analyses in the future.

The remainder of this paper is organized as follows. First, Chapter 2.2 outlines the SLR methodology and presents our literature-related findings. In Chapter 2.3, we describe the survey design, the ERI methodology, and the data collection process in detail. In Chapter 2.4, we analyze the collected data from 210 industry representatives. Afterward, in Chapter 2.5, we discuss the implications derived from our research for both researchers and practitioners. Chapter 2.6 summarizes our research and concludes with an outlook on potential future research.

2.2 Literature review

There is a large body of literature on organizational resilience, especially from the last few years, with multiple articles discussing approaches to measuring resilience (e.g., Hillmann & Guenther, 2021). Regarding the latter, we note that existing (quantitative) approaches (e.g., Hosseini et al., 2016) are quite complex and difficult to interpret. Instead, they require significant efforts from the respondents, support from researchers, and large amounts of data. Thus, the existing (quantitative) approaches are hard to scale and apply across multiple industries.

To provide a solid foundation for our resilience research, we conduct an SLR, as proposed by Denyer and Tranfield (2009) and Tranfield et al. (2003). The SLR ultimately helps us better understand the space of qualitative approaches, i.e., surveys, interviews, or case studies, to measuring organizational resilience. Hillmann and Guenther (2021) highlight that no standard

exists. We follow the five-step approach as depicted in Fig. 2.1 to gather and assess all essential information from existing research (Rousseau et al., 2008).




Fig. 2.1: Five-step systematic literature review (Denyer & Tranfield, 2009)

2.2.1 Question formulation

We intend to measure organizational resilience across many companies in the German manufacturing industry and, ideally, repeatedly over many years. As mentioned previously, we acknowledge that organizational resilience is an abstract attribute, with no consensus on its definition among researchers and practitioners. In this paper, we adopt the resilience definition of Tukamuhabwa et al. (2015, p. 5599): “the adaptive capability of a supply chain to prepare for and/or respond to disruptions, to make a timely and cost-effective recovery, and therefore progress to a post-disruption state of operations – ideally, a better state than prior to the disruption”. The term resilience level describes the concrete expression of resilience. On top, we find that measuring resilience is a complex endeavor. Nonetheless, researchers have tried to measure it before, and thus, we focus our efforts for the SLR on the following question: *What are qualitative organizational resilience measures?*

2.2.2 Locating studies

We use the Scopus database to gather the literature, which is widely used in research (see, e.g., Hosseini, Ivanov, et al., 2019; Kegyes et al., 2021). We collect 3,309 records as of April 2022 with the keywords “resilience index” OR “resilience benchmark” OR “resilience measurement”. We keep the keywords rather generic and purposefully do not add any further restrictions to start with a relatively large sample, which can be analyzed in more detail later. We restrict the search to journal articles, book chapters, or books in English, which are in the final publication stage. Besides, we limit the time horizon for publications and only focus on literature published between 2000 and 2022.

2.2.3 Study selection and evaluation

To follow a structured approach in the study selection and evaluation, we apply the well-established PRISMA 2020 process flow (Page et al., 2021). By defining a set of meaningful inclusion/exclusion criteria, we narrow the records successively to only include those most relevant to our review question (see Fig. 2.2).

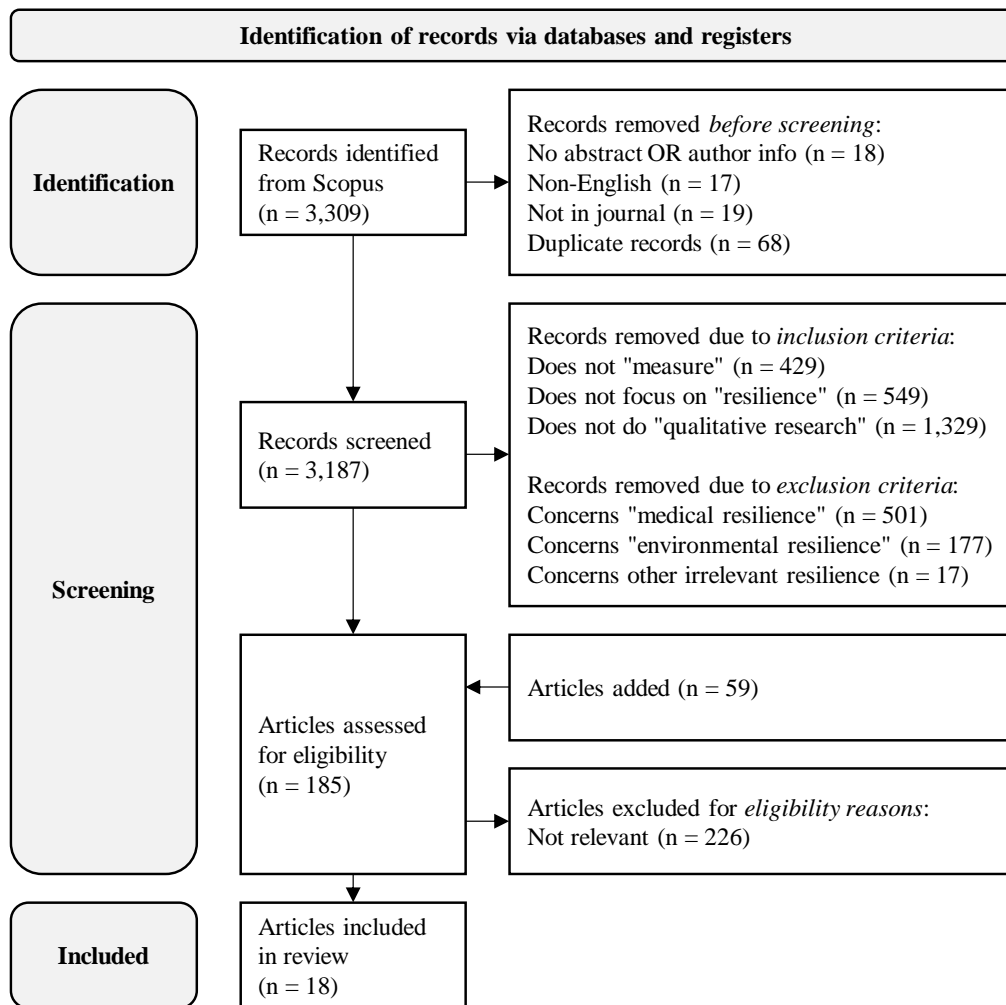


Fig. 2.2: Inclusion/exclusion funnel

First, we apply formal and hygienic criteria to exclude irrelevant records before the screening. We check again for English language and journal publications and only consider records provided with author and abstract information. Furthermore, we remove all duplicates.

Second, we screen the remaining 3,187 records in more detail and only include the ones that meet the following criteria: We assess, based on the abstract, title, and keywords, whether the records indeed (1) measure something (and not just conceptually describe a way to measure), (2) focus on resilience, and (3) use a qualitative research approach. This reduces the number of records to 880 records.

Third, we apply our exclusion criteria and remove all records coming from non-organizational resilience areas irrelevant to our review question, e.g., medical resilience (e.g., Grisanzio et al., 2018; Panter-Brick et al., 2018; Robertson et al., 2015; Sarkar & Fletcher, 2014; Shet et al., 2022; Windle et al., 2011), environmental resilience (e.g., de Moraes Sá et al., 2014; Gazol et al., 2017; Jones & Tanner, 2017; Owen, 2020; Waters & Adger, 2017), or ecological resilience (e.g., Cassidy & Barnes, 2012; Cutter et al., 2016; Fraccascia et al., 2017; Sharifi, 2016; Yazdani & Jeffrey, 2012). The application of the exclusion criteria yields 185 relevant records.

Finally, we assess the records' eligibility by downloading them and reading their abstracts, introductions, and conclusions. In this step, we add 59 articles via cross-referencing, a

commonly used technique in SLR. These 244 articles are then again analyzed in terms of relevance, i.e., if they measure organizational resilience qualitatively. As mentioned before, we are only interested in such articles that apply the resilience measure and not just conceptually develop it (different to, e.g., Hillmann & Guenther, 2021; Hosseini et al., 2016; Pettit et al., 2010; Zeng & Yen, 2017). After further inspection, we remove 226 articles that assess resilience in other fields or with other methods. Ultimately, we include 18 articles for analysis and synthesis.

2.2.4 *Analysis and synthesis*

We read the 18 articles thoroughly to understand better the proposed quantitative method and how the respective authors measure organizational resilience. While reading, we categorize the articles across multiple dimensions (Table 2.1). To ensure that the categorizations are valid, four authors coded each article independently from each other. Discrepancies were discussed in the larger group of authors.

Most importantly, we distinguish the different purposes behind measuring resilience: Some articles focus on providing a measure to quantify organizational resilience. We call this “Resilience Measurement”. Others focus on the antecedents of resilience, such as flexibility, collaboration, or adaptability, and try to measure the predictive power of these on organizational resilience (“Antecedent Analysis”). However, these papers fall short of explaining how to measure resilience itself. We are primarily interested in articles with the purpose of “Resilience Measurement” since this is closely related to the intention of our research. We do not intend to assess which antecedents impact the level of resilience and to what extent but instead want to determine the actual resilience level of an organization.

Ten of the 18 relevant articles measured organizational resilience (“Resilience Management”), whereas eight focused on the antecedents. In the following, we investigate the ten articles concerning “Resilience Measurement” in more detail.

Aleksić et al. (2013) developed a fuzzy mathematical model to assess the organizational resilience of SMEs within the process industry, which they tested as a case study in a Serbian processing company. They asked the management team, consisting of the top manager or owner, the quality manager, and the financial manager, to evaluate eleven resilience factors (categorized in internal actors, external factors, and enabling factors) across six central organizational processes (management, production, marketing and sales, purchase, design and development, and support processes). They find that the sample company has a medium resilience level.

Azevedo et al. (2016) suggested an index that measures not only the resilience but also the leanness, agility, and greenness of automotive companies and the corresponding supply chains. They use multiple management paradigms for each area (seven paradigms for resilience precisely), which are weighted using a Delphi approach. The index is then applied to and validated with six companies in the Portuguese automotive industry, showing a moderate resilience level of the companies assessed in this case study. This tool can further benchmark the companies and the supply chain.

Omidvar et al. (2017) also measured resilience in a single case study with an Iranian gas refinery using a fuzzy analytic hierarchy process. Their resilience assessment framework

consists of six indicators with multiple sub-indicators (43 in total) that are evaluated using expert pair-wise comparisons. The case study indicates medium to good resilience for the gas refinery.

Through interviews, Usher et al. (2019) assessed the resilience of the tourism industry in a coastal city in the United States. They developed a resilience framework based on five components (vulnerability, business planning and operations, preparation and recovery planning, communication, and workforce) and put it to practice via interviews with 32 participants from 42 businesses. Their findings indicate that the businesses do not feel vulnerable to natural hazards and demonstrate high levels of resilience, especially larger businesses.

Brown et al. (2019) focused on the hotel sector in disaster-prone regions. They surveyed hotel general managers and staff in two tourist regions in New Zealand with 72 questions, assessing the resilience level across multiple capitals (economic, social, human, physical, natural, and cultural) to understand their current resilience attributes but also focus on areas for future resilience-building activities. They received 77 responses from 33 hotels and found that the hotels showed positive predictors of resilience across all capitals but also need to work on a “more inclusive approach to disaster planning, add[ing] disaster management to the budget, and [...] improving back-up power generations systems” (N. A. Brown et al., 2019, p. 36).

Graveline and Grémont (2017) emphasize natural disasters and their impact on individual businesses. They propose an operational indicator for the economic resilience of businesses that governments can use to identify areas of low resilience. To show the indicator’s relevance, they applied it via a survey to the Urban Community of Central Martinique. One hundred eight businesses were asked to respond to two hypothetical hazards disrupting the drinking water supply or electricity. Their results indicate that the businesses are relatively more resilient to disruptions of the drinking water (86%) than to electricity shortages (75%). Moreover, they showed that turnover, the flexibility of working hours, and the flexibility of production processes are the factors that significantly contribute to economic resilience.

Brown et al. (2017) assessed the organizational resilience of infrastructure providers with their Benchmark Resilience Tool. It uses thirteen indicators of resilience that are evaluated using a series of corresponding statements. The survey was shared within a major lifeline utility group in New Zealand, and 219 responses were gathered from 18 organizations. Their results show that the organizations have effective partnerships but struggle to break silos and conduct stress tests, which might improve resilience. Since people of different seniority were surveyed, they can also show that senior managers have more positive views than their junior colleagues.

Utami et al. (2021) researched Indonesia’s micro, small, and medium enterprises (MSMEs). They provide a resilience measurement framework consisting of 25 questions. Using a mixed-method approach, they apply it as a structured questionnaire to 50 MSMEs from three sectors (food, salt farming, and craft) in four regencies of an Indonesian island. Their results highlight a low to moderate level of resilience across the three sectors, with salt farming showing the lowest level of resilience.

Table 2.1: Categorization of relevant articles

Authors	Purpose	Benchmark	Industry-agnostic	Definition of resilience score	Internal vs. external disruptions	Backward/forward-looking resilience	Country	Industry	Sample size	Longitudinal	Method	Resilience questions
Azadeh et al. (2017)	Antecedent Analysis		✓				Iran	Aluminum factory	97		Mixed method	30
Chowdhury and Quaddus (2016)	Antecedent Analysis		✓	✓			Bangladesh	Apparel	272		Survey	46
Kaviani et al. (2020)	Antecedent Analysis		✓	✓			Iran	Automotive	31		Survey	50+
Pettit et al. (2013)	Antecedent Analysis		✓	✓			USA	Mixed	170		Survey	Many
Rajesh (2019)	Antecedent Analysis		✓	✓			India	Electronics	1		Case study	23
Sambowo and Hidayatno (2021)	Antecedent Analysis		✓				Indonesia	Manufacturing	8		Interviews	No info
Sapeciay et al. (2017)	Antecedent Analysis	✓	✓				New Zealand	Construction	50		Mixed method	No info
Yazdanparast et al. (2021)	Antecedent Analysis		✓				Iran	Automotive	150		Survey	20
Aleksić et al. (2013)	Resilience Measurement		✓		✓		Serbia	Processing	1		Case study	No info
Azevedo et al. (2016)	Resilience Measurement	✓	✓	✓			Portugal	Automotive	6		Case study	7

(Continued on next page)

Table 2.1: Categorization of relevant articles (*continued*)

Authors	Purpose	Benchmark	Industry-agnostic	Definition of resilience score	Internal vs. external disruptions	Backward/forward-looking resilience	Country	Industry	Sample size	Longitudinal	Method	Resilience questions
N. A. Brown et al. (2019)	Resilience Measurement						New Zealand	Hotels	77		Survey	72
C. Brown et al. (2017)	Resilience Measurement	✓	✓				New Zealand	Infrastructure	219		Survey	No info
Graveline and Grémont (2017)	Resilience Measurement	✓	✓	✓			Martinique	Mixed	108		Survey	44-83
Omidvar et al. (2017)	Resilience Measurement		✓			✓	Iran	Gas refinery	1		Case study	No info
Shirali et al. (2013)	Resilience Measurement		✓				Iran	Manufacturing	88		Survey	61
Usher et al. (2019)	Resilience Measurement		✓				USA	Tourism	32		Interviews	No info
Utami et al. (2021)	Resilience Measurement	✓	✓	✓			Indonesia	Mixed	50		Mixed method	25
Wagner and Neshat (2012)	Resilience Measurement	✓	✓				Germany	Mixed	760		Survey	Many
Total (where applicable)		6	17	7	1	1				0		
<i>This study</i>	Resilience Measurement	✓	✓	✓	✓	✓	Germany	Mixed	199	(✓)	Survey	6

Shirali et al. (2013) approached the measurement of resilience from the angle of resilience engineering. They gathered 88 responses using a questionnaire (with six resilience indicators). They used principal component analysis to determine the importance of the indicators on resilience and the resilience score of 11 units of a company in the process industry in Iran. This way, units with low resilience (e.g., distillation or technical inspection) can be identified to enable managers to act and focus their attention.

Wagner and Neshat (2012) focus their work on supply chain vulnerability for various industries. Based on Normal Accident Theory and High Reliability Theory, they provide a measure for supply chain vulnerability. They administered their survey to a broad range of German firms and collected 760 responses. They can show that larger companies exhibit higher vulnerability than smaller ones and that companies with large series productions are more vulnerable than small series productions.

Table 2.2: Overview of purpose and research method

Purpose	Research method				Total
	Case study	Survey	Interviews	Mixed method	
Antecedent Analysis	1	4	1	2	8
Resilience Measurement	3	5	1	1	10
Total	4	9	2	3	18
Average sample size	2	208	20	66	118

Table 2.2 shows an overview of the purposes and the methods used in the articles. We see that most of the research is conducted via surveys with the intention of “Resilience Measurement”, which aligns with our review question. Case studies and interviews are less popular since they usually entail smaller sample sizes. 16 of the 18 articles use a sample size < 200 for their analysis, while Wagner and Neshat (2012) use a large sample size (Fig. 2.3).

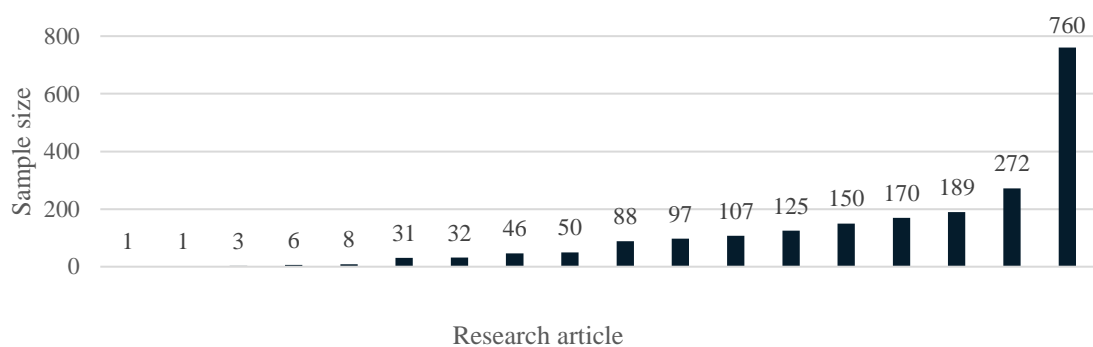


Fig. 2.3: Samples sizes of all 18 relevant papers

Our literature review reveals that only six articles intentionally use their measure as a benchmarking tool for organizations, but none in longitudinal studies. This is one of our index’s most essential features since we believe organizational resilience should be tracked over time and across organizations and industries to identify strengths and weaknesses and facilitate learning. Moreover, we believe that organizational resilience is a characteristic that should be tracked and monitored over time and be used to compare companies. Out of these six

benchmarking articles, only one resilience measurement (Wagner & Neshat, 2012) looks at the German economy. However, since it is from 2012, it is outdated and limited in its focus on the supply chain and logistics functions – a refreshed resilience measure in the German manufacturing industry is utterly needed.

Moreover, we see that the resilience definitions and the methodology used in the articles to compute the resilience levels are not always transparent: Only seven of the 18 articles in total and three of the six benchmarking articles provide sufficient information, which allows fellow researchers and practitioners to replicate the approach. However, we argue that practitioners need this transparency about the resilience definition and the methodology to apply the resilience measure in practice and trust the results, especially when comparing themselves to others. The practitioners we spoke with in the SPAICER project confirmed this. We also reveal that most authors provide industry-agnostic resilience measures, which are generic enough to be applied in various industries. Only Brown et al. (2019) present a measure designed specifically for the hotel industry; thus, it cannot be easily transferred to other industries such as manufacturing.

In terms of geography, we conclude that most of the research is done outside Europe: Only Wagner and Neshat (2012) study organizations in Germany, Azevedo et al. (2016) in Portugal, and Aleksić et al. (2013) in Serbia. Most resilience measures are done in Iran (five articles) and New Zealand (three articles).

Finally, we acknowledge that most of the articles (11/18) are more than five years old, highlighting the need for new data, given the recent advances in resilience research and events affecting organizations' resilience in the last few years.

2.2.5 *Reporting and using the results*

The results confirm that our research is indeed novel for many reasons: there is a lack of short and representative surveys for organizational resilience, especially in Europe. Most approaches are complex and require additional resources for data gathering, e.g., on-site researchers to support, explain the survey or conduct the interviews. Thus, they also take a long time to implement. Moreover, very few studies intentionally perform benchmarks, and none do longitudinally.

To close this gap, we conduct our own empirical study with a survey that can be easily repeated to track resilience on a large scale and over time.

2.3 Methodology

Given our SLR and our experience from the SPAICER project, we find that many measures for organizational resilience are rather complex and based on long and complicated surveys or interview protocols, hindering the large-scale implementation of such a measure. Thus, we want to tackle both issues and develop a new, easy-to-understand resilience measure based on a simple survey that can be implemented at scale while also providing a representative and reliable picture of the overall resilience level within an organization and industry⁵. Furthermore,

⁵ In our initial evaluation, however, we concentrate on the German manufacturing industry.

the simplicity of our approach should enable the collection of longitudinal data to track changes more easily over time.

2.3.1 Survey design

Resilience has many features, but they all have a common denominator: disruptions and the ability to handle them (Tukamuhabwa et al., 2015). After all, without disruptions, there would be no need for resilience. In our work, we shed light on two dimensions of disruptions and resilience specifically: *time* and *locus of control* (Fig. 2.4).

		Locus of Control	
		Internal	External
Time	Past	Q1: Past perceived internal resilience	Q2: Past perceived external resilience
	Present	Q3: Current perceived internal resilience	Q4: Current perceived external resilience
	Future	Q5: Anticipated need for internal resilience	Q6: Anticipated need for external resilience

Fig. 2.4: Dimensions of resilience

The time dimension is relevant since resilience is a dynamic capability and changes over time (Chowdhury & Quaddus, 2016; Pereira et al., 2014; Ponomarov & Holcomb, 2009). Companies most likely actively manage their resilience and adapt their resilience to past experiences and expectations for the future. To account for the dynamics, we establish a backward- and forward-looking perspective on resilience and differentiate between the past level of resilience, the current level of resilience, and the anticipated need for resilience. However, the time dimension is overlooked in most research. Only one other research in our SLR dataset (Omidvar et al., 2017) has established a backward- and forward-looking perspective on resilience. With the time dimension, we can analyze an organization's ability to assess expected resilience levels adequately. Ideally, next year's current resilience should align with this year's assessment of an organization's expected resilience.

On the other hand, the locus of control is also relevant because disruptions and sources of uncertainty can be manifold (e.g., internal vs. external (Aleksić et al., 2013)), but they all affect organizations. Thus, independent of the source, companies need resilience to manage them. Nonetheless, we expect that the type of disruption affects the (perceived) controllability of these disruptions. Such differentiation has not been captured in most resilience measures thus far (see Table 2.1).

The survey design has been discussed and validated with our industry partners from the German manufacturing industry taking part in the research project SPAICER. From a practical standpoint, many managers are concerned with these two resilience dimensions. They need to manage their organization's resilience over time while assessing the right level of future resilience to use resources efficiently. At the same time, they understand that the sources of

disruption and, thus, its controllability matter, especially in their communication of the disruption and their respective actions.

This focus allows us to keep the survey as short as possible while still getting a nuanced picture of the analyzed resilience levels. For each of the six resilience dimensions (i.e., resilience items), we ask the survey participants to rate the resilience level on a five-point Likert scale.

Additionally, we collect demographic information on the respondent's company's industry, the company's annual revenue, and the number of employees. We use this demographic information as control and clustering variables in our detailed data analysis (see Chapter 2.4 *Results*). We do not collect any more firm-specific data as we want to keep a maximum level of anonymity so that respondents are willing to share an honest evaluation of their perceived resilience levels. For the demographic information, respondents must indicate which category best fits their respective organization⁶.

2.3.2 *Enterprise Resilience Index*

We learned from our industry partners that managers ideally aim for a single meaningful indicator that is easy to understand and collect. While such an indicator should not be taken as the only tool for resilience management, it can prove valuable as one of the multiple tools in a resilience manager's toolkit. Especially in the case of managing resilience, having a simple, early indicator that can be obtained relatively quickly and more frequently than highly sophisticated resilience measures that might provide a more detailed view of specific resilience capabilities can be critical, especially in highly volatile markets and times of great uncertainty. Our goal is not to provide the perfect resilience indicator but one that strikes the right balance between complexity, predictive power, and practicability. That is why we developed the Enterprise Resilience Index (ERI).

In our quest to develop the ERI, we took inspiration from the ifo BCI. The BCI is one of the most important indicators for Germany's economic development and is published monthly (ifo, 2022). It sparked our interest for various reasons: First of all, the underlying survey is simple as it only consists of two components that jointly make up the index's score; Second, the index considers the current business situation and the businesses' expectations of their future course of business (Seiler & Wohlrabe, 2013), highlighting the focus on changes and not on absolute levels; Finally, it has proven to be an excellent early indicator for Germany's gross domestic product (GDP) (Sinn & Abberger, 2006).

For the ERI, we determine an overall propensity p_i for each item, i.e., whether respondents evaluate the respective resilience situation as positive or negative. Therefore, we calculate the weighted average (response options weighted by a number from 0 (disagree strongly) to 4 (agree strongly)) of responses and multiply the results by 25 to normalize the weighted average to a range from 0 to 100 (see Equation 1).

⁶ The entire survey can be found in the appendices A. 1 & A. 2.

Equation 1: Propensity per item

$$p_i = \frac{\sum_{j=0}^4 n_{i,j} \cdot j}{N} \cdot 25$$

$n_{i,j}$: number of responses per item (i) and response option (j)

N: total number of responses

j: strongly disagree ... strongly agree = 0 ... 4

With this propensity, we then calculate an overall score for the *past perceived resilience*, *current perceived resilience*, and *anticipated need for resilience* (see Equation 2 to Equation 4), considering both the internal and external perspectives. In turn, we use the past and current perceived resilience to calculate the ERI (see Equation 5). We base our calculations on the geometric rather than the arithmetic mean because extreme values do not influence the geometric mean as much as they do the arithmetic mean. We argue that a company (A) that has a resilience level of 100 in one year and 10 in the next is less resilient than a company (B) that consistently has a resilience level of 55 in both years. If the arithmetic mean were considered, both companies would have the same mean resilience level of 55, whereas company A has a mean resilience level of 31.6 when using the geometric mean.

Equation 2: Past perceived resilience

$$R_{past} = \sqrt{p_{past\ int} \cdot p_{past\ ext}}$$

Equation 3: Perceived current resilience

$$R_{current} = \sqrt{p_{current\ int} \cdot p_{current\ ext}}$$

Equation 4: Anticipated need for resilience (NFR)

$$NFR = \sqrt{p_{nfr\ int} \cdot p_{nfr\ ext}}$$

Equation 5: Enterprise Resilience Index (ERI)

$$ERI = \sqrt{R_{past} \cdot R_{current}}$$

Our ERI represents a new resilience measure that is easy to implement at scale, providing practitioners with a tool to benchmark their (perceived) organizational resilience level in cross-industry benchmarks. Since it considers both the past and the current resilience level, it can be used neatly to track resilience over time and to determine trends, like the BCI. We hope that the ERI can be added as another valuable tool to the toolkit of resilience managers. As this is the first time we have collected data on the ERI, it will remain unclear how much this index correlates with other indicators of economic development like the BCI or GDP. Still, as we have seen with the example of the BCI, this approach holds the potential to provide an important early indicator of perceived and actual resilience levels.

2.3.3 Data collection

We created an online survey. We used various channels to recruit as many relevant participants from the German manufacturing companies as possible. Since we developed the ERI as part of the SPAICER project, we contacted all consortium members and asked them to share the link to our online survey within their company and professional network. Furthermore,

we contacted industry experts directly and distributed the survey via mailing lists from industry associations involved in the SPAICER project.

This year’s data collection phase concentrates specifically on the German manufacturing industry. Germany has a particularly strong manufacturing industry that is responsible for large amounts of Germany’s GDP (World Bank, 2022). Furthermore, this industry has seen major disruptions recently, highlighting a need for resilience, which can only be adequately addressed if we understand the actual resilience level within an industry. Ultimately, we want to extend our evaluation to more industries and potentially other countries.

2.4 Results

Following, we lay out our *data preprocessing* strategy, present *descriptive statistics* for the preprocessed dataset, and perform statistical *data analysis*.

2.4.1 Data pre-processing

In total, 210 industry representatives completed our survey (corresponding to a completion rate of 80%). These were checked for response quality and inconsistencies (such as matching revenue figures and company sizes), resulting in nine responses being removed. Additionally, two responses were removed since the participants do not work in the manufacturing industry, our research focus. This resulted in a final dataset consisting of 199 responses, which is, compared to other studies, a relatively large sample size (see Literature Review).

To ensure the analyses are meaningful, i.e., to enable cross-industry cluster analyses, we further cluster the industries provided by the participants into six reasonable larger industry clusters: *Automotive, Chemicals, Food and Consumer Goods, Machinery, Manufacturing, and Others*.

2.4.2 Descriptive statistics

Table 2.3 gives an overview of the distribution of survey responses regarding their (1) industry cluster, (2) annual revenue, and (3) number of employees.

Across all categories, all but one sub-group has a sample size >10 , enabling further meaningful analyses. We see a relatively balanced distribution regarding industry clusters, emphasizing Machinery (55) and Manufacturing (44). Most respondents (109) come from “large enterprises” (LEs), defined as companies with >1 billion EUR in revenues or $>5,000$ employees (Insee, 2020). Yet, we also gathered a significant number of responses (38) from small and medium-sized enterprises (SMEs), defined as companies with <50 million EUR in revenues and <250 employees (European Commission, 2003).

Table 2.3: Survey participants by industry cluster, revenue, and size

Feature	Frequency	Percentage
Industry cluster		
Automotive	35	18
Chemicals	13	6
Food and Consumer Goods	18	9
Machinery	55	28
Manufacturing	44	22
Others	34	17
Total	199	100
Revenue, in million €		
< 10	16	8
10 – 50	25	13
50 – 100	16	8
100 – 500	16	8
500 – 1,000	17	8
1,000 – 5,000	46	23
> 5,000	63	32
Total	199	100
Number of employees		
< 100	23	12
100 – 250	22	11
250 – 500	11	5
500 – 1,000	9	4
1,000 – 5,000	25	13
5,000 – 10,000	19	10
> 10,000	90	45
Total	199	100

Fig. 2.5 visualizes all survey responses. Outliers in Q3 and Q5 are marked with a diamond symbol. Overall, we see high average levels of resilience, whether past or current, and internal or external perceived resilience. The boxplot also indicates a relatively large variance within the dataset, which we seek to analyze in the following chapter. We can also see that the current perceived resilience in both the internal and external dimensions (Q3 & Q4) is below past perceived internal and external resilience (Q1 & Q2). Furthermore, there appears to be a strong anticipated need for external resilience (Q6).




Fig. 2.5: Boxplot visualization of all responses

Table 2.4 presents a descriptive statistical analysis of the responses, including the median, mode, mean, and standard deviation for each survey item. It confirms the “strong”, i.e., median, mode, or mean greater than 2 (i.e., neutral), levels of resilience. However, a more nuanced picture is possible here: The past perceived internal and external resilience levels are slightly higher compared to the current levels in terms of the mean, whereas the median and mode are mostly the same (except for Q4). Indeed, the current perceived external resilience (Q4) shows the lowest median and mean compared to the items Q1, Q2, and Q3, indicating ongoing difficulties in dealing with external disruptions. This is reinforced by Fig. 2.6, which visualizes the distribution of the responses to items Q1, Q2, Q3, and Q4. Table 2.4 also shows that the anticipated need for external resilience (Q6) far exceeds the anticipated need for internal resilience (Q5). Fig. 2.7 shows that the answers to Q5 are almost the flipsides to the answers in Q6.

Table 2.4: Descriptive analysis of all responses

Item	Median	Mode	Mean	SD
Q1: Past perceived internal resilience	3	3	2.77	1.03
Q2: Past perceived external resilience	3	3	2.91	0.93
Q3: Current perceived internal resilience	3	3	2.47	1.01
Q4: Current perceived external resilience	2	3	2.15	0.98
Q5: Anticipated need for internal resilience	1	1	1.69	1.07
Q6: Anticipated need for external resilience	3	3	2.90	0.97

Note: “Disagree strongly” corresponds to a value of 0, “agree strongly” to 4

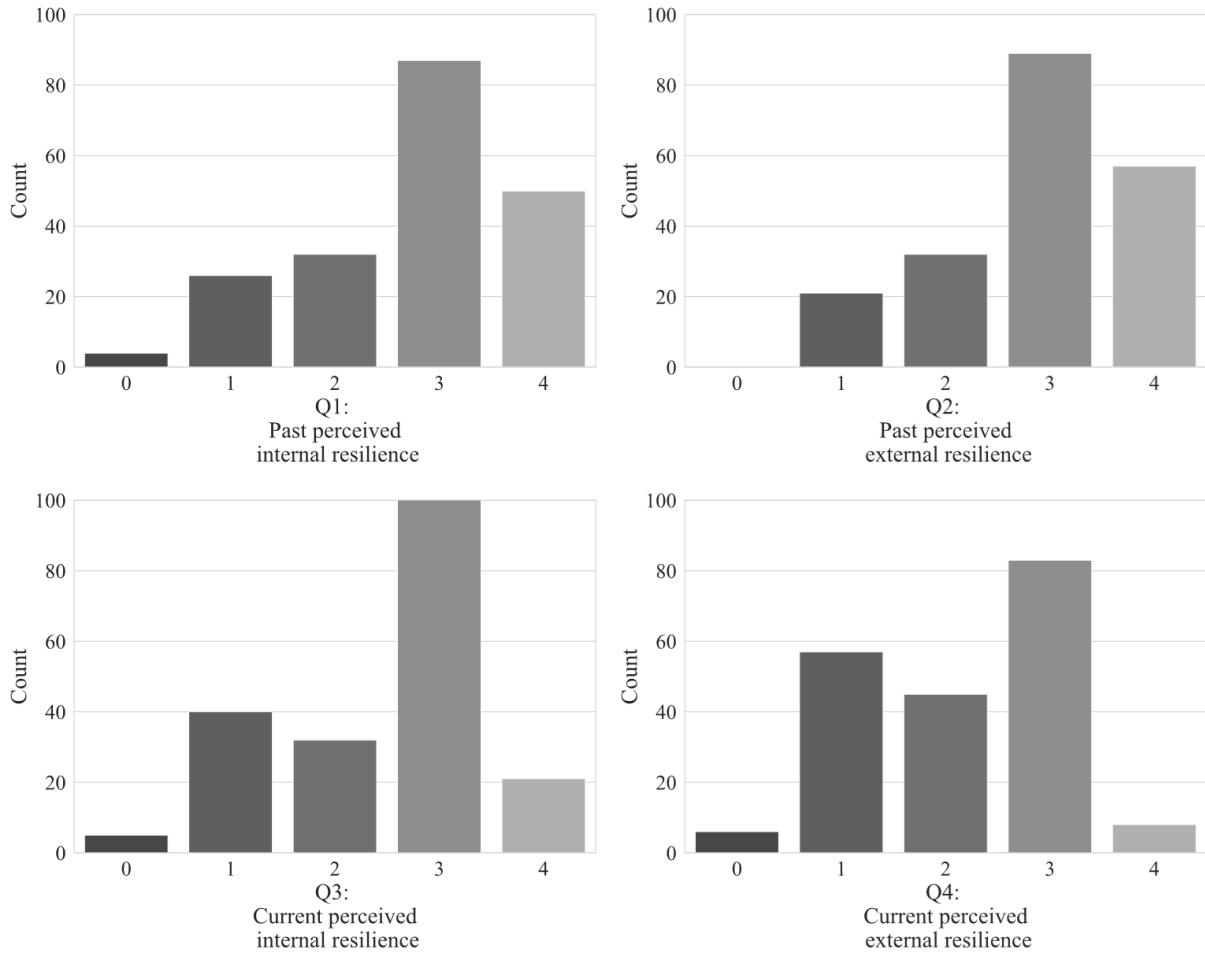


Fig. 2.6: Distribution of responses to items Q1, Q2, Q3, and Q4




Fig. 2.7: Distribution of responses to items Q5 and Q6

2.4.3 Data analysis

The descriptive statistics already suggest some interesting insights. In the following, we look deeper into the data using statistical tests such as Spearman’s rank correlation (Spearman, 1987) or the Mann-Whitney-U test (Mann & Whitney, 1947). First, we analyze the data in general, ignoring the demographic information on the respondent’s company’s industry, the

company's annual revenue, and the number of employees. Then, we take these into account and look for differences in sub-samples.

2.4.3.1 General analysis

Table 2.5 shows the correlation matrix with annotated Spearman's rank correlation coefficients ρ_{ij} . First, we look at the dimension *locus of control*. We see a significantly positive correlation between all internal and external items, i.e., between past perceived internal and external resilience ($\rho_{12} = 0.49$), current perceived internal and external resilience ($\rho_{34} = 0.34$), and the anticipated need for internal and external resilience ($\rho_{56} = 0.21$). Companies with lower internal resilience levels tend to show lower external resilience levels and vice versa. In other words, the internal resilience level spills over onto the external. This may be because both internal and external disruptions need to be managed simultaneously since resilience as a capability is holistic: It only works if all potential disruptions and uncertainties are addressed. Similarly, external disruptions can impact internal processes and vice versa; thus, both aspects must be managed.

Table 2.5: Spearman's rank correlation matrix

	Q1	Q2	Q3	Q4	Q5	Q6
Q1: Past perceived internal resilience		0.49**	0.63**	0.22**	-0.36**	0.02
Q2: Past perceived external resilience			0.40**	0.42**	-0.19**	0.06
Q3: Current perceived internal resilience				0.34**	-0.27**	0.08
Q4: Current perceived external resilience					-0.01	0.02
Q5: Anticipated need for internal resilience						0.21**
Q6: Anticipated need for external resilience						

Note: * shows significance at the 5%-level ($p < 0.05$), ** at the 1%-level ($p < 0.01$)

Secondly, we compare the *time* dimension, i.e., the past, current and anticipated levels. Here, we see that past perceived internal resilience is significantly positively correlated with the current perceived internal resilience ($\rho_{13} = 0.63$), and the same holds for the external view ($\rho_{24} = 0.42$). This seems intuitive since the past resilience levels should affect the current ones. Looking into the future, we can also determine significant correlations, but this time negative ones: the anticipated need for internal resilience is significantly negatively correlated with the past perceived internal resilience ($\rho_{15} = -0.36$), and the current perceived internal resilience ($\rho_{35} = -0.27$). Arguably, companies that experienced high levels of resilience in the past or presence anticipate only a small need for resilience in the future, maybe because they already are at such high levels or because they expect to be in control of the disruptions. Conversely,

companies with low levels in the past or presence expect higher levels in the future, showing their development need. The external perspective shows no correlation.

Next, we analyze the levels of the six items and compare these using the Mann-Whitney-U tests (Table 2.6). This test assesses whether the answering behavior for one item differs significantly from another. In Table 2.6, the null hypothesis for the test assumes that the underlying distributions of the respective items are the same. The alternative hypothesis suggests that the underlying distribution of item 1 is different from the underlying distribution of item 2.

Again, starting with the *locus of control*, we would expect that the internal dimension of each resilience item is greater than the external one, thinking that companies focus more on the internal level. The results confirm our expectations only partially: Whereas the perceived internal resilience is significantly higher than the external one in the presence ($Q3 > Q4$), we do not see significantly different resilience levels in the past ($Q1 > Q2$). Moreover, companies tend to have a higher anticipated need for external than internal resilience ($Q6 > Q5$). Potentially, they expect more external disruptions or believe their resilience levels are high enough to withstand internal disruptions already, whereas recent external disruptions have shown how vulnerable most companies are. Overall, the statistical analysis suggests that it is worthwhile to differentiate between internal and external resilience.

Table 2.6: Mann-Whitney-U test for various hypotheses

Hypothesis	Mean of item 1	Mean of item 2	p-value
$Q1 > Q2$	$\mu = 2.77$	$\mu = 2.91$	0.217
$Q3 > Q4$	$\mu = 2.47$	$\mu = 2.15$	0.001**
$Q6 > Q5$	$\mu = 2.90$	$\mu = 1.69$	0.000**
$Q3 > Q1$	$\mu = 2.47$	$\mu = 2.77$	0.002**
$Q4 > Q2$	$\mu = 2.15$	$\mu = 2.91$	0.000**

Note: * shows significance at the 5%-level ($p < 0.05$), ** at the 1%-level ($p < 0.01$)

Regarding the *time*, we assume that companies strive to improve their resilience. Thus, current resilience levels should be higher than past ones and future ones higher than current ones ($Q3 > Q1$ and $Q4 < Q2$). This argument would be in line with resilience being a dynamic capability that can change over time. Many authors suggest a growth phase in their resilience definitions (e.g., Pettit et al., 2010; Ponis & Koronis, 2012; Ponomarov & Holcomb, 2009; Tukamuhabwa et al., 2015) in terms of performance, which might also be applied to the underlying resilience capability. In contrast to our assumption, current perceived resilience levels fall significantly below past perceived resilience levels, especially for perceived external ones. This may be because the ongoing war in Ukraine is putting further pressure on companies after already troublesome years of the COVID-19 pandemic, mainly due to supply chain disruptions. Thus, companies realize they are not as resilient as they thought they would be. This finding further emphasizes the need for a methodology to gather data representing resilience levels over time, as suggested in this work.

2.4.3.2 Analysis of industry and company size

After the general analysis of our survey results on the perceived resilience levels across the two dimensions of *time* and *locus of control*, we now examine the resilience of sub-samples, considering the industry and the size of the participating companies. The ERI and NFR can be applied to get an aggregated view of multiple resilience items per sub-sample.

Table 2.7: Means of ERI and NFR per industry cluster

Industry cluster	ERI	NFR
Automotive	69.4	61.5
Chemicals	73.8	56.4
Food and Consumer Goods	61.4	53.9
Machinery	54.5	55.7
Manufacturing	69.2	47.4
Others	63.8	58.4
<i>Weighted mean</i>	<i>63.9</i>	<i>55.2</i>

In terms of ERI, Table 2.7 shows higher than average levels of resilience for the Automotive, Chemicals, and Manufacturing industries. In contrast, the Machinery industry is the only industry that has a lower ERI. The Machinery industry's low ERI is in line with PWC's most recent "Machinery Barometer" (PWC, 2022), indicating that six out of ten representatives from the Machinery industry see the need to increase their resilience as a top priority, which we cannot confirm with the Machinery industry's NFR.

For the NFR, we only see two outliers: the Automotive industry has a higher NFR than the average, whereas the Manufacturing industry has a lower one. This result shows a striking difference between the Automotive and the Manufacturing industry and potentially different future outlooks: While both have similar ERIs (69.4 vs. 69.2), the Automotive industry has a higher-than-average NFR whereas the Manufacturing industry one is lower. This might support various arguments: For instance, the Automotive industry might expect even more future disruptions and thus, already highlights the need to increase resilience. Perhaps, the appropriate level of resilience in the Automotive industry is higher than in the Manufacturing industry, justifying the additional need for resilience in the future. While all industry sectors were and are affected by severe disruptions (e.g., through the COVID-19 pandemic or the war in Ukraine), the Automotive industry was especially affected by these disruptions in recent years. On the contrary, the Manufacturing industry might have reached a resilience level that is too high and thus, not efficient anymore, thus showing a lower NFR.

Fig. 2.8 visualizes the ERI's two components (past and current perceived resilience) for each industry cluster. It is noteworthy that all industry clusters report lower current perceived resilience levels than past perceived resilience levels, which we already remarked on aggregate in Table 2.4. Furthermore, Fig. 2.8 demonstrates that the industry cluster Machinery shows the lowest measured levels of both current and past perceived resilience, whereas the Chemicals industry has by far the highest current perceived resilience and very high past perceived resilience.




Fig. 2.8: Visualization of the ERI per industry

Next, we analyze the difference in perceived resilience regarding the size of the companies. First, we use the Mann-Whitney-U test to compare the 109 LEs with the 38 SMEs in our data in terms of the six items.

The null hypothesis for the test assumes that the underlying distribution of the sample SMEs is the same as the underlying distribution of sample LEs for each item. The alternative hypothesis suggests that the underlying distribution of the sample SMEs is different from the underlying distribution of sample LEs.

Table 2.8: Item-wise results for the Mann-Whitney-U tests for SMEs vs. LEs, items Q1, Q2, Q3, and Q4

Item	p-value	Mean of SMEs	Mean of LEs	SD of SMEs	SD of LEs
Q1: Past perceived internal resilience	0.620	2.79	2.80	0.89	1.07
Q2: Past perceived external resilience	0.825	3.00	2.92	0.86	0.95
Q3: Current perceived internal resilience	0.048*	2.29	2.66	1.10	0.93
Q4: Current perceived external resilience	0.201	2.00	2.26	1.23	0.97
Q5: Anticipated need for internal resilience	0.033*	1.47	1.90	0.82	1.07
Q6: Anticipated need for external resilience	0.002**	2.58	3.17	1.02	0.82

Note: * shows significance at the 5%-level ($p < 0.05$), ** at the 1%-level ($p < 0.01$)

Table 2.8 presents the results of the Mann-Whitney-U test. Whereas the current perceived internal resilience levels of LEs significantly exceed the perceived resilience levels of SMEs (Q3), the same does not hold for past perceived resilience levels (Q3 & Q4) or the current

perceived external resilience (Q4). We might see higher levels of resilience at LEs because they have more resources to dedicate to resilience management, especially in the current period. Looking forward, LEs also show a significantly higher NFR than SMEs (Q5 & Q6), which may be due to complex and fragile supply chain structures and, thus, more uncertainty. These findings support Sapeciay et al. (2017), who find that LEs are more resilient than SMEs. Interestingly, Wagner and Neshat (2012) concluded that LEs are more vulnerable to (supply chain) disruptions than SMEs. Potentially, this explains our findings in terms of the NFR: Despite the LEs' relatively high levels of resilience, they are also aware of the potential disruptions and their vulnerabilities and thus, anticipate a higher NFR.

2.5 Discussion

Our results provide first and diverse insights into the perceived organizational resilience of companies in the German manufacturing industry. Given our large sample size, especially compared to sample sizes of other research contributions, we are confident that the results are relevant, reliable, and representative. Regarding the survey design, we received feedback from our industry partners in the SPAICER project, mentioning that they would prefer our short and easy-to-understand survey to existing concepts. The high number of responses and the 80% completion rate confirm that our survey design is generally well-accepted. We are convinced that the data collection can be reproduced quickly, for instance, for management reports or potential longitudinal analyses. Being able to paint a better picture of the state of resilience in an industry might be beneficial in itself for increasing resilience over time, considering the relevance of information sharing for, e.g., supply chain performance (Baah et al., 2022).

We learn that overall perceived organizational resilience is relatively high in the German manufacturing industry. Arguably, resilience has recently been in the media broadly and, thus, on top of many managers' minds (Hillmann & Guenther, 2021; Mena et al., 2020). The German manufacturing industry has been shaken by many disruptions in the last few years, be it the Covid-19 pandemic, the semiconductor shortage, or the ongoing war in Ukraine. However, past resilience levels are higher than current ones across all industries. Potentially, companies were in a state of "over-resilience" and thus, reduced or are planning to reduce their resilience levels. Alternatively, companies expect more disruptions or the challenges to change; therefore, they become aware that they are not as resilient as they thought (Lund et al., 2020). In any case, we see that resilience is a dynamic capability worth tracking over time.

Our work bears several managerial as well as theoretical implications. First, our survey enables industry professionals to benchmark themselves against peers. At this stage, the industries are still relatively high-level to allow for more granular statistical analysis, but, in the future, with larger sample sizes, more nuanced industry codes can be applied, thus facilitating the like-for-like comparability of companies. In addition, a cross-industry benchmark could also be of high value for practitioners as they could be encouraged to search for best practices in other industries. Potential future best practice sharing of "resilience leaders" can also enable "resilience laggards" to improve their resilience levels. Second, this work was set up to facilitate benchmarking over time with longitudinal studies. Since the survey is easy to understand and complete, companies should be able to integrate it easily into their business routines and

management reports. For example, the survey could potentially be used within a company to continuously and granularly track the perceived organizational resilience and identify issues early on. In addition, managers could use it as a valuable learning opportunity: they could compare the actual resilience with their assessment in hindsight and thus gradually improve their ability to assess their company's resilience. Furthermore, our and other analyses (e.g., OECD, 2021) suggest that demographic information such as industry characteristics (i.e., industry cluster) and company sizes (SME vs. LE) are essential in measuring resilience, including the past and current resilience levels or the anticipated need for resilience. However, we notice that not all research in this field routinely considers demographic information, e.g., when investigating the antecedents of resilience. Thus, we argue that demographic information should be integrated to create a more nuanced view of resilience in the future.

The survey design is a good start for future discussions and research efforts – further research can be executed to discuss complementary questions. For instance, deep-dive interviews with people from the Manufacturing industry might help understand if they feel over-resilient. Similarly, other industry professionals could share their detailed insights into their resilience journey, explaining if they have actively reduced their resilience levels and, if so, why. Since the survey is set up to be used longitudinally, promising results are also expected from the following data collection period. Then, we can see if the NFR tracks with the actual future resilience levels. The visualization of the ERI (see Fig. 2.8) will help us picture the development over time. Furthermore, after several collection rounds, the results can be correlated with other economic factors, e.g., the GDP, to check if the ERI and the NFR can be used as an early indicator for future disruptions, similar to the ifo BCI (Sinn & Abberger, 2006). Finally, this work needs to be replicated in more countries and industries to paint a more comprehensive picture of an economy's resilience.

This research also comes with its limitations. First, the regular limitations regarding survey research also apply here (see, e.g., Almeida et al., 2017; Coughlan et al., 2009). Specific to our work, we would like to square our resilience measure with a company's "actual resilience", measured as the performance development in times of disruption. Ideally, companies with high levels of perceived resilience would fare better than those with low levels. Unfortunately, this reconciliation with the ground truth requires a disruption in the real world, making this approach impractical. Alternatively, we could compare our resilience measure to other resilience measures in the field. While there is no standard measure, we believe that such a comparison with a small sample might help validate our resilience measure and potentially help the community make a step towards a standard measure for resilience. Finally, the strength of our survey approach also entails a major drawback: We do not know the underlying reasons for the results and can only speculate about them. Thus, further research, e.g., through interviews or case studies with focus groups, is needed to establish this link.

2.6 Conclusion

Recent events such as the COVID-19 pandemic, the war in Ukraine, or the all-time low water levels in the Rhine emphasize the need for resilience. To manage resilience, we need to measure it.

Our systematic literature review suggests that existing (qualitative) resilience measures are complex, challenging to interpret, and therefore hard to scale and apply across multiple industries. Thus, we establish a novel resilience measure, a survey-based index for (perceived) resilience levels across two dimensions, locus of control (internal, external) and time (past, current, future). Moreover, we develop the Enterprise Resilience Index (ERI), an easy-to-adapt indicator for the organizational resilience level.

We also collect empirical data using our newly established survey. Our results show that the German manufacturing industry perceives itself as relatively resilient, with significant differences between industries and company sizes. We also see that they anticipate a high need especially for external resilience across industries in the future. Most strikingly, the Machinery industry shows the lowest ERI levels while it anticipates a relatively high need for resilience, showing that this industry, in particular, needs to work on improving its resilience.

This work offers multiple avenues for future research: First, the survey should be repeated over time to leverage the longitudinal capability of this approach to track (and manage) resilience levels over time, ideally with an even larger sample size similar to the ifo business climate index. Second, our survey-based research can be complemented with a mixed-method research approach, e.g., by deep-dive interviews and targeted case studies. The mixed-method approach can help to explain causes for year-over-year changes in the ERI. Potentially, these changes can be linked to specific actions taken by a company. And finally, the perceived resilience levels can be compared to the “real” resilience, when another significant disruption happens since resilient companies or industries should fare better in times of crisis. This way, the measure can be validated to see how well individuals can judge their company’s resilience levels.

3 OPTIMIZING CONTROL OF WASTE INCINERATION PLANTS USING REINFORCEMENT LEARNING AND DIGITAL TWINS

This chapter is based on Schlappa et al. (2022).⁷

3.1 Introduction

As of today, 66 waste incineration plants (WIP) are active across Germany (Flamme et al., 2018), accounting for 3.7% of Germany’s end energy consumption (Weber et al., 2020). Despite the transition to renewable energy sources such as solar and wind, WIPs are here to stay for the foreseeable future because they offer predictable energy outputs, help meet consumers’ electricity and heating needs, and mitigate the environmental impact of landfills. While other “conventional” power sources such as lignite, coal, or nuclear have experienced a steady decline (Appunn et al., 2021) or have remained constant (Elhegazy & Kamal, 2022), the output from WIPs has increased slightly, mainly due to efficiency gains (NABU, 2019).

Yet, it has been argued that WIPs can be operated even more efficiently since many WIPs are at least in part operated manually by human operators (Zhan et al., 2021). Managing the combustion process is a complex task (Adams et al., 2021; Flynn, 2003); thus, highly skilled and experienced personnel are required (Zhang et al., 2010). Such personnel will likely be harder to find (Kaneko et al., 2019), and even today, major performance differences can be seen between human operators.

Various automation systems exist and are already applied to industrial systems such as WIPs. For instance, rule-based process automation is usually used to avoid or mitigate major breakdowns (e.g., shutdowns due to temperature or pressure), and proportional–integral–derivative (PID) controllers help adjust parameters around a specific set point (Takaghaj et al., 2014). Unlike a human operator, these systems operate 24/7 with no breaks, but they tend to only work well in simple, linear environments (Zhan et al., 2021; Zhuang et al., 2018). Unfortunately, the combustion process within a WIP is more complex, so frequent human intervention is needed to improve performance (Stephan et al., 2001).

More advanced controls such as model predictive control (MPC) can also be applied (Majanne, 2005). But these systems have their limitations: They usually require significant amounts of engineering and finetuning, a model of the process is needed, the computation time is online, and it can be long, applications of MPC in non-linear, dynamic processes are rare, and in general, MPC is not very flexible (Di Cairano, 2012; Forbes et al., 2015; Kubosawa et al., 2019; Mayne, 2014; Nian et al., 2020; Qin & Badgwell, 2003; Schubnel et al., 2020). Thus, MPC, which originated in the realm of industrial chemistry (Di Cairano, 2012), is only rarely (if not at all) used in power plant environments (Qin & Badgwell, 2003).

Reinforcement learning (RL) provides a new approach to automate and optimize the WIP using “learning by doing” and thus, overcoming many of the abovementioned issues. An agent

⁷ This paper is based on Schlappa, M., Spinler, S., & Hegemann, J. (2022). Optimizing Control of Waste Incineration Plants using Reinforcement Learning and Digital Twins. *IEEE Transactions on Engineering Management*

is trained via “trial and error”, given a reward function that considers the operational trade-offs of maximizing steam output while keeping emission levels low. This way, performance levels can be reached that exceed human performance. Since training within the real environment is neither possible nor desirable (Roberts & Rousseau, 1989; Steege et al., 2010), a data-driven simulation (“digital twin”) of the WIP is built using a neural network architecture.

Thus, this research contributes to the application of RL in complex, industrial systems, an area that is still vastly under-researched compared to more classic approaches such as MPC, especially in the energy sector (Perera & Kamalaruban, 2021). It uses a significant amount of original plant data, proposes a highly data-driven process to build a digital twin of the WIP, and discusses a procedure for training an RL model in an offline mode. Thus, it helps bridge the gap between theoretical research on control problems and real-life industrial issues, demonstrating that RL can be used to operate and optimize a WIP.

The rest of this paper is organized as follows: Chapter 3.2 discusses applications of RL in industrial systems such WIPs. Chapter 3.3 presents the setup of the digital twin of the WIP, the RL agent’s design, and the RL agent’s interaction with the digital twin. Chapter 3.4 demonstrates the quantitative and qualitative performance of the RL agent and the limitations of this approach. Chapter 3.5 highlights the main conclusions from this application, possible future research, and the managerial implications.

3.2 Literature review

RL has been researched for decades (Gosavi, 2009; Yuxi Li, 2018; Schmidhuber, 2015; Sutton & Barto, 2018; Tesauro, 1994), and recently, it has experienced a significant increase in interest due to “benefiting from big data, powerful computation, new algorithmic techniques, mature software packages and architectures, and strong financial support” (Yuxi Li, 2018, p. 5). The interest has been amplified by the super-human performance in Atari Games (Mnih et al., 2013, 2015) and the victory of Google DeepMind’s AlphaGo against Lee Sedol in March 2016 (Silver et al., 2016; Sutton & Barto, 2018). Please refer to Nian et al. (2020) for a more detailed timeline. Outside the realm of video games, RL has been applied to learn various tasks: for instance, Behnke and Bennewitz (2005) teach a robot to play soccer, Arel et al. (2010) use RL for traffic light control, whereas Grissom et al. (2014) use it for simultaneous machine translation, Zhou et al. (2017) optimize chemical reaction with RL and Zheng et al. (2018) train a model to provide personalized online news. Generally, most real-life applications seem to come from the area of robotics (Duan et al., 2017; Gu et al., 2017; Kober et al., 2013), autonomous driving (Kiran et al., 2021), system control (Glavic, 2019), and related problems.

3.2.1 Reinforcement learning in thermal power plants

Relative to the research of RL in games or robotics, the research on RL in industrial applications is more limited since “generally, academic research groups have only limited access to real industrial data and applications” (Hein et al., 2018, p. 1). Nonetheless, thermal power plants present a great use case for the application of RL because lots of data are logged for regulatory and safety reasons, plant operations and the combustion dynamics tend to be very complex, thus overwhelming other automation systems such as PID controllers or MPC, and

efficiency gains can have a significant impact on the environment simply due to the sheer amount of power plants worldwide.

As an early example, Stephan et al. (2001) apply multi-agent RL to control a hard-coal power plant. Four agents, represented by neural function approximators, are used to manipulate 12 controls with an action space of $[-1, 0, 1]$. They show that the efficiency factor (inversely related to the air consumption) can be increased and NO_x-emissions reduced while satisfying several limits such as steam and gas temperatures. More recently, Cheng et al. (2018) use the deep Q-network algorithm (DQN) to optimize flue gas emissions, boiler efficiency, and coal consumption of a coal-fired boiler. The boiler is simulated using data from the distributed control system (DCS) and a combination of a long short-term memory (LSTM) module with a convolutional neural network. Four controls are manipulated with an action space of $[-1, 1]$. The results show a reduction in flue gas emissions of $\sim 25\%$, an increase in boiler efficiency of $\sim 0.2\%$, and a reduction in coal consumption of $\sim 5\%$ compared to the baseline. Zou et al. (2018) follow a similar approach to Cheng et al. (2018) to optimize the control of a thermal power plant with a combination of simulator and optimizer. However, they use a dense neural network architecture for their simulator and the deep deterministic policy gradient (DDPG) as their RL optimizer, showing that DDPG outperforms the genetic algorithm. Fu et al. (2020) focus on denitrification efficiency. They use LSTM to predict the efficiency of selective catalytic reduction and an asynchronous advantage actor-critic (A3C) model to optimize the control using six controls. The results show that the A3C model enables a denitrification efficiency above the target level in $>96\%$ of the time. In the most recent studies, Adams et al. (2021) use an advantage actor-critic (A2C) model to optimize a commercial circulating fluidized bed (CFB) power plant for a good balance between steam output and emission levels. The CFB power plant is simulated using deep neural networks. They show that electricity generation can be increased by $\sim 2\%$ while at the same time reducing emissions by $\sim 1.5\%$. Zhan et al. (2021) introduce DeepThermal, a data-driven system that increases combustion efficiency, reduces pollutant emission, and controls operational risks (e.g., safety constraints). To simulate the power plant, they propose several deep recurrent neural networks (RNN) for each stage of the combustion process. They develop a new RL algorithm called MORE (“Model-based Offline RL with Restrictive Exploration”), essentially an actor-critic approach with restrictive exploration during training. In contrast to the research mentioned before, Zhan et al. implement the solution in “4 coal-fired thermal power plants in China” (Zhan et al., 2021) after offline training, showing that DeepThermal can improve the combustion efficiency by up to $\sim 0.5\%$.

While impressive progress has been made primarily on coal-fired plants, coal-fired plants and WIPs differ in fuel and boiler design, affecting the problem’s complexity. WIPs usually use untreated, heterogeneous waste, and thus, the energetic value of this fuel varies a lot (Sasaki et al., 2020). In contrast, the energetic value of coal is stipulated in the procurement contract according to the needs of a specific coal-fired plant. Next, waste is burned straight on a grate firing system, whereas in coal-fired plants, the coal is first pulverized and then sprayed into the boiler to ensure consistent combustion. These differences mean that there is more uncertainty in the combustion process of a WIP. Thus, the control problem is significantly more complex since, for instance, varying energetic fuel value affects the temperature within the boiler, which needs to be adjusted to meet regulatory limits and optimal combustion levels.

3.2.2 Reinforcement learning in waste incineration plants

Research on applications of RL in WIPs is minimal: Only two research papers can be found from 2004 and 2010. However, research in this field is necessary since, in contrast to coal-fired plants, WIPs are here to stay for the following decades, and thus, optimized control using RL can provide significant upside in the long run.

Stephan et al. (2004) use action dependent heuristic dynamic programming (ADHDP) to control the airflow in a WIP to avoid high CO emissions. The controller is pre-trained to mimic the behavior of a human operator and then optimized online. They show that CO peaks can be forecasted and mitigated but not avoided. Steege et al. (2010) present a neural fitted Q-iteration (NFQ) controller that is trained offline using seven days of data from a PID controller and, after training, implemented in a WIP for real-life experiments. The results are promising, showing more stable operations and reduced emission levels in terms of CO and NO_x compared to the PID controller. Interestingly, the experiment without artificial data had to be canceled due to dangerously poor performance, highlighting that additional data about rare, unique states and the proper control action based on expert knowledge are needed to make this approach work.

While both papers show promising results, the shortfalls also highlight the need for further research: There is a lack of research using modern RL algorithms and fully data-driven models to control a WIP and optimize several key parameters.

3.3 Methodology

The typical goal in WIP operations is to optimize the combustion process such that the optimal amount of steam is produced while emissions (e.g., CO and CO₂) are kept at a minimal level. However, operational goals can be highly specific to a particular plant, depending on external requirements from the heat/electricity grid, composition or mix of waste, boiler design, or flue gas treatment, to name a few. In most cases, the human operator can influence the combustion process via specific controls, e.g., in the case of steam/O₂ optimization, by adjusting the airflows and the waste throughput.

		Environment	
		Real environment	Digital twin
Operator	Human	A (Data source)	B
	RL agent	C	D (Research setup)

Fig. 3.1: Operators and environments

This control process can be learned within an RL framework. Since training of an RL agent is neither possible nor desirable within the real environment (i.e., the actual WIP), offline learning must be applied within a simulation environment (see Fig. 3.1, option D). Thus, a digital twin is a prerequisite for this application of RL. A similar approach is used by Schubnel et al. (2020) in building facilities, by Cheng et al. (2018), Zou et al. (2018), and Fu et al. (2020)

in coal-fired boilers, by Zhuang et al. (2018) for the treatment of paper-making wastewater, and by Li et al. (2020) for the cooling of data centers.

The digital twin is built using the data collected at the WIP (option A) and a data-driven neural network model. One could also employ a physical model, but we refrain from doing so due to the lack of scaling potential of such a solution. To test the digital twin further than we did in our research, one could ask a human operator to operate the digital twin and give feedback on the “feel” of the model (option B). And finally, the trained RL agent could be implemented in the real environment to prove the performance (option C). However, options B and C are not part of this work and remain aspects for future research. The comprehensive implementation process is described in Fig. 3.2.

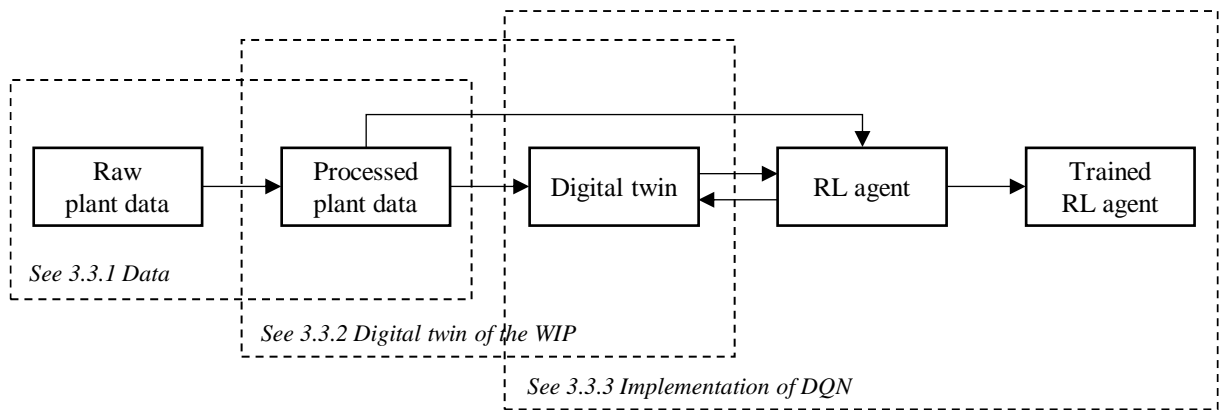


Fig. 3.2: Overarching design of the implementation approach

3.3.1 Data

This research is supported by Uniper Technologies GmbH, which has collected the data at a WIP in Germany. Data come from the plant’s DCS and is available from multiple days across three years, approximately 330 hours of operations in 30-second frequency, resulting in over 40,000 data points. In total, 32 variables are measured, clustered in three categories (similar to Cheng et al. (2018) and Zou et al. (2018)): 12 control variables, two performance variables, and 18 observation variables (see the Appendix B. 1 for further details). The control variables are used to influence the combustion process, more specifically, the performance variables steam output and oxygen levels: Steam output is crucial for the plant operator to produce electricity and heat. Too much steam, however, is bad since it negatively affects the machinery. The oxygen level indicates the “quality” of the combustion: As a general rule of thumb, the lower the oxygen level, the better the combustion and the lower the emissions.

No fuel or oxygen is left in perfect, stoichiometric combustion because their amounts exactly match. However, in an actual boiler, things are not perfect: for instance, local fires occur, fuel is not homogeneous, or the air is not moving in perfect planar currents. Having less oxygen than needed leads to emissions of CO. Therefore, typically, more oxygen is required than the optimal stoichiometric amount to compensate for all inhomogeneities. On the other hand, too much air in the system reduces the efficiency since the used air needs to be moved and pre-warmed. Thus, there is an optimal amount of oxygen in the system. The observation

variables help describe the state of the boiler (e.g., various pressure and temperature measurements).

The data of the human operator itself reveal some of the critical operational issues: (1) human operators only use a few of the control levers; (2) human operators act only infrequently, leading to a long time of inactivity followed by rather “big interventions”; (3) the performance is rarely optimal nor even close to it.

(1) Control levers: 12 control levers are available in the process. Table 3.1 shows how many times a control lever was adjusted on average per hour. Given the 30-second frequency of data collection, a maximum of 120 actions per hour can be logged. Three control levers are never used, and four more are rarely used (i.e., <5 actions per hour). Of the remaining five control levers, the secondary air flaps are not controlled manually but automatically by a conventional controller, explaining why they are changed more than 80 times per hour. Thus, the human operators mainly operate the power plant with only three levers (the two air flaps at the end of the grate (air flap 4 and air flap 5) and the load factor), which are adjusted on average 6-11 times per hour. Note that these three levers are “fast” levers, i.e., they directly impact the combustion process without long delays, which is why human operators prefer them.

Table 3.1: Analysis of control variables

Control variables	Average actions per hour (max. = 120)
Secondary air flap (front)	88
Secondary air flap (back)	80
Air flap 4	11
Air flap 5	8
Load factor	6
Swirl flap primary airflow 1	4
Swirl flap primary airflow 2	4
Air flap 3	2
Load regulator	1
O2 controller	-
Air flap 1	-
Air flap 2	-

In theory, the RL approach could help overcome this shortfall since all control levers could be given to the RL agent as possible actions. Unfortunately, this is not possible in our setup since the digital twin is built upon this data. Thus, these seven control levers cannot be adequately simulated because only too little (or no) data are available for them. Therefore, the RL agent will only be trained to use the three control levers that the human operators use the most (air flap 4, air flap 5, and the load factor).

(2) Infrequent actions: Table 3.1 already shows how limited human operators intervene in the process. To put it differently: On average, human operators change one of the control levers 36 times per hour, usually changing multiple controls simultaneously. Thus, they effectively only intervene 23 times per hour. Interestingly, these actions are not equally spread over time but happen in intervals. The human operator operates the plant for several minutes, followed by periods of inactivity (breaks between such intervention periods last, on average, 16 minutes). Fig. 3.3 shows this typical behavior. The fact that actions are performed rarely

entails that the actions tend to be rather significant interventions even though more frequent but smaller interventions would be preferred from a process perspective. Performing an optimal action would mean that the process is in an optimal state after the action has been propagated through the system. The more time passes after an optimal action has been performed, the more the system has potentially drifted away from the optimal state. Thus, continuous optimal actions are necessary to *maintain* an optimal state in the system.

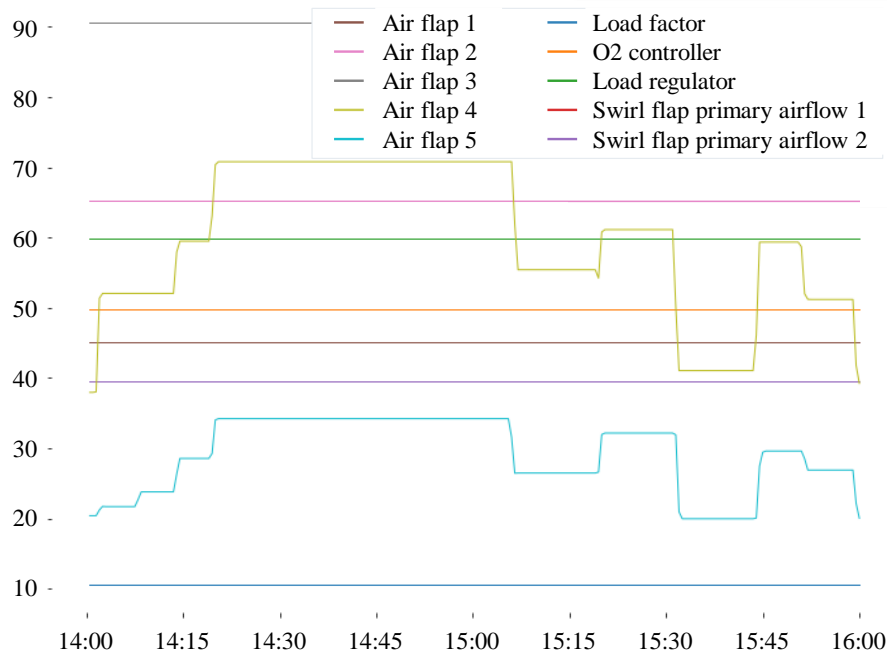


Fig. 3.3: Typical behavior of the human operator;
secondary air flaps are excluded since they are not manually controlled

(3) Sub-optimal performance: A closer look at the two performance variables (i.e., steam output [t/h] and oxygen level [%]) shows that the combustion is volatile and that the targets are rarely met (Fig. 3.4). The oxygen target ($\pm 5\%$) is only met during 3% of plant operations, while the steam target ($\pm 5\%$) is met during 25%. Meeting both targets simultaneously is even rarer (1% of plant operations), revealing a general problem: The operational goals for the performance variables do not necessarily fall within a process regime that is perfectly suitable for steady operation. Often, operational goals are formulated based on economic or political interests. In a second stage, the goals are then adjusted by process engineers to meet the technical requirements of the plant and the process. Naturally, the final conditions are not in perfect agreement with the plant's design, which typically leads to problems in operations.

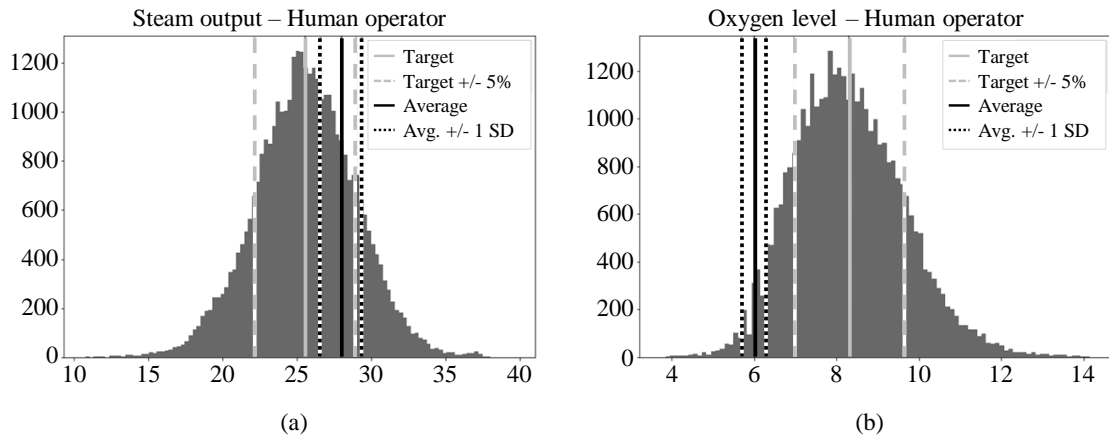


Fig. 3.4: Analysis of performance variables: (a) steam output and (b) oxygen level

Before the data are used for the digital twin or the RL agent, all pressure variables are filtered using 2-minute rolling averages to smoothen short-term fluctuations. For each control variable, an action variable is added, representing the delta between the respective control variable in t_0 vs. t_1 . This is needed since the RL agent outputs an action and not the target value of the control variable. Additional data features are engineered and added, such as averages over various periods in the past and past data points. The values are normalized using min-max-scaling to $[0, 1]$ to facilitate the training process, a commonly used technique in machine learning (Chakraborty et al., 2016). Since this increases the dimensionality of the input vector from $[1 \times 32]$ to $[1 \times 2156]$, the vector is reduced using principal component analysis (PCA) transformation to $[1 \times 253]$. This way, the relevant information in the data is kept while the overall amount of the data is reduced to facilitate the computation. After the data preparation, the entire data set consists of $\sim 36,000$ data points which are randomly split into training (70%), validation (10%), and test set (20%). The training set is used to train the models, the validation set to optimize the models, and the test set to provide the accuracy of the optimized, final model.

3.3.2 Digital twin of the WIP

A digital twin is needed to simulate the WIP since the real plant cannot be used for the training of the RL agent for safety reasons. This approach is prevalent in RL research on vital infrastructure such as power plants (Glavic, 2019), facilitating cheap and fast data collection.

In some areas, specialized digital twins are deployed when the relationships in the environments can be adequately modeled and information is fully available, e.g., virtual image synthesis and transformation for autonomy (VISTA) for autonomous driving (Amini et al., 2020), simulation of urban mobility (SUMO) for traffic simulation (Lopez et al., 2018), vinyl acetate monomer (VAM) plant model for chemical plants (Kubosawa et al., 2019) or Shadow Dexterous Hand for robotic manipulation (Andrychowicz et al., 2020), to name a few. Unfortunately, the combustion process of the WIP is non-linear. Thus it can only be approximated since many dynamics are too complex, and not all information can be collected (Z.-H. Zheng et al., 2020; L. Zhou et al., 2022). In such cases, neural networks have shown good performances as a basis for digital twins (Chakraborty et al., 2020; Cheng et al., 2018; Chow et al., 2002; Elhegazy et al., 2022; Fu et al., 2020; Jammeli et al., 2021; Kumar & Gururaj,

2019; Lv & Ren, 2019; Schubnel et al., 2020; Shin et al., 2020; Zhuang et al., 2018; Zou et al., 2018).

Our digital twin is designed in a way that it takes the current values of the variables (i.e., the “current state”) and an action (i.e., a change of the control variables) to simulate the following values of the variables (i.e., the “next state”) (Fig. 3.5). This reflects well what is happening during real-world operations: the human operator observes the data (i.e., the current state) and acts accordingly (i.e., action). The action influences the combustion process, and new data can be observed (i.e., next state). Since data are collected every 30 seconds, we also think in 30 seconds intervals for our digital twin. Notably, the digital twin only predicts the performance and observation variables since the control variables can be computed directly (control variable_{t₀} + action = control variable_{t₁}).




Fig. 3.5: Inputs for and outputs of the digital twin

For the training, the training set of ~25,000 data points is used and validated with ~3,500 data points using 5-fold cross-validation. Mean-squared error (MSE) is used as the loss function, and the training process is terminated as soon as the performance in terms of MSE of the validation set stalls (“early stopping”). To determine the best-performing model, a manual grid search tests several architectures of fully connected dense neural networks, the number of input variables, and some hyperparameters, e.g., learning rate or batch size.

We observe that complex architectures (in terms of the number of layers and neurons) outperform simple structures in terms of MSE. However, this comes at the cost of increasingly higher computation times and overfitting. Thus, a balanced network architecture with 3 hidden layers and 1500 neurons each is chosen. The final network architecture is described in the Appendix B. 2.

Regarding the number of input variables, i.e., history needed to predict the future state, discussions with the human operators indicate that it takes up to 30 minutes for an operator’s action to translate into visible changes since the system dynamics are relatively slow due to large spatial scales. Thus, various input vectors are tested, from using no past data to using the last 60 data points (i.e., 30 min of operations). On top of that, averages of different periods (e.g., last three data points, last ten data points) are added since it is known that momentum (as represented by the averages) is critical in the combustion process. Overall, >30 scenarios are tested, and the final digital twin is built using the last 60 data points and six averages.

Finally, other hyperparameters such as batch size and learning rate are tuned to determine the best-working setup. The detailed settings of the training of the digital twin can be found in the Appendix B. 3.




Fig. 3.6: Training and test loss for the best performing settings (see the Appendix B. 3 for details)

An accurate digital twin is quintessential for the RL agent to ensure that the policies learned within the digital twin can be transferred to the real environment. Fig. 3.6 shows that the digital twin predicts the data well (as measured by the MSE) but also generalizes well (as seen by the difference in the two loss curves). This best-performing model shows an MSE of $1.7e-5$ in the training and $7.6e-5$ in the test set.

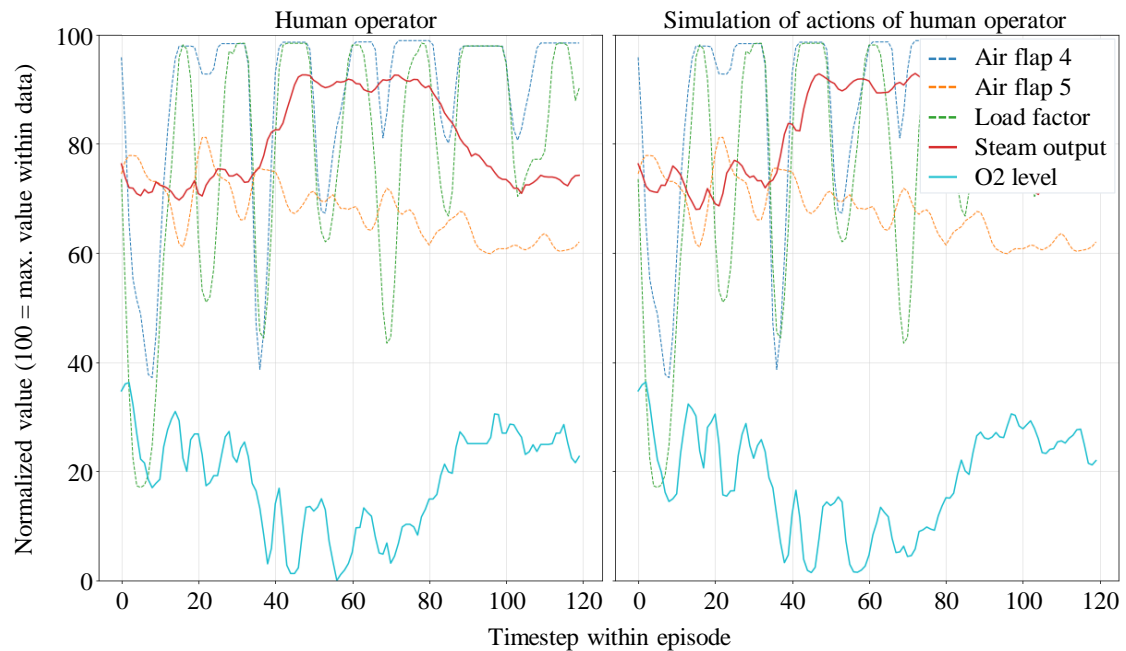


Fig. 3.7: Comparison of real and simulated oxygen and steam levels

For a final check, the model is put to the test by comparing real to simulated data (similar to Cheng et al. (2018) and Fu et al. (2020)). A random data point from the test set is used as the starting point of a simulation. Then, only the actions of the human operator in the following hour are copied and input into the digital twin. Fig. 3.7 shows that the simulated data are very

close to the original data, meaning that the digital twin does a good job processing the actions of the human operator.

3.3.3 Implementation of DQN

With the digital twin in place, the RL agent can be implemented. For this application, DQN (Mnih et al., 2015) is chosen since the action space is small and discrete, the state space continuous, and it is known to offer better sample efficiency (François-Lavet et al., 2018; Sutton & Barto, 2018). Apart from fitting with the properties of the problem at hand, DQN is also chosen because it is more widely researched (compared to other RL algorithms), and it is used in other automation problems (Cheng et al., 2018; Lai et al., 2021).

The RL algorithm enables the agent (i.e., the artificial operator of the WIP) to make sequential decisions to control the power plant, optimizing steam output while keeping emissions low. To do so, the RL agent receives a state and outputs an action that maximizes the reward function.

The state space represents what the RL agent “sees”. Similar to previous work (Cheng et al., 2018; Mnih et al., 2013), not only the current values of the 32 variables are fed into the RL agent, but also the last three points in time (so-called “stacking”). Other stacks sizes, e.g., 2 or 8, show worse performance. Thus, the RL agent works with a stack of four points in time (4×32), representing the last 1:30 min. The values are again normalized using min-max-scaling to $[0, 1]$ to facilitate the training process.

The action space describes what the RL agent can “do”. As described above, only three control levers are used. In real life, the action space for each control variable is continuous since each lever can be moved within a range of $[0, 100]$, 0 being fully closed and 100 fully opened. Yet, to keep the action space size as small as possible, the action space is discretized and limited to $[-1, 0, 1]$, similar to, e.g., Stephan et al. (2001). The three control variables are changed simultaneously, and thus, the action space has a size of $3^3 = 27$. The remaining nine control variables are kept at their respective levels from the beginning of the episode.

The goal of the RL agent is to maximize steam output while keeping emission levels low. The target value for the steam output is set at $S_T = 28$ t/h, and both less and more steam is bad since less steam means lower revenues, and more steam means a higher risk of exceeding the plant’s design limit. For that reason, it is also beneficial to reduce fluctuations because this allows for increasing the set point in a subsequent step. The oxygen level is used as a proxy for the emission levels since the oxygen level illustrates the combustion quality. The target value is $O_T = 6\%$. For both elements, the reward is the negative quadratic distance to the respective target value, and thus, the range of rewards is $[-\infty, 0]$. On top of that, several constraints are introduced as “punishments” to ensure that the RL agent only performs “legal” actions (similar to Schubnel et al. (2020)) since the control variables, i.e., air flaps (AF) and load factor (LF), are generally only defined within $AF_1 = AF_2 = LF = [0, 100]$. Punishments are introduced to actively discourage bad behaviors, whereas rewards encourage good ones (Sutton & Barto, 2018). Thus, if the air flap AF_1 is already at 100 and the RL agent suggests opening it further, this action (a) is not executed, and the RL agent is punished with $P = -100$. The overall reward function r is given by:

$$r = \begin{cases} -(S_T - S_{t_0})^2 - (O_T - O_{t_0})^2 & \text{if } 0 < (AF_1 + a_{AF1}), (AF_2 + a_{AF2}), (LF + a_{LF}) < 100 \\ -(S_T - S_{t_0})^2 - (O_T - O_{t_0})^2 - 1 * P & \text{if 1 control lever is outside } [0, 100] \\ -(S_T - S_{t_0})^2 - (O_T - O_{t_0})^2 - 2 * P & \text{if 2 control levers are outside } [0, 100] \\ -(S_T - S_{t_0})^2 - (O_T - O_{t_0})^2 - 3 * P & \text{if 3 control levers are outside } [0, 100] \end{cases}$$

An episode consists of a 60-minute time window of the WIP, and actions are executed every 30 seconds, totaling 120 actions (“steps”) per episode. After each action, the RL agent receives a reward. Thus, the optimal cumulative reward over the 60 minutes would be 0 if an episode started on an optimal level and the RL agent kept this optimal level during the whole episode, which is virtually impossible.

The training is conducted in Python using the Stable Baselines package (Hill et al., 2018). An episode starts with an initial state, a random data point from the training set. The three previous data points are added, and this stack is fed to the RL agent to make a step. An action is selected based on an epsilon-greedy policy, i.e., some actions are chosen based on the trained policy (“greedy”), and some are random. This action is paired with the current state (t_0) and its additional features and then fed into the digital twin, which simulates the next state (t_1). Then, the values of the control variables in t_1 are computed and added to create the full next state (t_1). The step rewards are calculated, and the tuple (current state, action, next state, reward) is stored in the replay buffer (Fig. 3.8). The next state serves as the current state for the next step. After 120 steps, the episode ends, and the total episode reward is computed. The next episode starts with another random starting point, and this process is repeated for as many episodes as required to train the RL agent.

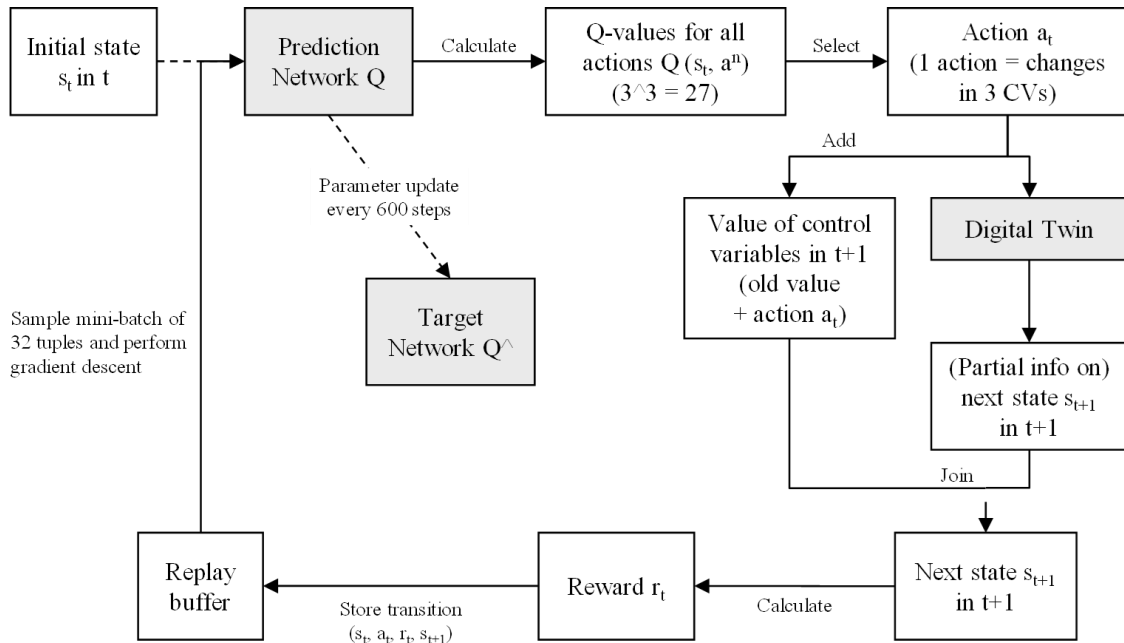


Fig. 3.8: Training process using DQN and the digital twin

Similar to the training of the neural network, the network structure and various hyperparameters of the DQN algorithm are tuned. For the DQN network, a dense neural network with 3 hidden layers, 1500 neurons each, is chosen, confirming Cheng et al.’s (2018)

observation that deeper networks show better performance. The best-performing set of hyperparameters (in terms of rewards) can be found in the Appendix B. 4.

3.4 Results and discussion

The performance of the RL agent can be determined based on the reward function. With the reward function described above, the rewards for each step can be calculated and accumulated to determine the episode rewards (1 episode equals 120 steps in 30-second increments).

Table 3.2: Mean episode results (for 100 episodes from the test set) after training a selected number of episodes

Training period (episodes)	Training period (steps)	Training time	Mean episode reward	Number of punishments
167	~20,000	1:58h	3,017	177
334	~40,000	3:54h	2,268	2
500	60,000	5:48h	2,313	65
667	~80,000	7:46h	2,175	67
833	~100,000	9:43h	2,666	226

The RL agent is trained for various numbers of episodes (Table 3.2). Three hundred thirty-four episodes are chosen as the optimal training period since, at this point, the episode rewards seem to have converged, the agent has learned not to go outside the set limits, and the risk of catastrophic forgetting (Kirkpatrick et al., 2017) due to “too much training” is limited.

To judge the RL agent’s performance, the RL agent’s rewards are compared to the human operator (Table 3.3). On average, the RL agent outperforms the human operator by -32%. Table 3.3 also shows that perfect performance (i.e., rewards ≈ 0) cannot be achieved since this would imply that first, an episode starts with steam and oxygen levels at their respective target levels, and second, this level can be maintained with no deviation for an extended period, which is unlikely due to the combustion processes. The best total reward for the RL agent and the human operator in the test set are -307 and -465, respectively.

Table 3.3: Mean episode rewards (for 100 episodes from the test set)

	Mean episode reward (human)	Mean episode reward (RL)	Difference
Count	100	100	
Mean	-3,325	-2,268	-32%
SD	2,288	1,484	-35%
Min.	-9,398	-5,758	-39%
Max.	-465	-307	-34%

To better understand the actual performance difference, it is helpful to look at the underlying values, i.e., the actual oxygen and steam levels. Using 100 random episodes from the test set, we see that the RL agent achieves a reduction of oxygen levels by 3.6% and an increase of steam levels by 7.4% (Table 3.4). Overall, the mean steam output is relatively close to the target level ($S_T = 28$). The mean oxygen level, while improved, is still further away (relatively and absolutely) from its target ($O_T = 6$). There might be several reasons for this

result: The oxygen target might be generally more ambitious, explaining why we rarely see performance close to the target level in the original data (Fig. 3.4). Alternatively, the design of the reward structure might have led the RL agent to optimize the steam first, given the, on average, bigger (squared) distance of steam levels from their target. Nonetheless, even with this potential focus on optimizing steam output, performance in both metrics is better.

Table 3.4: Mean oxygen and steam levels (for 100 episodes from the test set)

	Oxygen (human)	Oxygen (RL)	Oxygen difference	Steam (human)	Steam (RL)	Steam difference
Count	100	100		100	100	
Mean	8.26	7.96	-3.6%	25.35	27.24	7.4%
SD	0.93	1.09	17.1%	2.67	2.65	-0.5%
Min.	5.66	5.82	2.9%	20.24	22.08	9.1%
25%	7.54	7.13	-5.4%	23.57	25.23	7.0%
50%	8.19	7.80	-4.7%	25.34	27.10	6.9%
75%	8.76	8.93	2.0%	26.96	29.47	9.3%
Max.	11.28	10.10	-10.4%	33.91	33.36	-1.6%

The performance of the RL agent can also be analyzed qualitatively by comparing the actions of the human operator to the suggestions of the RL agent in a given episode. This implies that the actions of the human operator are reasonable in that specific situation. This approach is very tedious and time-consuming, yet it helps spark a conversation with the engineers on-site who are used to looking at charts of various variables instead of the reward function.

In Fig. 3.9, air flap 4 is analyzed within a 2-hour timeframe. In the first 1.5 hours, the RL agent suggests closing the air flap 4 whereas the human operator keeps it at a constant level. Unfortunately, we cannot tell from the RL agent's suggestion what absolute level it was expecting for the air flap 4, only that it should be lower. After 1.5h, we see that the human operator does close the air flap 4 significantly, and in turn, the RL agent changes the recommendation suggesting that it was closed too much. After around 3:05 hours, the oxygen level rises, and the steam output sinks. In turn, the RL agent suggests closing the air flap 4, which the human operator does several minutes later. Again, the RL agent perceives the level as too low, suggesting opening the air flap 4, and the human operator follows only a few minutes later. Overall, this episode indicates that the RL agent suggests the correct movements. But the comparison is challenging because the RL agent can only open or close the air flaps in increments of 1. Thus, it needs to plan well ahead to reach certain levels in time, whereas the human operator usually makes bigger changes but less frequently so.

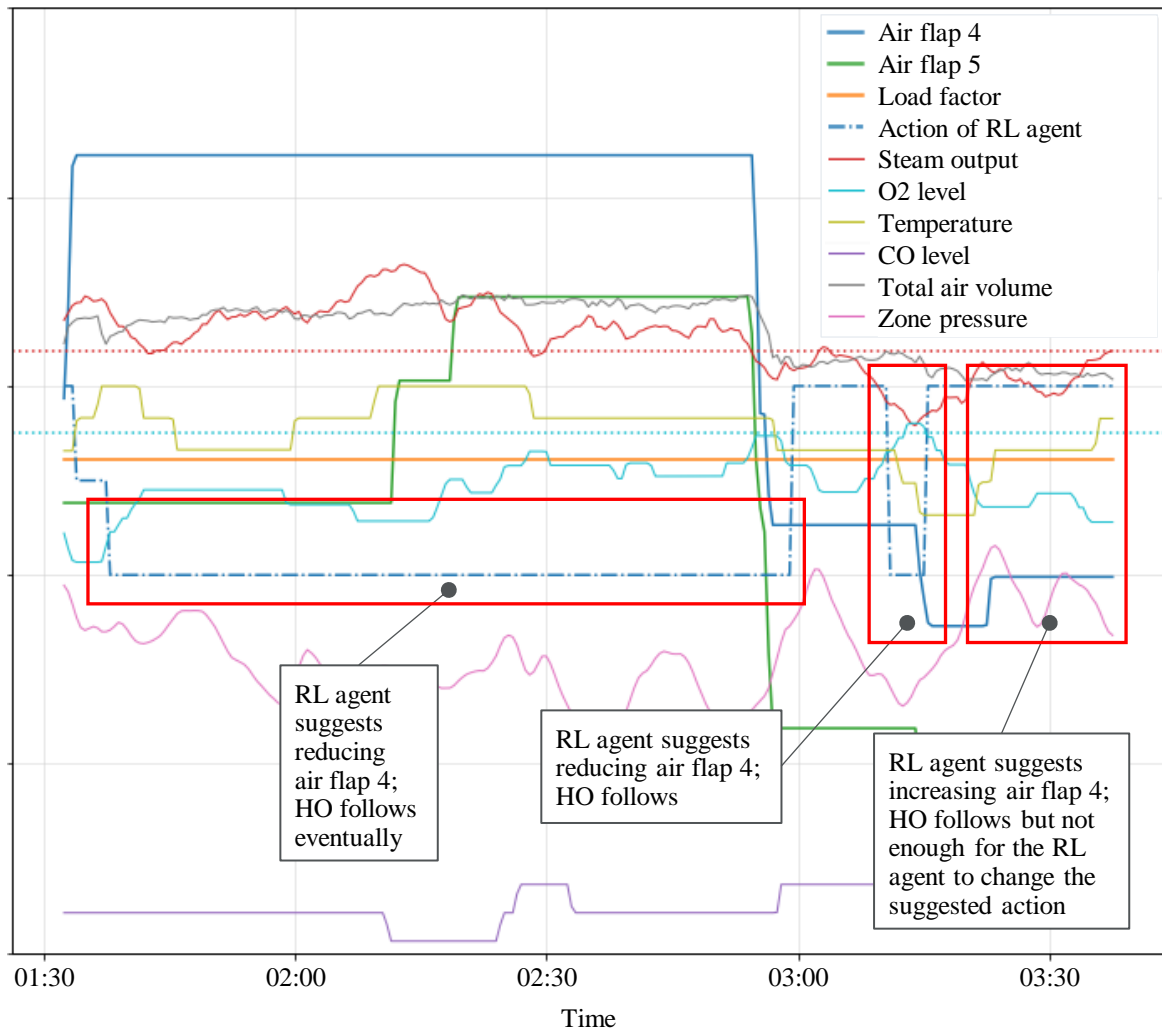


Fig. 3.9: Comparison of the original behavior of the human operator and the suggested behavior by the RL agent for 1 episode and 1 control lever (air flap 4)

While certainly, some challenges exist, these are very much in line with previous research in the field of RL applications in the real world (Dulac-Arnold et al., 2019, 2020).

For instance, the digital twin of the WIP is fully data-driven, and thus, it cannot perfectly capture the real-life combustion process, leading to a simulation bias. A bottom-up approach, e.g., based on fluid dynamics, physics, and chemical reactions, might alleviate this issue. Yet this would defy the purpose of a generalizable, data-driven modeling approach (Flynn, 2003). Moreover, we see in other fields that even after putting tremendous effort into a simulation environment such as CARLA, “a simulation gap still exists” (Amini et al., 2020, p. 1147). Nonetheless, the digital twin could be improved using more data (e.g., the composition and quality of the waste, which introduces a lot of uncertainty into the process) or different modeling approaches (e.g., “semi-physical simulation” (Schoukens & Ljung, 2019) or “hybrid learning methods” (Ota et al., 2021; Schubnel et al., 2020)).

Second, the performance benchmark might be unfair in two ways: First, a WIP is complex, and the reward function (as described above) might not capture everything the human operator is trying to balance. Whereas the human operators certainly follow target levels for steam production and emission levels, it remains unclear to what extent they try to achieve them given

that, e.g., no performance-based incentives are in place. Second, the performance comparison itself is not entirely fair since the rewards for the RL agent are derived from the digital twin (see Fig. 3.1, option D), whereas the rewards from the human operator come from the real machinery (option A). Even though the digital twin works quite well, for a better comparison, the RL agent would need to prove itself in a WIP (option C) or the human operator in the digital twin (option B) to avoid the simulation bias. Unfortunately, this could not be done given the current technical limitations of the WIP.

And finally, this approach is computation-heavy due to the iterative approach of RL in general. Specialized IT knowledge and hardware would be needed to speed up the training process. However, this applies to the training phase only, and when fully trained and implemented, the RL agent decides in a fraction of a second.

3.5 Conclusion

Waste incineration plants are highly complex yet lack state-of-art control and optimization. Whereas expert systems exist, they tend to be very specialized, requiring accurate models of the process or expert knowledge (Stephan et al., 2001). Thus, only limited control systems are used in practice, leading to shortfalls in performance.

This research suggests significant upside potentials for WIPs due to the application of RL, benefitting the plant operator in financial terms and the environment in terms of reduced emission levels, which are two relevant factors for the continuing acceptance of WIPs (Kheybari et al., 2021). This approach might be extended to tackle further operational challenges in the future. For instance, the temperature of the boiler might be controlled to reduce cleaning costs caused by slagging and fouling or costs for oil to boost temperatures or efficiency in terms of overall airflow increased. Since more of the operations can be automated this way, the engineers can focus on different tasks at the plant where they can provide greater value. We are convinced that the upside potential can be even expanded when new RL algorithms are tested, the action space of the RL agent is extended to give a greater degree of freedom, and the reward function is finetuned to include other targets and balance the respective reward components.

Moreover, this research shows that a digital twin of a WIP can be built using only original plant data, which are readily available, and neural networks, where off-the-shelf models yield good results, as demonstrated in the present research. This approach reduces the burden for future research on WIPs since the digital twin of the underlying environment usually is the main deterrent. In the future, this approach can be improved as suggested in Chapter 3.4 *Results and discussion* or by including additional hardware to collect novel data, e.g., video footage (Kaneko et al., 2019) of the waste or the inside of the combustion chamber.

RL technology paired with digital twins can be applied beyond power plants. Ideal areas for applications are fields in which digital twins of the real environments (e.g., the machinery or processes) are readily available, i.e., industries with high degrees of digitization. There, RL can be applied for complex, non-linear control problems that are not well covered by today's automation and control mechanisms.

4 OVERCOMING BARRIERS TO AI ADOPTION: A CASE STUDY FROM THE WASTE INCINERATION INDUSTRY

This chapter is based on Schlappa, Hegemann et al. (2022).⁸

4.1 Introduction

Artificial intelligence (AI) is finding its way into many industries, from aerospace and agriculture to hotels and the public sector (Bughin et al., 2017; Chui et al., 2018; Nam et al., 2021; Tao et al., 2019). Over 80% of large organizations have adopted some form of AI (Ghosh et al., 2019). A McKinsey report shows that 50% of the surveyed companies “have adopted AI in at least one business function” (Balakrishnan et al., 2020, p. 2). This rate of adoption seems natural since AI is creating significant benefits in different functions such as manufacturing or risk (Chui et al., 2018) and by different means like process automation, increased efficiencies, decreased failure rates, and new business models (Ammanath et al., 2020; Davenport, 2018).

The above-mentioned adoption rates are quite high and demonstrate the broad applicability of AI in general. “Nevertheless, the majority of companies still use AI tools as point solutions” (Ghosh et al., 2019, p. 1), and “only 8% of firms engage in core practices that support widespread adoption” (Fontaine et al., 2019, p. 2). In addition, high adoption rates do not reflect the success rates and the actual value AI generates in individual use cases. Thus, whereas the perceived adoption of AI is relatively high, implementations with a measurable impact are still rare, many use cases fail to scale, and many applications are as yet experimental (Benbya et al., 2020; Brock & von Wangenheim, 2019; Correani et al., 2020; Fontaine et al., 2019; Holmström, 2021).

We are convinced that one of the main reasons for this observation is the lack of AI readiness, a set of crucial properties for an organization to adopt AI successfully. Many companies jump into AI adoption (step 2) before doing the groundwork on AI readiness (step 1). However, organizations need to be ready for AI to adopt AI (AlSheibani et al., 2018; Issa et al., 2021; Jöhnk et al., 2021; Pumplun et al., 2019). The properties of AI readiness are more likely to be found in, e.g., big tech, whereas “full production implementation of AI technology is relatively scarce outside of the largest and most capable firms” (Davenport, 2021, p. 168). Most companies are not high-tech, especially in old economies such as the German one, and thus, they are not ready for AI (AlSheibani et al., 2018). This leads us to our research question:

Which strategies facilitate AI adoption in traditional, i.e., not-AI-ready industries?

To investigate our research question further, we set out to conduct an in-depth case study in an exemplary market that is not AI-ready but could and potentially should adopt AI technologies: the market for waste incineration plants (WIPs). WIPs are an excellent example since the 66 WIPs in Germany are, on average, over 30 years old and thus use old technology

⁸ This unpublished working paper with the title “Overcoming Barriers to AI Adoption: A Case Study from the Waste Incineration Industry” was written by Martin Schlappa, Jonas Hegemann, Tobias Mathur, and Stefan Spinler.

(Flamme et al., 2018; Weber et al., 2020); most employees are traditionally trained engineers; the IT hardware used is built for reliability rather than computing performance; the software used is tailormade and designed to be a critical safety net rather than a dynamic optimizer; regulations are often intended to disincentivize innovation. Similar circumstances can be found in other “old economy” industries, such as manufacturing.

We follow the qualitative case study approach as suggested by Jöhnk et al. (2021) and AlSheibani et al. (2020) to explore further which factors influence AI adoption. More specifically, we investigate Uniper’s efforts in implementing their AI solution “Operate” across the waste incineration industry and assess interactions and discussions with over 160 clients. Based on these interactions, we distill five strategies that counter typical AI readiness barriers, thus increasing AI readiness. With increased AI readiness, chances for AI adoption rise (AlSheibani et al., 2018; Jöhnk et al., 2021; Weiner, 2009).

This research is relevant for many reasons: We see few empirical studies exist on strategies for successful AI adoption, especially when faced with not AI-ready organizations. The “lack of research on the mechanisms of AI adoption constitutes a significant research gap” (Neumann et al., 2022, p. 2). As Jöhnk et al. point out: “[...] Further [exploration of] organizations’ specifics for AI adoption, for instance, through in-depth case studies [is needed]” (2021, p. 17). Our empirical results are needed to address the specific challenges of practitioners in the AI adoption process. We intend to go beyond theoretical considerations and provide practitioners with a set of strategies that are easy to use and do not require the transformation of the entire business. Making the entire organization entirely AI-ready is a meaningful goal. From a practitioner’s standpoint, however, it is essential to transform the business just as much as necessary to make it ready for the AI technology that is supposed to be implemented. While the results are derived from one specific industry (WIP), we believe they can be transferred to other industries well. We are convinced that practitioners will benefit greatly from this case study as it guides future AI implementation and rollout efforts.

4.2 Literature review

We first investigate the current body of literature relevant to our research. This paper builds upon research from two related fields: organizational AI readiness and adoption. These two fields are, in turn, grounded in other theories such as the technology acceptance model (TAM) (Davis, 1985), the technology-organization-environment (TOE) framework (Tornatzky et al., 1990), the diffusion of innovation theory (DOI) (Rogers, 1995) or the unified theory of acceptance and use of technology (UTAUT) (Venkatesh et al., 2003). First, we discuss (conceptual) frameworks and theories regarding AI readiness and adoption itself. Second, we conduct a systematic literature review (SLR) to distill articles putting these frameworks into practice and determining organizations’ AI readiness and adoption.

4.2.1 *Organizational AI readiness and adoption*

AI readiness and adoption can be seen as an area of the more general field of innovation or technology adoption, which has been extensively studied, be it on an individual or organizational level (e.g., Chen et al., 2021; Gfrerer et al., 2021; Neumann et al., 2022; Oliveira & Martins, 2011). Nonetheless, AI readiness and adoption require a particular focus since AI

has a much higher implementation complexity than conventional technologies. It involves intensive and continued efforts across various functions, sometimes with the support of external parties, and it entails substantial adjustments in other areas such as the company culture and data policies (AlSheibani et al., 2020; Chatterjee et al., 2021; Davenport & Ronanki, 2018; Issa et al., 2021; Jöhnk et al., 2021).

For organizations to adopt AI solutions, they need to be AI-ready (Jöhnk et al., 2021). Incidentally, there is no standard definition of AI readiness (Jöhnk et al., 2021). We follow Holmström's definition of AI readiness, being "an organization's abilities to deploy and use AI in ways that add value to the organization" (Holmström, 2021, p. 2). Research in this field is "still in its infancy" (Jöhnk et al., 2021, p. 8), and it is noted that AI readiness and adoption have not been studied exhaustively (Chatterjee et al., 2021; Jöhnk et al., 2021).

Most of the research focuses on developing AI readiness and adoption frameworks of various forms which can be used to assess organizations. Even though they differ in semantics, many are grounded in similar underlying theories (e.g., TOE, TAM, DOI), and they include similar precedents to AI readiness and adoption.

Jöhnk et al. (2021) argue that AI readiness is an essential precursor to AI adoption. They deduce 18 AI readiness factors along five categories (strategic alignment, resources, knowledge, culture, and data) from literature and interviews with 25 experts. These factors enable organizations to assess their readiness before adopting AI technologies and, thus, to proactively identify potential shortfalls within the organization, reducing the risk of failure regarding AI adoption. Moreover, their interviews reveal that AI readiness and adoption are not merely connected unilaterally, but rather that they are intertwined and "mutually reinforce (or restrict) one another" (Jöhnk et al., 2021, p. 16). Due to its clarity and comprehensiveness, we will use this framework for organizational AI readiness in our assessment (see Table 4.3).

Similarly, AlSheibani et al. "determine factors that influence an organization's readiness for AI adoption" (2018, p. 6). They analyze organizational AI adoption through the lenses of TOE and DOI, merging aspects of these theories to develop their framework for AI adoption at the firm level. Readiness is thus assessed on three levels and seven sub-levels: (1) technological (relative advantage and compatibility), (2) organizational (top management, organization size, and resources), and (3) environmental (competitive pressure and government regulatory issues). They test this framework empirically, showing that all factors but organization size and competitive pressure significantly affect AI adoption.

Issa et al. (2021) maintain that AI technologies have unique features, and thus, they entail different challenges, necessitating the development of a specific adoption framework for AI. To account for this, they adapt the perceived characteristics innovating theory and develop a comprehensive AI readiness and adoption framework. They use a mixed-method (survey and interviews) approach to show that three strategic components (mobility, interactivity, and autonomy) are key to AI readiness, which is, in turn, a mediator for AI adoption.

Baabdullah et al. (2021) examine the antecedents and consequences of successful acceptance of AI practices, specifically by business-to-business (B2B) SMEs in Saudi Arabia. They use the TOE framework to develop a conceptual model that relates, amongst others, AI readiness to acceptance of AI practices. Using structural equation modeling of survey data, they show that infrastructure and awareness (as part of AI readiness) are significantly correlated with

the acceptance of AI practices whereas technicality (i.e., the ease of use of AI technologies) is not.

Pumplun et al. (2019) recognize that the TOE framework can be applied to describe AI adoption. However, multiple expert interviews suggest that additional factors are needed due to the specific requirements of AI. Thus, they complement the standard TOE framework by, e.g., customer readiness and industry requirements. Moreover, they shed light on the different views on AI readiness and adoption of AI technology users and providers. They argue that transparency over the (lack of) AI readiness is needed to enable successful AI implementation.

Finally, Nortje & Grobbelaar (2020) argue that while AI can potentially transform businesses, many businesses face challenges in terms of implementation (i.e., adoption) of AI. Thus, they develop a framework to assess AI readiness, supporting the implementation of AI in businesses. Based on a grounded theory methodology and a conceptual framework analysis, they conduct an SLR to deduct their framework, which consists of the following dimensions: (1) employee and culture, (2) technology management, (3) organizational governance and leadership, (4) strategy, (5) infrastructure, (6) knowledge and information, and (7) security. Companies may use this framework to calculate their readiness and prioritize efforts to improve their AI readiness.

4.2.2 *Case studies on AI readiness and adoption*

While the body of research on organizational AI readiness and adoption is not immense (see above), all available literature suggests that some organizational AI readiness is needed to adopt AI technologies successfully. This readiness, in turn, can be split into multiple categories, depending on the underlying theory. Whereas some of the frameworks are validated empirically through interviews and small surveys, these concepts seem to have rarely been put into practice at scale.

Thus, we use an SLR proposed by Denyer & Tranfield (2009) and Tranfield et al. (2003) to compile all articles assessing AI readiness and adoption of organizations empirically. Here, we are intentionally not interested in conceptual articles (such as the ones mentioned above). Using the search string “TITLE-ABS-KEY ((“AI read*” OR “AI adopt*” OR “artificial intelligence read*” OR “artificial intelligence adopt*”) AND (“org*”) AND (“empiric*” OR “case stud*”))” we only find 27 records in the Scopus database (<https://www.scopus.com>) as of June 2022. Scopus has been used for similar SLRs (see, e.g., Kabalisa & Altmann, 2021). We purposefully do not restrict the initial study selection with additional keywords or time horizons, given that the sample size is already small. Out of these 27 records, only seven records are relevant based on the abstracts and other hygiene factors such as the language of the article (English only). This list is expanded by seven records through cross-references, out of which only three turn out to be relevant based on the abstract. The final ten articles are read in full, and another two articles are removed since, upon further inspection, they are not relevant to our research question. Only eight articles put the concepts of AI readiness and adoption in practice, i.e., measuring an organization’s AI readiness or adoption (Table 4.1).

Table 4.1: Overview of articles assessing organizational AI readiness and adoption empirically

Authors	Focus	Methodology	Geography	Industry	Status of AI readiness/adoption	Main challenges	Strategies to overcome main challenges
Bettoni et al. (2021)	AI readiness & adoption	Survey	Europe	Mixed	Unprepared to adopt AI	AI expertise; Data strategy; Organizational structure; Organization's culture	<i>not provided</i>
Chatterjee et al. (2021)	AI readiness & adoption	Survey	India	Manufacturing, production	Preliminary stage	Lack of perceived usefulness and ease of use of AI	AI adoption strategy; Collaboration; Communication; Customer centricity; Education; High-quality data; Top management support
Ellefsen et al. (2019)	AI adoption	Survey	Norway, Poland	Logistics, production	AI novice	AI expertise; Financial constraints	<i>not provided</i>
Holmström (2021)	AI readiness	Case study	<i>Not provided</i>	Insurance	Low to moderate	AI expertise; Confidence to work with AI technologies	AI tools and frameworks; Collaboration; Education; Mindset shifts
Nam et al. (2021)	AI readiness & adoption	Case study	Dubai	Hotels	AI-ready	AI expertise; Implementation; Integration of AI with existing technologies	Collaboration
Neumann et al. (2022)	AI adoption	Case study	Switzerland	Public sector	Mixed	AI expertise	Collaboration

(Continued on next page)

Table 4.1: Overview of articles assessing organizational AI readiness and adoption empirically (*continued*)

Authors	Focus	Methodology	Geography	Industry	Status of AI readiness/adoption	Main challenges	Strategies to overcome main challenges
Tao et al. (2019)	AI adoption	Case study	China	Public sector	Still in its infancy	AI expertise; Bureaucracy; Data sharing and security; Ethics; Legislation	<i>not provided</i>
Vuong et al. (2019)	AI readiness	Literature review	Vietnam	Healthcare	Lack of AI readiness	Data management and infrastructure; Financial constraints; Organizational structure	<i>not provided</i>
<i>This study</i>	AI readiness & adoption	Case study	Europe	Power plants	Low	AI expertise; Implementation	Education; Trust; Customer centricity; Focus; Collaboration

The SLR confirms that only limited research exists on applying AI readiness or adoption frameworks. Out of these, only a few authors follow the case study methodology. We note that across various geographies and industries, organizations indeed have low levels of AI readiness or AI adoption, which was also suggested by AlSheibani et al. (2018). Only Nam et al. (2021) provide a case with high AI readiness. Arguably, this can be ascribed to the fact that they study high-class hotels in Dubai, a relatively advanced and modern industry and city. We can also see that all authors use different terminology: Some authors define different levels of AI readiness and adoption and classify organizations accordingly (Ellefsen et al., 2019; Holmström, 2021; Neumann et al., 2022), whereas most authors use plain words to describe the level of AI readiness and adoption (or lack thereof). No standards exist.

In terms of challenges, we see recurring themes (see Fig. 4.1). For instance, most authors (6/8) highlight “AI expertise” as one of the main challenges to AI adoption. As Tao et al. point out: “[...] The department finds it challenging to fully take on the role of managing and supervising AI application, even at its early stage” (2019, p. 11). Because of this lack of AI expertise, organizations seek external support, further weakening their internal capabilities.

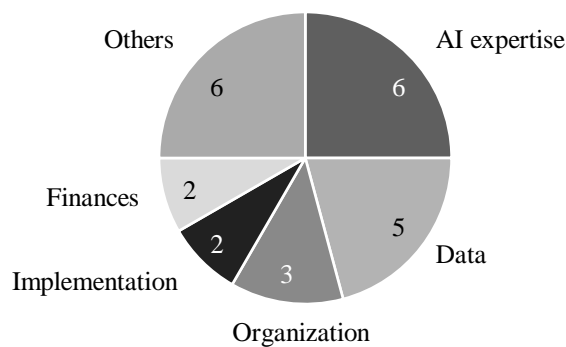


Fig. 4.1: Main challenges for AI readiness and adoption (frequency of mentioning)

Issues around “data” are also apparent. Here, organizations not only struggle with their data strategy but also with their data infrastructure and management. Because of this, they fail to exploit or even introduce AI technologies since the data needed are not readily available: “Without data, AI cannot be implemented” (Bettoni et al., 2021, p. 705). Moreover, “difficulties in data sharing and security measurement restrict the effective use of data” (Tao et al., 2019, p. 11).

Interestingly, only half of the articles offer strategies to overcome these challenges. They overwhelmingly suggest “collaboration” as a strategy, mainly because this directly impacts the main challenge: the lack of internal AI expertise. “Due to the lack of competency, [...] advanced AI-based solutions [...] need to be produced through close collaborations” (Nam et al., 2021, p. 571). Initially, external AI experts can be used to introduce AI technologies (Neumann et al., 2022). Over time, these external AI experts can build internal AI capabilities, facilitating further AI adoption (Chatterjee et al., 2021). Moreover, Holmström (2021) and Chatterjee et al. (2021) mention the education of the workforce as another critical success factor: “[...] Managers should try to appropriately train the employees of the organization to enhance their expertise

and skill, [...] facilitating the organizational authority to adopt AI in an easier way” (Chatterjee et al., 2021, p. 10).

The SLR confirms that further empirical research is needed on AI readiness and adoption (AlSheibani et al., 2020; Jöhnk et al., 2021). It strengthens our efforts to close this research gap by conducting a case study.

4.3 Methodology

This research uses a case study (see, e.g., Eisenhardt, 1989; Yin, 2018) to close the previous research gaps, following Eisenhardt’s (1989) eight-step process for case study research (see Fig. 4.2). This approach was chosen to get insights into a complex phenomenon with qualitative data, which might not be revealed by other types of data (Eisenhardt & Graebner, 2007).

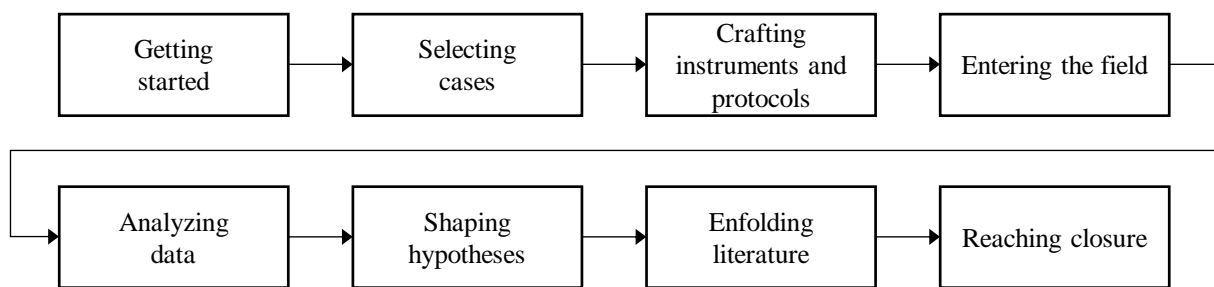


Fig. 4.2: Process for case study research (Eisenhardt, 1989)

4.3.1 Getting started

“An initial definition of the research question, in at least broad terms, is important in building theory from case studies” (Eisenhardt, 1989, p. 536). As described in our introduction, we want to investigate *which strategies facilitate AI adoption in traditional, i.e., not AI-ready industries*. We believe that many industries struggle to adopt AI technologies because they are not AI-ready in the first place. Thus, we want to understand what they are struggling with and which strategies can help with these challenges. Specifically, we are looking for strategies that make organizations just *AI ready enough* to adopt AI technologies since, from our experience, that is what practitioners need the most. This case study is explorative, with the intention to potentially generate new theories while at the same time keeping in mind the existent body of literature on AI readiness and adoption. We hope that our learnings from one industry can be used to generalize strategies that can be applied in other industries.

4.3.2 Selecting cases

To study our research question, we collaborate with Uniper and its Operait product. We do not follow the approach of statistical sampling but rather a theoretical sampling since we intentionally want to choose cases that are rich in information to gain insights and extend the theory (Eisenhardt, 1989).

Operait is a state-of-the-art AI tool to optimize WIPs in many ways, be it reducing emissions, increasing outputs, stabilizing operations, or other operational challenges. Operait is a new product, and the team essentially operates as a corporate start-up pioneering the field

of AI-driven optimization in power plants. The small team of approximately ten employees covers the whole business process from customer interviews, use case identification and assessment, design of custom AI solutions, commercial contracting, technical optimization, and after-sales service. The product itself is made of a standardized core that can be quickly adapted to various use cases. In a nutshell, it is a controller implementing a neural network that takes sensor measurements as inputs and generates control signals as outputs. In principle, it serves as a digital twin of the human operator, who originally controls the process. Typically, many projects are executed in parallel, such that synergies and cross learnings can be generated. Though advanced process control technologies have been available for decades, adoption is still not ubiquitous, and existing solutions are often expensive, mainly due to time-consuming detail engineering. Here, AI solutions can shine due to their data-driven nature, enabling efficient and standardized project execution as one key to a profitable business model. Since Operate is developing and implementing simultaneously, we can monitor the process in real-time and gather customer feedback, the challenges, and the approaches to overcome these from the sidelines, similar to Tao et al. (2019).

Operate's clients are typically waste and biomass incineration companies, mainly in Germany and other European countries. Both waste and biomass are considered "green" fuels for power plants since the landfill of waste causes substantial environmental damage. WIPs, in particular, have a long history and have been subsidized in the last decades, such that many municipalities have built their own plants because they offer a two-sided benefit: (1) they eliminate most of the waste that is fed into it, and (2) they produce energy in terms of electricity or heat. On the other hand, large energy companies also operate fleets of waste or biomass incineration plants. These companies often have more expertise and are keener on optimizing their processes. These power plants are usually automated with conventional automation mechanisms from the field of mechanical engineering, such as proportional–integral–derivative (PID) controllers, fuzzy controllers, or model predictive control (MPC). AI technologies, however, are usually not present on site.

The waste incineration industry is representative of other traditional companies, and as a sector, it is big enough to conduct good research with a large sample size. It has a meaningful impact on the economy overall and is vital for future sustainability goals. Thus, this case study offers the ideal, real-world example for assessing our research question.

4.3.3 *Crafting instruments and protocols*

The discussions between Uniper employees and the customers are regular business negotiations and conversations. Thus, they are confidential and should not be disturbed by researchers in the room. The researchers do not interfere in any way before or after the interactions. Thus, we rely for this research on qualitative data from various sources: discussions with Uniper employees, information from ERP systems, customer feedback, emails, and other archives, where available. These are typical qualitative data sources (Eisenhardt, 1989; Rashid et al., 2019), and we follow a qualitative data-only approach (Yin, 2018). In some of these client interactions, multiple people are present, increasing the validity of the notes and findings (Eisenhardt, 1989). The findings are synthesized by the researchers in several workshops with relevant Uniper employees. Such a collaboration with "within the

company” employees is not common, but it enables access to confidential and proprietary data that otherwise cannot be accessed (Correani et al., 2020).

4.3.4 Entering the field & Analyzing data

As Glaser and Strauss (1999) suggest, we combine the two phases *Entering the field* and *Analyzing data* since the data collection overlaps with the analysis of the notes. Data are collected for two years, from May 2020 to May 2022. No specific formatting is required for documentation to make the process as easy as possible for the Uniper employees. Instead, available data are analyzed and discussed with the respective employees.

Over 160 clients are approached (see Table 4.2). For confidentiality reasons, no client names can be shared. The majority of the clients operate WIPs, and most of the clients are in Germany. We estimate that this research covers more than 75% of the WIP market in Germany (Flamme et al., 2018). Beyond that, other European countries are included in this research and some other plant types, such as biomass incineration plants.

Table 4.2: Split of customers by plant type and country

	Frequency	Percentage
Plant type		
Waste incineration	105	65
Biomass incineration	36	22
Others	21	13
Total	162	100
Country		
Germany	85	52
United Kingdom	28	17
Netherlands	12	7
Belgium	7	4
Finland	7	4
Austria	5	3
Switzerland	5	3
Others	13	8
Total	162	100

This relatively big sample size goes beyond what is typically suggested for case study research (Eisenhardt, 1989; Nam et al., 2021), but we intended to cover as much of the industry as possible. We believe that different from some more targeted case studies with smaller sample sizes (e.g., Chatterjee et al., 2021; Ellefsen et al., 2019; Holmström, 2021; Nam et al., 2021), each sample in our case study does not provide that much detail since only second-hand data are available that was not intentionally gathered for research but rather for business purposes.

On Uniper’s side, the employees interacting with the customers were mainly engineers or product owners. On the customer side, primarily people in charge of the technology adoption, i.e., plant operators or engineers, people involved in day-to-day operations, or, if available, people familiar with AI or automation systems, were involved.

4.3.5 *Shaping hypotheses & Enfolding literature*

Please refer to Chapter 4.4 *Results* for the description of the hypotheses and the comparison with existing literature. Here, we also put the findings into a broader context and compare them with other cases (Macurová & Jurásková, 2013).

4.3.6 *Reaching closure*

Eisenhardt (1989) and Yin (2018) propose that data should be collected until no new information emerges from the data. We reach closure after two years of data collection, yielding over 160 samples. We decided to stop here for many reasons: First, the initial agreement between the researchers and Uniper was to collect data for two years unless circumstances changed. Second, the composition of samples by plant type and country was well-balanced and representative of the market at this point. Third, we were at a point where the sources exhibited consistency and good quality regarding challenges and strategies to overcome these. And fourth, Uniper's Operait product has developed significantly over the collection period, and the team transitioned from an experimentation and business development phase to an implementation phase with their customers. Thus, the number of new customer outreaches decreased since the focus shifted to different business aspects.

4.4 **Results**

During the consultation and implementation of the Operait AI solution, barriers were encountered on multiple levels: general skepticism of various stakeholders towards AI, especially being a “black-box” approach; limited access to data due to restrictive systems that are not optimized for extensive data exports; compatibility issues due to non-standardized DCS; problems in the daily operation that distract plant engineers from optimization tasks; and bad prior experiences with other advanced control methods like fuzzy control. These barriers align with previous research (e.g., Chui et al., 2018; Chui & Malhotra, 2018).

The most critical goal in a typical power plant is to keep it running since any shutdown due to technical problems or maintenance is way more expensive than any marginal performance optimization. This key goal determines the organizational structure, the mindset, as well as the daily operation. Taking risks in operation, e.g., by going very close to the design limit of the power plant, is thus not common. Keeping the plant running also involves ordering and organizing the delivery of the waste or biomass, catching up with the latest regulatory constraints, updating hard- and software, solving mechanical issues, planning annual revisions, selling the produced energy, organizing the side processes such as recycling, depositing the burnt waste, and many more activities.

With all these activities, core employees often find themselves lost in daily operations. They are under pressure to find ad-hoc solutions to keep the plant running and maintain a steady process as far as possible. Thus, for this group of employees, it is difficult to get a holistic view of the plant operations and to understand global optimization goals that often originate at the management level. In addition, since many people have been working in the plant for decades, different mindsets and strategies evolve in different shifts, leading to heterogeneous outcomes in different shifts.

Nonetheless, many power plants are still willing to invest in process optimization and adopt the latest AI technologies, but due to their limited expertise and resources, they need to rely on this expertise from external sources. These external partners then need to deal with the internal barriers to achieve a successful implementation.

Table 4.3: Assessment of AI readiness in WIPs using organizational AI readiness factors (Jöhnk et al., 2021)

Category	Factor	Assessment
Strategic alignment	AI-business potentials	High
	Customer AI readiness	<i>Not applicable</i>
	Top management support	Medium
	AI-process fit	High
	Data-driven decision making	Medium
Resources	Financial budget	Medium
	Personnel	Low
	IT infrastructure	Medium
Knowledge	AI awareness	Low
	Upskilling	Low
	AI ethics	Low
Culture	Innovativeness	Low
	Collaborative work	Low
	Change management	Low
Data	Data availability	Medium
	Data quality	High
	Data accessibility	Low
	Data flow	Low

In summary, we see that WIPs are not AI-ready (Table 4.3), like many other industries (see Table 4.1). While there is a strategic alignment and data are at least in part usable, resources are restricted, knowledge about AI is limited, and the company culture is not geared towards innovation. Nonetheless, this does not mean that AI cannot be adopted under no circumstances. Jöhnk et al. (2021) argue that each company has a specific target level for AI readiness (depending on, e.g., the industry or customers) and that AI readiness and adoption are “intertwined concepts” (Jöhnk et al., 2021, p. 15), i.e., that companies iteratively work on both. To tackle the barriers mentioned above and reach the right level of AI readiness, we develop a set of five strategies, targeting the factors with low AI readiness (Fig. 4.3). We think that our strategies work well across industries and geographies since these barriers are arguably hindering proper AI adoption in many traditional markets.

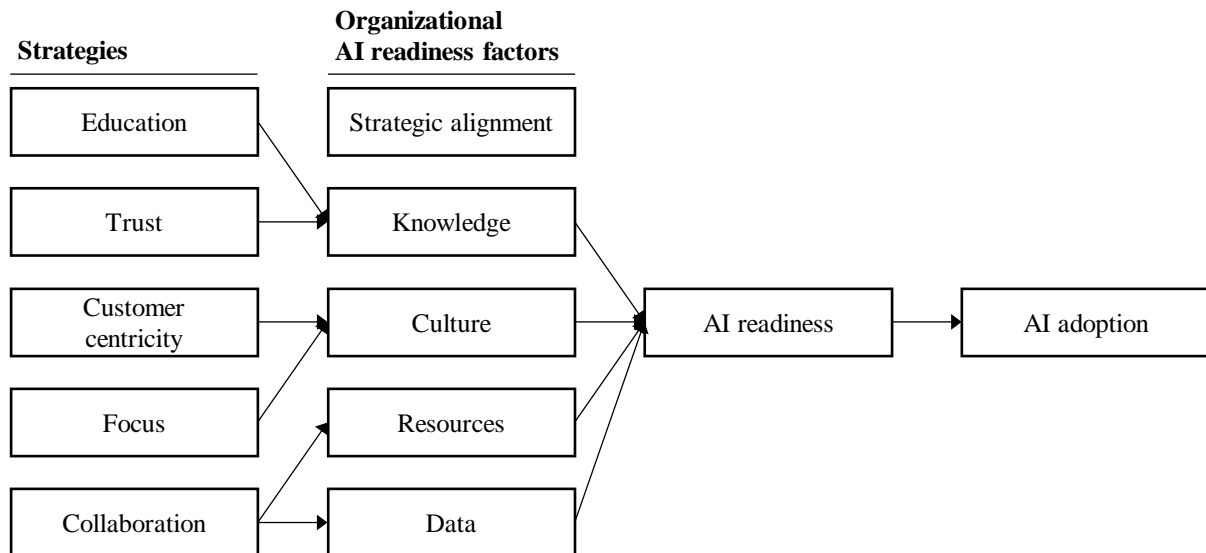


Fig. 4.3: Strategies to increase organizational AI readiness factors

(1) Educate stakeholders more about AI: Discussions with stakeholders show that AI is still seen as a black box that cannot be explained, leading to skepticism and aversion toward its implementation, especially in safety-critical environments (Hallinan & Striphas, 2016; Holmström, 2021; Pasquale, 2015). Thus, there is a desire for explainability that also affects the choice of technology: Instead of the newest and fanciest technology, a good compromise of performance, robustness, and explainability is needed (Baabdullah et al., 2021). The CTO of an energy utility company highlighted: *“For us, the human operators must always be on top of their game. Therefore, they must fully understand what is happening in the plant’s automation system.”* Education also entails realistic expectation management and a clear scope of the AI since some stakeholders might overestimate the effects of AI (Brock & von Wangenheim, 2019; Davenport & Abbey, 2018; Davenport & Westerman, 2018).

(2) Build trust between the stakeholders and the AI: AI may arouse fear in stakeholders (Benbya et al., 2020; Davenport & Abbey, 2018), and AI solutions tend to fundamentally change how, for instance, WIPs are operated. AI automation interacts directly with the plant’s control system; thus, a high level of trust is needed for full-scale implementation. This change may be softened by a staged implementation and continuous communication to achieve a smooth transition from conventional to AI control. To build this trust, a 3-step approach proved helpful (Fig. 4.4) that slowly transfers control from the human to the AI solution. After implementing the AI solution, a regular training and feedback loop with the operators in the control room has proven effective.

One operations manager of a WIP remarked: *“The AI Assistant phase really helped our operators to gain trust. In the beginning, most operators have been skeptical. However, after watching the AI for a couple of weeks, they now enjoy turning it on because they see how it helps them and that it is reliable.”*

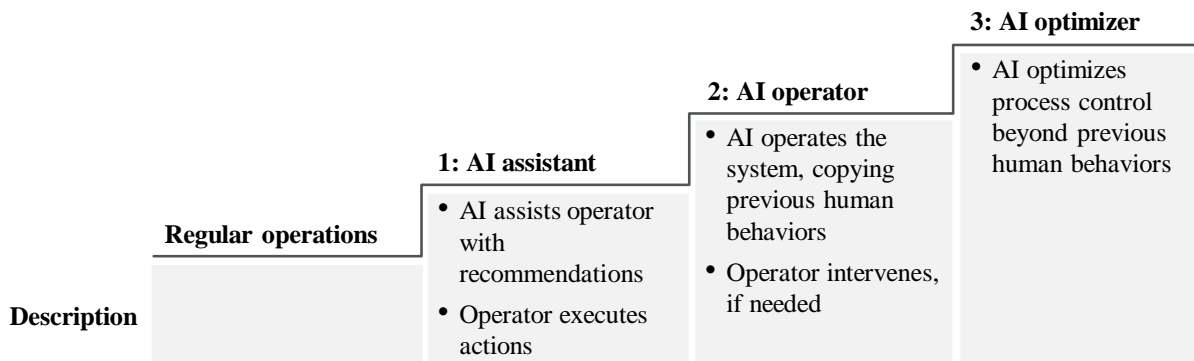


Fig. 4.4: Staged implementation of AI solution

(3) Put the customers at the center of development: AI product development is complex. Companies might shy away from close collaboration with stakeholders when the education need is high, and the focus is on quick AI implementation. However, we think that close cooperation with the customers leads to a better product-market fit (Sjödín et al., 2022). Thus, we suggest agile product development that facilitates proximity to the customers. For instance, Operaito engineers included plant operators in their development who were happy with this approach: *“The feedback was taken seriously. We were able to regularly influence the final solution”*. Early prototyping, together with customer engineers and operators, helps build a solution with the expected performance level and usability. Employing the staged implementation (Fig. 4.4), prototypes should also be implemented fast to “learn while doing” and get real process feedback, especially in the “AI assistant” phase, where the AI does not directly intervene. Here, it is essential to pick the right customers, particularly for pilots to test new technology or products (Brock & von Wangenheim, 2019; Wouters et al., 2018). Again, we emphasize that there are often quite heterogeneous opinions in different shifts of plant operators. It is imperative to gather and address the complete feedback from all the different streams to enable a higher degree of adoption in the next iteration cycle.

(4) Focus on the selected customers: Focusing on one industry might seem like the right level of focus, but our experience suggests that an even more narrow focus is needed, especially when faced with traditional industries (Brock & von Wangenheim, 2019). Even within one industry, customers tend to focus on different aspects or face different regulatory boundary conditions. For instance, one customer might ask for an optimization of the emission levels while another is struggling with safety shutdowns. Naturally, a wide variety of use cases emerges (Wouters et al., 2018). Thus, it is important to prioritize them for technical and commercial feasibility while verifying that they are “AI-only” use cases.

(5) Collaborate with hardware providers for implementation: AI adoption is usually not a software-only approach, and it also entails a hardware and system integration component (Benbya et al., 2020). In many cases, customers cannot provide easy access to data, and sensitive operating technology systems restrict them. On top of that, compatibility issues, for instance, due to non-standardized DCS across WIPs, limit the easy and fast implementation of AI solutions. This is especially troublesome since various DCS are used in a highly fragmented market (Deutsche Bank Research, 2017). Integration with the production/operating environment is one of the biggest obstacles (Davenport, 2021). Thus, hardware and interface providers need to be included in the development process from the start to facilitate the

implementation, and flexible, technical solutions must be found that are compatible with different setups. A plant maintenance manager mentioned: *“It is crucial for us that the AI system is integrated into our DCS system. We do not want to run systems from different OEMs. This creates interface and procurement issues.”* Collaboration is one of the essential factors for AI adoption (Dewi & Ahamat, 2018; Jöhnk et al., 2021; Sjödin et al., 2022). Whereas some customers depend entirely on original equipment manufacturers (OEM) to apply changes or extensions to their DCS, others have the knowledge to train employees on site and independently and quickly adopt newer (AI) technology. To increase implementation speed, these independent customers should be preferred, e.g., for pilots. DCS expertise on site is also preferable from an economic point of view since OEM service is usually costly, particularly for small power plants.

4.5 Discussion

We see that the waste incineration industry is not AI-ready (yet) for many reasons (see Table 4.3). Nonetheless, they can be made just AI-ready enough to enable AI adoption with the support of external partners (Neumann et al., 2022). We have distilled five best practices that can help overcome barriers when bringing technology to markets that are not AI-ready. They should be part of the rollout strategy and be considered when designing business processes.

Making a company fully AI-ready across all factors can be costly since organizations need to invest in the necessary initial resources and capabilities and develop adequate organizational structures and business processes (Ellefsen et al., 2019; Neumann et al., 2022). Thus, from a business perspective, making a company just “AI-ready enough” to enable AI adoption is essential, thus keeping investments in AI readiness low. This can also be done iteratively (Jöhnk et al., 2021). The five suggested strategies are all relatively cheap. Nonetheless, they can help companies save money along the AI adoption process since the most expensive resource during such a process is the time and commitment of the employees involved. AI adoption, in many cases, is not only a technical challenge but a social one. The employees need to know how to use the AI, be confident around it (Holmström, 2021), and see the benefits in the specific use case (Chatterjee et al., 2021). If the employees’ frustration during the implementation process is too high, it will be difficult to reverse these feelings; from our experience, such projects tend to die.

For education purposes, it is crucial to have a high-level and easy-to-understand pitch. To give a holistic view of the AI product, the pitch should (1) provide the context of the technical solution, i.e., compare to other technologies customers are familiar with, (2) emphasize the benefits of AI in this context, and (3) manage expectations and honestly say what AI cannot do. Based on this, it is easier to dive deeper into the technical properties of AI solutions and make the stakeholder understand why it can be valuable. Supporting material might be case studies of reference implementations. When the project is ongoing, additional steps can be taken to educate the employees, such as speaking in front of larger groups of employees (including ones not directly involved with the implementation or usage of the AI technology) to spread the word across the company.

To build trust between stakeholders and AI, we suggest a staged implementation. We start with an assistant mode that produces recommendations before the solution actively and

autonomously manipulates the system. During this phase, training sessions for the operators can support adoption and make them feel comfortable with the AI. Taking upcoming constructive feedback during these sessions seriously and using it for further optimization is beneficial for the implemented solution itself as well as for the customer because it makes them feel involved. Thus, the customers also put more trust in the solution. To make this happen, it is vital to listen carefully to the customer, beginning with the first meeting and understanding the specific pain points. Only a clear problem statement that everybody agreed upon can lead to fruitful discussions in later stages of the project.

Suppose, at some point during the project, it turns out that the customer has a specific problem that does not seem scalable, or there are external boundary conditions by the municipality, OEM, or management that cannot easily be solved. In that case, it should be seriously considered to stop the project. A good time to stop the project is after the first interviews and the data assessment when most of the information is available to judge the success rate. Having a list of critical questions that should be clarified early in the project can avoid major blockage in later stages. Companies need to keep in mind that AI cannot solve all kinds of problems and that it might not be the right solution to a specific issue at hand.

When external hardware providers or OEMs need to be involved in implementation, it might help to establish a separate and direct communication channel with them. Frequent alignment on specific projects will lead to more efficient project execution. Likewise, it will reduce the feeling of competition for the OEM because they will also benefit from more orders.

Even though many conversations with customers and employees were used to understand AI readiness and adoption in WIPs, the findings are primarily related to this industry. Arguably, it is peculiar since it is a safety-critical infrastructure that might have cultivated a distinct company culture. Beyond that, most plants are unique and tailor-made, thus hindering the scalability of general AI technologies.

Following the strategies suggested above does not imply successful AI adoption. It might merely help increase AI readiness. As pointed out earlier: With increased AI readiness, chances for AI adoption rise (AlSheibani et al., 2018; Jöhnk et al., 2021; Weiner, 2009). Individual organizations might face barriers in other areas that we do not impact. For instance, many companies mention issues with their data as a central challenge to successful AI adoption (Bettoni et al., 2021; Björkdahl, 2020; Chui & Malhotra, 2018; Vuong et al., 2019), for which other strategies are needed. In our case study with WIPs, data are usually readily available since it must be monitored for regulatory reasons and is crucial to provide control mechanisms and maintenance. Moreover, none of our strategies touch upon strategic alignment as one of the critical factors for AI readiness (Jöhnk et al., 2021) because we did not observe this as a challenge in our case study. Nonetheless, we recognize that this can be a significant barrier (Chui & Malhotra, 2018). AI readiness and adoption are complex matters that will require many targeted strategies from managers to implement AI technologies successfully.

4.6 Conclusion

We set out to find strategies that facilitate AI adoption in traditional, i.e., not-AI-ready industries. While there is much focus on the big tech companies (e.g., Google, Apple) and how they apply AI technologies in innovative ways, traditional, engineering-driven markets are

usually overlooked. Nonetheless, these markets still represent a significant portion of many economies with great potential for improvement. We used the waste incineration industry as an exemplary not-AI-ready industry and conducted a 2-year case study with over 160 client interactions. We see that many companies intend to adopt AI technologies without being AI-ready. The challenges they face align with the body of literature, even though empirical work such as case studies are generally rare. Companies must do much more than sell AI solutions to introduce AI solutions to such traditional markets. Thus, we propose five strategies to enable a sufficient level of AI readiness. Focusing from the start on the people and the eventual implementation beyond the product development is critical.

This study has several contributions: First, we summarize the literature on case studies related to organizational AI readiness and adoption. This is, to our knowledge, the first review of such kind, showing how under-researched this area is. Second, we determine the AI readiness of the waste incineration industry, representative of various other old, engineering-driven industries. Applying one of the AI readiness frameworks to an entire industry is unique. Third, we use the interactions between Uniper and its customers to develop five strategies that can help companies become just AI-ready enough for AI adoption. The idea of becoming just “good enough” differs from most research approaches (especially the conceptual ones), developing a long list of factors and strategies which quickly overwhelm companies. Our small set of strategies has proven to be successful.

In the future, we see multiple research needs. We believe that similarly large case studies must be conducted in other industries and with other AI technologies to validate the results. For instance, exploring industries with data challenges would be interesting since this might entail very different strategies. Moreover, we see a need for an overarching AI adoption framework. Even though some authors suggest AI readiness or AI adoption assessments (e.g., AlSheibani et al., 2018; Holmström, 2021; Jöhnk et al., 2021), these only help companies assess their current levels and potentially show gaps. We believe that companies need more than that, for instance, assessing AI readiness depending on AI technology and suggesting strategies to overcome barriers and close AI readiness and adoption gaps.

We are convinced that these strategies can be applied in various environments. As our literature review shows, many environments are not AI-ready; thus, the companies are partnering with external AI experts, like in our case study. We believe that more strategies must be shared to enable faster AI adoption. We see that many prototypes and pilots are tested in the market, and we perceive an appetite for AI solutions from traditional markets. However, repeated failures continue to strengthen skepticism and fear toward AI. Thus, the solutions proposed in this paper should help practitioners succeed in their AI endeavors in traditional markets.

5 CONCLUSION

Disruptions are and will be part of our lives. Because of this inevitability, we need strategies to avoid or handle them effectively. This dissertation looked empirically at two such ways: resilience and artificial intelligence (AI). We first assess existing resilience measures, develop our own resilience measure, and benchmark resilience levels in a large part of the German economy: the manufacturing industry (Chapter 2). We then use a state-of-the-art AI approach (reinforcement learning) to control and optimize a waste incineration plant (WIP), increasing performance as well as enabling more stable operations (Chapter 3). Finally, we examine challenges with AI adoption via an in-depth case study and suggest strategies to overcome these (Chapter 4).

5.1 Summary

In Chapter 2, we focus on the perceived organizational resilience of the German manufacturing industry. First, we conduct a systematic literature review (SLR) and realize that (1) no recent data for the German market exist, (2) there are only a few simple measures for resilience, and (3) the surveys seem impractical for practitioners. These realizations further strengthen our research motivation. Thus, we develop a resilience survey and a novel resilience measure, the Enterprise Resilience Index (ERI). We collect responses from ~200 professionals and investigate their company's past, current and anticipated resilience levels. The results show a medium resilience level ($ERI = 63.9$) across the manufacturing industry. However, we see differences between sub-industries: For instance, the Chemicals industry ($ERI = 73.8$) seems to be relatively more resilient than the Machinery ($ERI = 54.5$) and the Foods and Consumer Goods industries ($ERI = 61.4$). We also find that the companies perceive their current resilience levels to be below past (i.e., last year's) levels while at the same time gathering a strong anticipated need for (primarily external) resilience, arguably because the German manufacturing industry expects disruption from external sources. This need is especially pronounced in large enterprises with significantly higher levels than small and medium-sized enterprises. Moreover, we see strong correlations between internal and external perceived resilience levels across all time horizons.

As the sample is sufficiently large, we are convinced that the results are representative of the manufacturing industry. Nonetheless, we would like to increase the sample size further to better capture trends at the sub-industry level using more granular industry codes. The survey is well-suited to be used across industries and geographies, and our feedback suggests that the participants appreciated the shortness of the survey.

Our survey provides a first-of-its-kind overview of perceived organizational resilience across multiple companies and industries in Germany. Practitioners can use the tools (i.e., the survey and the ERI) to benchmark their organizational resilience against their peers', thus enabling all companies to become more resilient through transparency. We hope the results will spark conversations between resilience experts in various fields. Potentially, "resilience leaders" can share their best practice to enable "resilience laggards" to improve their resilience, which in turn could increase the resilience of an entire supply chain due to its dependencies. The differences in past, present and future organizational resilience highlight the need for

continuous resilience measuring. Thus, we believe that in the future, our survey should be used to continuously and regularly benchmark (parts of) the economy in terms of perceived organizational resilience. This will create a longitudinal view of resilience that might inspire further research.

In Chapter 3, we bring state-of-the-art AI technology, i.e., reinforcement learning (RL), to the traditional waste incineration industry. We use raw plant data from the distributed control system to build a digital twin of the waste incineration plant (WIP) since the actual equipment, i.e., the real WIP, cannot be used as the training environment. The data include key characteristics such as the control parameters (e.g., air flap position), boiler characteristics (e.g., temperature, pressure), and the performance parameters (steam output and oxygen levels). These parameters are chosen together with industry experts who use these to control the process in the day-to-day operations. The digital twin uses a neural network with 3 hidden layers and 1500 neurons each, and it is trained with a set of ~25,000 data points using 5-fold cross-validation. The best-performing model shows a mean-squared error (MSE) of $1.7e-5$ in the training and $7.6e-5$ in the test set. Within this offline environment, the RL agent is trained to control and optimize the combustion process of the WIP using multiple optimization criteria. We use the deep Q-network (DQN) algorithm for the training. The RL agent controls three control parameters (two air flaps and the load factor). The goal is to keep the steam output at a specific target level while keeping emission levels low. These goals (along with some other operational targets) are implemented in the reward function. The agent is trained for 334 60-minute-long episodes, each episode consisting of 120 steps in which the agent takes actions. In terms of rewards, the RL agent outperforms the human operator by 32% while simultaneously reducing the volatility of the process (measured by the standard deviation) by 35%. In terms of the performance parameters, we see that the RL agent achieves a reduction of oxygen levels by 3.6% and an increase of steam levels by 7.4%.

Since the approach is fully data-driven and with limited expert input, it can be transferred to other control processes, where original plant data are readily available. Mainly using off-the-shelf models, this approach reduces the burden for future research on WIPs since the digital twin of the underlying environment usually is the main deterrent. Nonetheless, RL is still best suited for applications in which digital twins of the real environments are available, i.e., industries with high degrees of digitization. RL can be applied for complex, non-linear control problems that today's automation and control mechanisms do not handle well.

This research suggests significant upside potentials for WIPs due to the application of RL, benefitting the plant operator financially, environmentally, and operationally: (1) it increases profits of the WIP by increasing the average steam output, which is converted to electricity and heat but also by increasing waste throughput, a second source of revenue of a WIP; (2) it decreases emissions in an environment where regulations are getting tougher; (3) it decreases the dependence on highly skilled and scarce human operators which in turn can be used for other value-adding tasks. This approach can be used to solve further typical issues at WIPs, e.g., staining of the boiler due to improper temperatures and high emission levels (“slagging and fouling”) or unplanned shutdowns. In a broader sense, this application can be seen as a general-purpose framework, showcasing the power of RL to control complex processes without the need for very specialized models of the process or expert knowledge. The fully data-driven

approach using many “off-the-shelf” tools facilitates the adoption in environments where data scientists may be lacking.

In Chapter 4, we look at the barriers to AI adoption. Specifically, we dive deeper into the waste incineration market to understand better which strategies facilitate AI adoption in traditional, i.e., not-AI-ready industries. To investigate our research question, we conduct an SLR on previous efforts to empirically study AI readiness and adoption. Since only limited research exists, we administer our own qualitative case study based on Uniper’s efforts in implementing their AI solution “Operate” across the waste incineration industry. We assess interactions and discussions with over 160 clients and synthesize five strategies that counter typical AI readiness barriers, thus increasing AI readiness: education, trust, customer centricity, focus, and collaboration. These strategies focus on transforming the businesses just as much as necessary to prepare them for the AI technology that is supposed to be implemented. With increased AI readiness, chances for AI adoption rise.

While the results are derived from one specific industry (WIP), we are convinced that these strategies can be applied in various environments. As our literature review shows, many environments are not AI-ready. Many industries, especially in developed countries such as Germany, are grounded on technologies from the last centuries, and thus, they have not fully adopted AI. We are convinced that practitioners will benefit greatly from this case study as it guides future AI implementation and rollout efforts.

This work contributes to the literature on AI readiness and adoption in many ways: it provides a comprehensive overview of the current state of the literature; it gives empirical evidence on AI readiness and adoption, applying an AI readiness framework to an entire industry; and it suggests a prioritized set of proven strategies to facilitate the mechanism of AI adoption whereas most research only provides general strategies for AI adoption. Managers benefit from this work since it goes beyond theoretical considerations and provides practitioners with a set of strategies that are easy to use and do not require the transformation of the complete business. Making the whole organization entirely AI-ready is a meaningful goal. From a practitioner’s standpoint, however, it is essential to transform the business just as much as necessary to make it ready for the AI technology that is supposed to be implemented.

In conclusion, the three chapters use various research methods to answer numerous research questions regarding organizational resilience and AI technologies. They give empirical evidence and expand the literature on organizational resilience and benchmarking, reinforcement learning and digital twins, and AI readiness and adoption.

5.2 Outlook

This dissertation makes multiple contributions to the academic literature and has many managerial implications. Nonetheless, we still see ways to take this research further and tackle some of its limitations.

The survey conducted in Chapter 2 is designed to be used as a longitudinal study. Our survey results show significant differences in perceived organizational resilience over two years. It would be interesting to see how the perceived organizational resilience develops and how effective managers are at assessing their future need for resilience or past resilience levels, which would become possible with another round of data collection. Additionally, this survey

could be extended systematically to more industries and countries since the literature suggests such differences. Finally, the concern with all ex-ante resilience measures is their validity. Thus, a real disruption is needed to see if the companies that perceived themselves as resilient fare better than the non-resilient companies.

For Chapter 3, we also see multiple routes for future extensions. The current setup with the DQN algorithm can be expanded step by step, for instance, by enlarging the action space to enable greater step sizes. Moreover, the reward function can be adapted to include more operational parameters that might be of second-order relevance to WIPs. To take the approach beyond DQN, different algorithms can be applied and benchmarked. The literature on RL makes fast advancements, and new algorithms are regularly developed and tested in game environments. It would be exciting to see their application with real process data. Next, we would like to see the implementation of our RL agent in the real environment to confirm its superior performance. Finally, RL should be brought to more control problems beyond WIPs, for instance, in other power plants with combustion processes.

The case study in Chapter 4 provides hands-on guidance for AI adoption. Since it was gathered across companies from one industry, we would like to see similar analyses for other industries that might have other challenges with their AI readiness, such as their data strategy. Moreover, implementing other AI technologies should be studied since AI technologies differ vastly in terms of resource needs and complexity. These research paths could be combined to create an overarching AI readiness and adoption framework that helps professionals understand their current level of AI readiness and the needed level based on certain types of technologies.

Finally, while we looked at organizational resilience and AI separately, we also see that AI technologies can enable resilience. Future research could be specifically devoted to directly testing the link between AI and resilience in companies, industries, or economies. AI might help solve some of the complex problems that bring uncertainty and disruptions to companies.

In conclusion, we see many exciting ways to extend this work, highlighting the importance of the presented topics and the need for further research. We believe that resilience and AI are more than just trendy topics: They are here to stay and will affect many areas of our lives.

A APPENDIX TO CHAPTER 2

A. 1: Survey questions, in original language *[answering options]*

1. Mein Unternehmen ist im letzten Jahr mit internen Störungen (Mitarbeiterfluktuationen, Maschinenausfälle, ...) gut umgegangen.
[stimme nicht zu; stimme eher nicht zu; neutral; stimme eher zu; stimme voll zu]
2. Mein Unternehmen ist im letzten Jahr mit externen Störungen (Lieferkettenstörungen, politische Konflikte, ...) gut umgegangen.
3. Mein Unternehmen ist auf interne Störungen (Mitarbeiterfluktuationen, Maschinenausfälle, ...) gut vorbereitet.
4. Mein Unternehmen ist auf externe Störungen (Lieferkettenstörungen, politische Konflikte, ...) gut vorbereitet.
5. Mein Unternehmen erwartet im nächsten Jahr erhebliche interne Störungen (Mitarbeiterfluktuationen, Maschinenausfälle, ...).
6. Mein Unternehmen erwartet im nächsten Jahr erhebliche externe Störungen (Lieferkettenstörungen, politische Konflikte, ...).
7. Welcher Industrie würden Sie Ihr Unternehmen zuordnen?
[Herstellung von Kraftwagen und Kraftwagenteilen; Herstellung von Metallernzeugnissen (inkl. Metallherzeugung und -bearbeitung); Herstellung von Datenverarbeitungsgeräten, elektronischen und optischen Erzeugnissen, elektrischen Ausrüstungen; Herstellung von pharmazeutischen Erzeugnissen; Herstellung von Nahrung, Getränken und Futtermitteln; Herstellung von Textilien, Bekleidung, Schuhen und Lederwaren; Herstellung von chemischen Erzeugnissen; Maschinenbau; Mineralölverarbeitung; Herstellung von Möbeln; Herstellung von Glas und Glaswaren, Keramik, Verarbeitung von Steinen und Erden; Herstellung von sonstigen Waren]
8. Wie hoch war der Umsatz Ihres Unternehmens im letzten Jahr?
[< 10 Mio. €; 10-50 Mio. €; 50-100 Mio. €; 100-500 Mio. €; 500 Mio. € - 1 Mrd. €; 1-5 Mrd. €; >5 Mrd. €]
9. Wie viele Beschäftigte hat Ihr Unternehmen?
[< 100; 100-250; 250-500; 500-1,000; 1,000-5,000; 5,000-10,000; > 10,000]

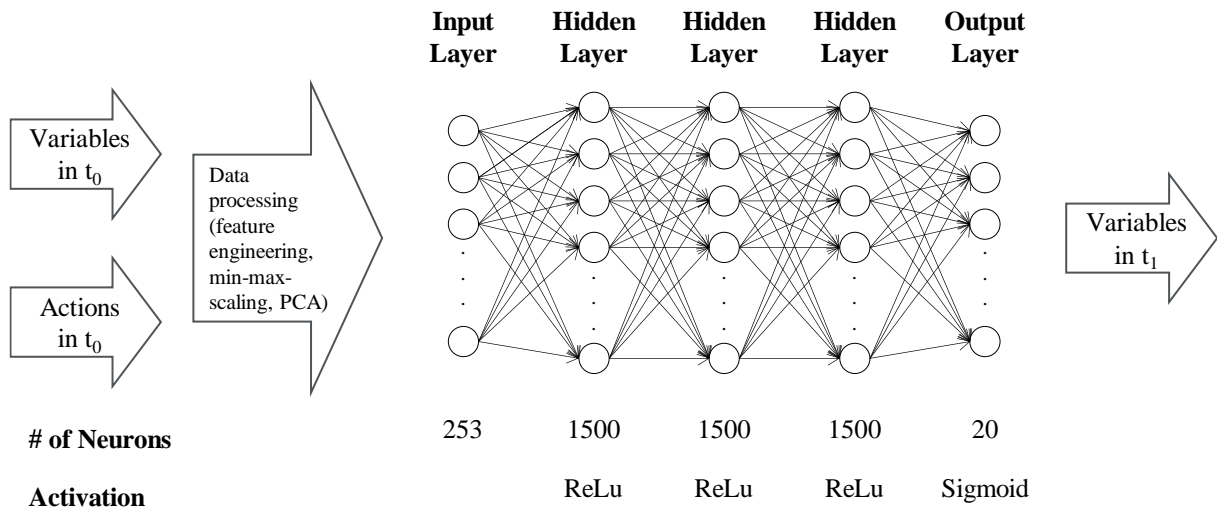
A. 2: Survey questions, translated [answering options]

1. My company has dealt well with internal disruptions (employee fluctuations, machine breakdowns, ...) in the past year.
[disagree strongly; disagree somewhat; neutral; agree somewhat; agree strongly]
2. My company has dealt well with external disruptions (supply chain disruptions, political conflicts, ...) in the past year.
3. My company is well prepared for internal disruptions (employee fluctuations, machine breakdowns, ...).
4. My company is well prepared for external disruptions (supply chain disruptions, political conflicts, ...).
5. My company expects significant internal disruptions (employee fluctuations, machine breakdowns, ...) in the next year.
6. My company expects significant external disruptions (supply chain disruptions, political conflicts, ...) in the next year.
7. To which industry would you classify your company?
[Manufacture of motor vehicles; manufacture of fabricated metal products (incl. basic metals and fabricated metal products); manufacture of computers, electronic and optical products, electrical equipment; manufacture of pharmaceuticals; manufacture of food products, beverages and tobacco; manufacture of textiles, apparel, footwear and leather products; manufacture of chemicals and chemical products; machinery and equipment; manufacture of refined petroleum products; manufacture of furniture; manufacture of glass and glass products, ceramics and related products; manufacture of other products].
8. What was the turnover of your company last year?
[< 10 m. €; 10-50 m. €; 50-100 m. €; 100-500 m. €; 500 m. € - 1 bn. €; 1-5 bn. €; >5 bn. €]
9. How many employees does your company have?
[< 100; 100-250; 250-500; 500-1,000; 1,000-5,000; 5,000-10,000; > 10,000]

B APPENDIX TO CHAPTER 3**B. 1: Variables in the data set and their ranges**

Variables	Mean	Min.	Max.	Unit
Control variables				
Air flap 1	45.2	1.3	47.3	%
Air flap 2	65.5	1.2	81.9	%
Air flap 3	83.6	28.9	100	%
Air flap 4	40.4	1.1	100	%
Air flap 5	20.5	0.9	100	%
Secondary air flap (front)	5,100	4,431	8,432	
Secondary air flap (back)	3,144	2,736	5,328	
Swirl flap primary airflow 1	38.1	14.9	61.1	%
Swirl flap primary airflow 2	38.1	14.9	61.1	%
Load factor	58.4	0	94.0	%
Load regulator	10.4	8.9	11.6	%
O2 controller	49.7	49.7	50.0	%
Performance variables				
O2 level	8.3	3.7	14.2	%
Steam output	25.6	10.8	39.6	t/h
Observation variables				
Primary air volume	40,025	26,788	50,139	m ³ /h
Secondary air volume (front)	5,114	0	8,659	m ³ /h
Secondary air volume (back)	3,151	0	5,422	m ³ /h
Primary air pressure	3.3	-0.4	6.6	mbar
Secondary air pressure	79.3	2.3	99.2	mbar
Drum pressure	41.9	39.4	45.0	bar
Dynamic pressure	-0.6	-1.7	1.1	mbar
Air pressure roller 1	-0.5	-2.0	1.1	mbar
Air pressure roller 2	1.1	-1.4	5.0	mbar
Air pressure roller 3	1.0	-1.1	5.0	mbar
Air pressure roller 4	-0.7	-2.5	4.2	mbar
Air pressure roller 5	-1.5	-2.5	2.6	mbar
Boiler temperature	981	826	1,091	°C
O2 in boiler	10.0	0.2	15.2	%
CO in boiler	15.9	0	675	mg/m ³
CO in chimney	22.7	0	300	mg/m ³
Steam pressure	38.0	36.0	38.9	bar
Steam temperature	396	364	413	°C

B. 2: Final network architecture



B. 3: Hyperparameters of the digital twin

Hyperparameter	Value
PCA factor	0.995
Batch size	256
Optimizer	RMSProp
Learning rate	0.001
Patience	100
Best epoch	402

B. 4: Hyperparameters for training of DQN algorithm

Hyperparameter	Value
Training episodes	334
Episode length (i.e., steps)	120
Batch size	32
Gamma	0.99
Learning rate	0.0005
Final epsilon	0.02
Exploration fraction	0.1
Replay buffer size	50,000
Training frequency	4
Target network update (every XX steps)	600
Stacks	4
Learning starts after XX steps	1200

REFERENCES

- Adams, D., Oh, D. H., Kim, D. W., Lee, C. H., & Oh, M. (2021). Deep reinforcement learning optimization framework for a power generation plant considering performance and environmental issues. *Journal of Cleaner Production*, 291, 125915. <https://doi.org/10.1016/j.jclepro.2021.125915>
- Aleksić, A., Stefanović, M., Arsovski, S., & Tadić, D. (2013). An assessment of organizational resilience potential in SMEs of the process industry, a fuzzy approach. *Journal of Loss Prevention in the Process Industries*, 26(6), 1238–1245. <https://doi.org/10.1016/j.jlp.2013.06.004>
- Almeida, F., Superior, I., Gaya, P., Queirós, A., & Faria, D. (2017). Strengths and Limitations of Qualitative and Quantitative Research Methods. *European Journal of Education Studies*, 3(9), 369–387. <https://doi.org/10.5281/zenodo.887089>
- AlSheibani, S., Cheung, Y., & Messom, C. (2018). Artificial intelligence adoption: AI-readiness at firm-level. *Proceedings of the 22nd Pacific Asia Conference on Information Systems - Opportunities and Challenges for the Digitized Society: Are We Ready?, PACIS 2018*.
- AlSheibani, S., Cheung, Y., & Messom, C. (2020). Re-thinking the competitive landscape of artificial intelligence. *Proceedings of the Annual Hawaii International Conference on System Sciences*, 5861–5870. <https://doi.org/10.24251/hicss.2020.718>
- Amini, A., Gilitschenski, I., Phillips, J., Moseyko, J., Banerjee, R., Karaman, S., & Rus, D. (2020). Learning Robust Control Policies for End-to-End Autonomous Driving from Data-Driven Simulation. *IEEE Robotics and Automation Letters*, 5(2), 1143–1150. <https://doi.org/10.1109/LRA.2020.2966414>
- Ammanath, B., Hupfer, S., & Jarvis, D. (2020). State of AI in the Enterprise, 3rd Edition: Thriving in the era of pervasive AI. *Deloitte Insights*, 1–25.
- Andrychowicz, O. A. M., Baker, B., Chociej, M., Józefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., Schneider, J., Sidor, S., Tobin, J., Welinder, P., Weng, L., & Zaremba, W. (2020). Learning dexterous in-hand manipulation. *International Journal of Robotics Research*, 39(1), 3–20. <https://doi.org/10.1177/0278364919887447>
- Appunn, K., Haas, Y., & Wettengel, J. (2021). *Germany's energy consumption and power mix in charts*. Clean Energy Wire. <https://www.cleanenergywire.org/factsheets/germanys-energy-consumption-and-power-mix-charts>
- Arel, I., Liu, C., Urbanik, T., & Kohls, A. G. (2010). Reinforcement learning-based multi-agent system for network traffic signal control. *IET Intelligent Transport Systems*, 4(2), 128–135. <https://doi.org/10.1049/iet-its.2009.0070>
- Azadeh, A., Salmanzadeh-Meydani, N., & Motevali-Haghighi, S. (2017). Performance optimization of an aluminum factory in economic crisis by integrated resilience engineering and mathematical programming. *Safety Science*, 91, 335–350. <https://doi.org/10.1016/j.ssci.2016.08.030>
- Azevedo, S. G., Carvalho, H., & Cruz-Machado, V. (2016). LARG index: A benchmarking tool for improving the leanness, agility, resilience and greenness of the automotive supply chain. *Benchmarking: An International Journal*, 23(6), 1472–1499. <https://doi.org/10.1108/BIJ-07-2014-0072>
- Baabdullah, A. M., Alalwan, A. A., Slade, E. L., Raman, R., & Khatatneh, K. F. (2021). SMEs and artificial intelligence (AI): Antecedents and consequences of AI-based B2B practices. *Industrial Marketing Management*, 98(September), 255–270. <https://doi.org/10.1016/j.indmarman.2021.09.003>
- Baah, C., Opoku Agyeman, D., Acquah, I. S. K., Agyabeng-Mensah, Y., Afum, E., Issau, K., Ofori, D., & Faibil, D. (2022). Effect of information sharing in supply chains: understanding the roles of supply chain visibility, agility, collaboration on supply chain performance. *Benchmarking: An International Journal*, 29(2), 434–455. <https://doi.org/10.1108/BIJ-08-2020-0453>
- Balakrishnan, T., Chui, M., & Henke, N. (2020). Global Survey: The State of AI in 2020. *McKinsey Digital*, November, 1–14. <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/global-survey-the-state-of-ai-in-2020>
- Bauer, D., Böhm, M., Bauernhansl, T., & Sauer, A. (2021). Increased resilience for manufacturing systems in supply networks through data - based turbulence mitigation. *Production Engineering*, 0123456789. <https://doi.org/10.1007/s11740-021-01036-4>
- Behnke, S., & Bennewitz, M. (2005). Learning to Play Soccer using Imitative Reinforcement. *ICRA 2005 Workshop "Social Aspects of Robot Programming through Demonstration."* http://www.ais.uni-bonn.de/nimbro/papers/ICRA05ws_behnke_bennewitz.pdf
- Benbya, H., Davenport, T. H., & Pachidi, S. (2020). Artificial Intelligence in Organizations: Current State and Future Opportunities. *MIS Quarterly Executive*, 19(4). <https://doi.org/10.2139/ssrn.3741983>
- Bettoni, A., Matteri, D., Montini, E., Gladysz, B., & Carpanzano, E. (2021). An AI adoption model for SMEs: A conceptual framework. *IFAC-PapersOnLine*, 54(1), 702–708. <https://doi.org/10.1016/j.ifacol.2021.08.082>
- Björkdahl, J. (2020). Strategies for Digitalization in Manufacturing Firms. *California Management Review*,

- 62(4), 17–36. <https://doi.org/10.1177/0008125620920349>
- Brock, J. K. U., & von Wangenheim, F. (2019). Demystifying AI: What digital transformation leaders can teach you about realistic artificial intelligence. *California Management Review*, 61(4), 110–134. <https://doi.org/10.1177/1536504219865226>
- Brown, C., Seville, E., & Vargo, J. (2017). Measuring the organizational resilience of critical infrastructure providers: A New Zealand case study. *International Journal of Critical Infrastructure Protection*, 18, 37–49. <https://doi.org/10.1016/j.ijcip.2017.05.002>
- Brown, N. A., Rovins, J. E., Feldmann-Jensen, S., Orchiston, C., & Johnston, D. (2019). Measuring disaster resilience within the hotel sector: An exploratory survey of Wellington and Hawke’s Bay, New Zealand hotel staff and managers. *International Journal of Disaster Risk Reduction*, 33, 108–121. <https://doi.org/10.1016/j.ijdrr.2018.09.014>
- Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlström, P., Henke, N., & Trench, M. (2017). Artificial Intelligence: The Next Digital Frontier? *McKinsey Global Institute*.
- Cassidy, L., & Barnes, G. D. (2012). Understanding Household Connectivity and Resilience in Marginal Rural Communities through Social Network Analysis in the Village of Habu, Botswana. *Ecology and Society*, 17(4), art11. <https://doi.org/10.5751/ES-04963-170411>
- Chakraborty, D., Elhegazy, H., Elzarka, H., & Gutierrez, L. (2020). A novel construction cost prediction model using hybrid natural and light gradient boosting. *Advanced Engineering Informatics*, 46(April), 101201. <https://doi.org/10.1016/j.aei.2020.101201>
- Chakraborty, D., Elzarka, H., & Bhatnagar, R. (2016). Generation of accurate weather files using a hybrid machine learning methodology for design and analysis of sustainable and resilient buildings. *Sustainable Cities and Society*, 24, 33–41. <https://doi.org/10.1016/j.scs.2016.04.009>
- Chatterjee, S., Rana, N. P., Dwivedi, Y. K., & Baabdullah, A. M. (2021). Understanding AI adoption in manufacturing and production firms using an integrated TAM-TOE model. *Technological Forecasting and Social Change*, 170(May), 120880. <https://doi.org/10.1016/j.techfore.2021.120880>
- Chen, H., Li, L., & Chen, Y. (2021). Explore success factors that impact artificial intelligence adoption on telecom industry in China. *Journal of Management Analytics*, 8(1), 36–68. <https://doi.org/10.1080/23270012.2020.1852895>
- Cheng, Y., Huang, Y., Pang, B., & Zhang, W. (2018). ThermalNet: A deep reinforcement learning-based combustion optimization system for coal-fired boiler. *Engineering Applications of Artificial Intelligence*, 74(July), 303–311. <https://doi.org/10.1016/j.engappai.2018.07.003>
- Chow, T. T., Zhang, G. Q., Lin, Z., & Song, C. L. (2002). Global optimization of absorption chiller system by genetic algorithm and neural network. *Energy and Buildings*, 34(1), 103–109. [https://doi.org/10.1016/S0378-7788\(01\)00085-8](https://doi.org/10.1016/S0378-7788(01)00085-8)
- Chowdhury, M. M. H., & Quaddus, M. (2016). Supply chain readiness, response and recovery for resilience. *Supply Chain Management: An International Journal*, 21(6), 709–731. <https://doi.org/10.1108/SCM-12-2015-0463>
- Christopher, M., & Lee, H. (2004). Mitigating supply chain risk through improved confidence. *International Journal of Physical Distribution & Logistics Management*, 34(5), 388–396. <https://doi.org/10.1108/09600030410545436>
- Christopher, M., & Peck, H. (2004). Building the Resilient Supply Chain. *The International Journal of Logistics Management*, 15(2), 1–14. <https://doi.org/10.1108/09574090410700275>
- Chui, M., & Malhotra, S. (2018). AI adoption advances, but foundational barriers remain. *McKinsey Analytics*, November, 11. <https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain>
- Chui, M., Manyika, J., Miremadi, M., Henke, N., Chung, R., Nel, P., & Malhotra, S. (2018). Notes From the AI Frontier: Insights From Hundreds of Use Cases. *McKinsey Global Institute*, 1–36.
- Correani, A., De Massis, A., Frattini, F., Petruzzelli, A. M., & Natalicchio, A. (2020). Implementing a Digital Strategy: Learning from the Experience of Three Digital Transformation Projects. *California Management Review*, 62(4), 37–56. <https://doi.org/10.1177/0008125620934864>
- Coughlan, M., Cronin, P., & Ryan, F. (2009). Process and limitations. *International Journal of Therapy and Rehabilitation*, 16(1), 9–16.
- Cutter, S. L., Ash, K. D., & Emrich, C. T. (2016). Urban–Rural Differences in Disaster Resilience. *Annals of the American Association of Geographers*, 106(6), 1236–1252. <https://doi.org/10.1080/24694452.2016.1194740>
- Davenport, T. H. (2018). *The AI Advantage: How to Put the Artificial Intelligence Revolution to Work*. MIT Press.
- Davenport, T. H. (2021). Enterprise Adoption and Management of Artificial Intelligence. *Management and Business Review*, 1(1), 165–172.
- Davenport, T. H., & Abbey, J. (2018). Can We Solve AI’s “Trust Problem.” *MIT Sloan Management Review*.

- Davenport, T. H., & Ronanki, R. (2018). Artificial Intelligence for the Real World. *Harvard Business Review*, 1(January-February), 108–117.
- Davenport, T. H., & Westerman, G. (2018). Why so many high-profile digital transformations fail. *Harvard Business Review*, 2–6. <https://hbr.org/2018/03/why-so-many-high-profile-digital-transformations-fail>
- Davis, F. (1985). A technology acceptance model for empirically testing new end-user information systems: theory and results. In *Massachusetts Institute of Technology*. <https://doi.org/10.1126/science.146.3652.1648>
- de Moraes Sá, J. C., Tivet, F., Lal, R., Briedis, C., Hartman, D. C., dos Santos, J. Z., & dos Santos, J. B. (2014). Long-term tillage systems impacts on soil C dynamics, soil resilience and agronomic productivity of a Brazilian Oxisol. *Soil and Tillage Research*, 136, 38–50. <https://doi.org/10.1016/j.still.2013.09.010>
- Denyer, D., & Tranfield, D. (2009). Producing a Systematic Review. In *The SAGE Handbook of Organizational Research Methods* (pp. 671–689).
- Deutsche Bank Research. (2017). *Winners and losers of the Industrial Internet*. [http://www.fullertreacymoney.com/system/data/files/PDFs/2017/April/7th/0900b8c08c914e35 \(1\).pdf](http://www.fullertreacymoney.com/system/data/files/PDFs/2017/April/7th/0900b8c08c914e35%20(1).pdf)
- Dewi, S., & Ahamat, A. (2018). The role of entrepreneurial orientation in achieving organization performance through business model innovation and asset relational collaboration. *Humanities and Social Sciences Reviews*, 6(2), 100–105. <https://doi.org/10.18510/hssr.2018.6212>
- Di Cairano, S. (2012). An Industry perspective on MPC in large volumes applications: Potential benefits and open challenges. In *IFAC Proceedings Volumes (IFAC-PapersOnline)* (Vol. 4, Issue PART 1). IFAC. <https://doi.org/10.3182/20120823-5-NL-3013.00040>
- Duan, Y., Andrychowicz, M., Stadie, B., Ho, J., Schneider, J., Sutskever, I., Abbeel, P., & Zaremba, W. (2017). One-shot imitation learning. *Advances in Neural Information Processing Systems, 2017-Decem(Nips)*, 1088–1099.
- Dulac-Arnold, G., Levine, N., Mankowitz, D. J., Li, J., Paduraru, C., Gowal, S., & Hester, T. (2020). *An empirical investigation of the challenges of real-world reinforcement learning*. <http://arxiv.org/abs/2003.11881>
- Dulac-Arnold, G., Mankowitz, D., & Hester, T. (2019). *Challenges of Real-World Reinforcement Learning*. <http://arxiv.org/abs/1904.12901>
- Eisenhardt, K. M. (1989). Building Theories from Case Study Research. *The Academy of Management Review*, 14(4), 532. <https://doi.org/10.2307/258557>
- Eisenhardt, K. M., & Graebner, M. E. (2007). Theory Building From Cases: Opportunities And Challenges. *Academy of Management Journal*, 50(1), 25–32. <https://doi.org/10.5465/amj.2007.24160888>
- Elhegazy, H., Chakraborty, D., Elzarka, H., Ebid, A. M., Mahdi, I. M., Aboul Haggag, S. Y., & Abdel Rashid, I. (2022). Artificial Intelligence for Developing Accurate Preliminary Cost Estimates for Composite Flooring Systems of Multi-Storey Buildings. *Journal of Asian Architecture and Building Engineering*, 21(1), 120–132. <https://doi.org/10.1080/13467581.2020.1838288>
- Elhegazy, H., & Kamal, M. (2022). Implementing nuclear power plants (NPPs): state of the art, challenges, and opportunities. *Innovative Infrastructure Solutions*, 7(1), 11. <https://doi.org/10.1007/s41062-021-00611-z>
- Ellefsen, A. P. T., Oleśków-Szlapka, J., Pawłowski, G., & Toboła, A. (2019). Striving for excellence in AI implementation: AI maturity model framework and preliminary research results. *Logforum*, 15(3), 363–376. <https://doi.org/10.17270/J.LOG.2019.354>
- European Commission. (2003). *SME definition*. Internal Market, Industry, Entrepreneurship and SMEs. https://single-market-economy.ec.europa.eu/smes/sme-definition_en
- Fiksel, J., Polyviou, M., Croxton, K. L., & Pettit, T. J. (2015). From risk to resilience: Learning to deal with disruption. *MIT Sloan Management Review*, 56(2), 79–86.
- Flamme, S., Hanewinkel, J., Quicker, P., & Weber, K. (2018). Energieerzeugung aus Abfällen. *Umweltbundesamt*, 159. <http://www.umweltbundesamt.de/publikationen>
- Flynn, D. (2003). Thermal Power Plant Simulation and Control. *Thermal Power Plant Simulation and Control*. <https://doi.org/10.1049/pbpo043e>
- Forbes, M. G., Patwardhan, R. S., Hamadah, H., & Gopaluni, R. B. (2015). Model predictive control in industry: Challenges and opportunities. *IFAC-PapersOnLine*, 28(8), 531–538. <https://doi.org/10.1016/j.ifacol.2015.09.022>
- Fountaine, T., McCarthy, B., & Saleh, T. (2019). Building the AI-Powered Organization. *Harvard Business Review*.
- Fraccascia, L., Giannoccaro, I., & Albino, V. (2017). Rethinking Resilience in Industrial Symbiosis: Conceptualization and Measurements. *Ecological Economics*, 137, 148–162. <https://doi.org/10.1016/j.ecolecon.2017.02.026>
- François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., & Pineau, J. (2018). An Introduction to Deep Reinforcement Learning. *Foundations and Trends in Machine Learning*, II(3–4), 1–140. <https://doi.org/10.1561/22000000071.Vincent>

- Fu, J., Xiao, H., Wang, H., & Zhou, J. (2020). Control Strategy for Denitrification Efficiency of Coal-Fired Power Plant Based on Deep Reinforcement Learning. *IEEE Access*, 8, 65127–65136. <https://doi.org/10.1109/ACCESS.2020.2985233>
- Gazol, A., Ribas, M., Gutiérrez, E., & Camarero, J. J. (2017). Aleppo pine forests from across Spain show drought-induced growth decline and partial recovery. *Agricultural and Forest Meteorology*, 232, 186–194. <https://doi.org/10.1016/j.agrformet.2016.08.014>
- Gfrerer, A., Hutter, K., Füller, J., & Ströhle, T. (2021). Ready or Not: Managers' and Employees' Different Perceptions of Digital Readiness. *California Management Review*, 63(2), 23–48. <https://doi.org/10.1177/0008125620977487>
- Ghosh, B., Wilson, H. J., Burden, A., & Daugherty, P. R. (2019). Taking a Systems Approach to Adopting AI. *Harvard Business Review*.
- Glaser, B. G., & Strauss, A. L. (1999). *The Discovery of Grounded Theory*. Routledge. <https://doi.org/10.4324/9780203793206>
- Glavic, M. (2019). (Deep) Reinforcement learning for electric power system control and related problems: A short review and perspectives. *Annual Reviews in Control*, 48, 22–35. <https://doi.org/10.1016/j.arcontrol.2019.09.008>
- Gosavi, A. (2009). Reinforcement learning: A tutorial survey and recent advances. *INFORMS Journal on Computing*, 21(2), 178–192. <https://doi.org/10.1287/ijoc.1080.0305>
- Graveline, N., & Grémont, M. (2017). Measuring and understanding the microeconomic resilience of businesses to lifeline service interruptions due to natural disasters. *International Journal of Disaster Risk Reduction*, 24(September 2016), 526–538. <https://doi.org/10.1016/j.ijdrr.2017.05.012>
- Grisanzio, K. A., Goldstein-Piekarski, A. N., Wang, M. Y., Rashed Ahmed, A. P., Samara, Z., & Williams, L. M. (2018). Transdiagnostic Symptom Clusters and Associations With Brain, Behavior, and Daily Function in Mood, Anxiety, and Trauma Disorders. *JAMA Psychiatry*, 75(2), 201. <https://doi.org/10.1001/jamapsychiatry.2017.3951>
- Grissom, A. C., Boyd-Graber, J., He, H., Morgan, J., & Daumé, H. (2014). Don't until the final verb wait: Reinforcement learning for simultaneous machine translation. *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, 1342–1352. <https://doi.org/10.3115/v1/d14-1140>
- Grover, P., Kar, A. K., & Dwivedi, Y. K. (2022). Understanding artificial intelligence adoption in operations management: insights from the review of academic literature and social media discussions. In *Annals of Operations Research* (Vol. 308, Issues 1–2). Springer US. <https://doi.org/10.1007/s10479-020-03683-9>
- Gu, S., Holly, E., Lillicrap, T., & Levine, S. (2017). Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. *Proceedings - IEEE International Conference on Robotics and Automation*, 3389–3396. <https://doi.org/10.1109/ICRA.2017.7989385>
- Hallinan, B., & Striphas, T. (2016). Recommended for you: The Netflix Prize and the production of algorithmic culture. *New Media & Society*, 18(1), 117–137. <https://doi.org/10.1177/1461444814538646>
- Hein, D., Depeweg, S., Tokic, M., Udluft, S., Hentschel, A., Runkler, T. A., & Sterzing, V. (2018). A benchmark environment motivated by industrial control problems. *2017 IEEE Symposium Series on Computational Intelligence, SSCI 2017 - Proceedings, 2018-Janua*, 1–8. <https://doi.org/10.1109/SSCI.2017.8280935>
- Henry, D., & Ramirez-Marquez, J. E. (2016). On the Impacts of Power Outages during Hurricane Sandy—A Resilience-Based Analysis. *Systems Engineering*, 19(1), 59–75. <https://doi.org/10.1002/sys.21338>
- Hill, A., Raffin, A., Ernestus, M., Gleave, A., Kanervisto, A., Traore, R., Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., & Wu, Y. (2018). Stable Baselines. In *GitHub repository*. GitHub.
- Hillmann, J., & Guenther, E. (2021). Organizational Resilience: A Valuable Construct for Management Research? *International Journal of Management Reviews*, 23(1), 7–44. <https://doi.org/10.1111/ijmr.12239>
- Hohenstein, N.-O., Feisel, E., Hartmann, E., & Giunipero, L. (2015). Research on the phenomenon of supply chain resilience. *International Journal of Physical Distribution & Logistics Management*, 45(1/2), 90–117. <https://doi.org/10.1108/IJPDLM-05-2013-0128>
- Holmström, J. (2021). From AI to digital transformation: The AI readiness framework. *Business Horizons*. <https://doi.org/10.1016/j.bushor.2021.03.006>
- Hosseini, S., Barker, K., & Ramirez-Marquez, J. E. (2016). A review of definitions and measures of system resilience. *Reliability Engineering & System Safety*, 145, 47–61. <https://doi.org/10.1016/j.ress.2015.08.006>
- Hosseini, S., & Ivanov, D. (2020). Bayesian networks for supply chain risk, resilience and ripple effect analysis: A literature review. *Expert Systems with Applications*, 161, 113649. <https://doi.org/10.1016/j.eswa.2020.113649>
- Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of quantitative methods for supply chain resilience analysis. *Transportation Research Part E: Logistics and Transportation Review*, 125(December 2018), 285–307. <https://doi.org/10.1016/j.tre.2019.03.001>

- Hosseini, S., Morshedlou, N., Ivanov, D., Sarder, M. D., Barker, K., & Khaled, A. Al. (2019). Resilient supplier selection and optimal order allocation under disruption risks. *International Journal of Production Economics*, 213(August 2018), 124–137. <https://doi.org/10.1016/j.ijpe.2019.03.018>
- ifo. (2022). *ifo Business Climate Index for Germany*. <https://www.ifo.de/en/survey/ifo-business-climate-index-germany>
- Insee. (2020). *Large enterprise*. Definitions. <https://www.insee.fr/en/metadonnees/definition/c1035>
- Issa, H., Jabbouri, R., & Palmer, M. (2021). An Artificial Intelligence (Ai)-Readiness and Adoption Framework for Agritech Firms. *SSRN Electronic Journal*, 1–37. <https://doi.org/10.2139/ssrn.3971188>
- Jammeli, H., Ksantini, R., Ben Abdelaziz, F., & Masri, H. (2021). Sequential Artificial Intelligence Models to Forecast Urban Solid Waste in the City of Sousse, Tunisia. *IEEE Transactions on Engineering Management*, 1–11. <https://doi.org/10.1109/TEM.2021.3081609>
- Jöhnk, J., Weißert, M., & Wyrтки, K. (2021). Ready or Not, AI Comes— An Interview Study of Organizational AI Readiness Factors. *Business and Information Systems Engineering*, 63(1), 5–20. <https://doi.org/10.1007/s12599-020-00676-7>
- Jones, L., & Tanner, T. (2017). ‘Subjective resilience’: using perceptions to quantify household resilience to climate extremes and disasters. *Regional Environmental Change*, 17(1), 229–243. <https://doi.org/10.1007/s10113-016-0995-2>
- Jüttner, U., & Maklan, S. (2011). Supply chain resilience in the global financial crisis: an empirical study. *Supply Chain Management: An International Journal*, 16(4), 246–259. <https://doi.org/10.1108/13598541111139062>
- Kabalisa, R., & Altmann, J. (2021). AI Technologies and Motives for AI Adoption by Countries and Firms: A Systematic Literature Review. *Lecture Notes in Computer Science*, 13072 LNCS(December), 39–51. https://doi.org/10.1007/978-3-030-92916-9_4
- Kamalahmadi, M., & Parast, M. M. (2016). A review of the literature on the principles of enterprise and supply chain resilience: Major findings and directions for future research. *International Journal of Production Economics*, 171, 116–133. <https://doi.org/10.1016/j.ijpe.2015.10.023>
- Kaneko, T., Tsurumine, Y., Poon, J., Onuki, Y., Dai, Y., Kawabata, K., & Matsubara, T. (2019). Learning deep dynamical models of a waste incineration plant from in-furnace images and process data. *IEEE International Conference on Automation Science and Engineering*, 873–878. <https://doi.org/10.1109/COASE.2019.8842972>
- Kaviani, M. A., Tavana, M., Kowsari, F., & Rezapour, R. (2020). Supply chain resilience: a benchmarking model for vulnerability and capability assessment in the automotive industry. *Benchmarking: An International Journal*, 27(6), 1929–1949. <https://doi.org/10.1108/BIJ-01-2020-0049>
- Kegeyes, T., Süle, Z., & Abonyi, J. (2021). The Applicability of Reinforcement Learning Methods in the Development of Industry 4.0 Applications. *Complexity*, 2021, 1–31. <https://doi.org/10.1155/2021/7179374>
- Kheybari, S., Mahdi Rezaie, F., & Rezaei, J. (2021). Measuring the Importance of Decision-Making Criteria in Biofuel Production Technology Selection. *IEEE Transactions on Engineering Management*, 68(2), 483–497. <https://doi.org/10.1109/TEM.2019.2908037>
- Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Sallab, A. A. A., Yogamani, S., & Perez, P. (2021). Deep Reinforcement Learning for Autonomous Driving: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, 1–18. <https://doi.org/10.1109/TITS.2021.3054625>
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13), 3521–3526. <https://doi.org/10.1073/pnas.1611835114>
- Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. *International Journal of Robotics Research*, 32(11), 1238–1274. <https://doi.org/10.1177/0278364913495721>
- Kubosawa, S., Onishi, T., & Tsuruoka, Y. (2019). *Synthesizing Chemical Plant Operation Procedures using Knowledge, Dynamic Simulation and Deep Reinforcement Learning*. 2–5. <http://arxiv.org/abs/1903.02183>
- Kumar, P. S., & Gururaj, S. (2019). Conceptual cost modelling for sustainable construction project planning- A levenberg-marquardt neural network approach. *Applied Mathematics and Information Sciences*, 13(2), 201–208. <https://doi.org/10.18576/AMIS/130207>
- Lai, Y. H., Wu, T. C., Lai, C. F., Yang, L. T., & Zhou, X. (2021). Cognitive Optimal-Setting Control of AIoT Industrial Applications with Deep Reinforcement Learning. *IEEE Transactions on Industrial Informatics*, 17(3), 2116–2123. <https://doi.org/10.1109/TII.2020.2986501>
- Li, Yuanlong, Wen, Y., Tao, D., & Guan, K. (2020). Transforming Cooling Optimization for Green Data Center via Deep Reinforcement Learning. *IEEE Transactions on Cybernetics*, 50(5), 2002–2013. <https://doi.org/10.1109/TCYB.2019.2927410>
- Li, Yuxi. (2018). Deep Reinforcement Learning. *Lecture Notes in Networks and Systems*, 16, 426–440. <http://arxiv.org/abs/1810.06339>

- Lopez, P. A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flotterod, Y. P., Hilbrich, R., Lucken, L., Rummel, J., Wagner, P., & Wiebner, E. (2018). Microscopic Traffic Simulation using SUMO. *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC, 2018-Novem*, 2575–2582. <https://doi.org/10.1109/ITSC.2018.8569938>
- Lund, S., Manyika, J., Woetzel, J., Barriball, E., Krishnan, M., Alicke, K., Birshan, M., George, K., Smit, S., Swan, D., & Hutzler, K. (2020). Risk, resilience, and rebalancing in global value chains. *McKinsey Global Institute, August*.
- Lv, Y., & Ren, X. (2019). Approximate Nash Solutions for Multiplayer Mixed-Zero-Sum Game with Reinforcement Learning. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(12), 2739–2750. <https://doi.org/10.1109/TSMC.2018.2861826>
- Macurová, P., & Jurásková, K. (2013). Analysis of risks generated by suppliers during the period of economic fluctuations. *Amfiteatru Economic*, 15(33), 27–43.
- Majanne, Y. (2005). Model predictive pressure control of steam networks. *Control Engineering Practice*, 13(12), 1499–1505. <https://doi.org/10.1016/j.conengprac.2005.03.008>
- Mann, H. B., & Whitney, D. R. (1947). On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*, 18(1), 50–60. <https://doi.org/10.1214/aoms/1177730491>
- Mayne, D. Q. (2014). Model predictive control: Recent developments and future promise. *Automatica*, 50(12), 2967–2986. <https://doi.org/10.1016/j.automatica.2014.10.128>
- Mena, C., Melnyk, S. A., Baghersad, M., & Zobel, C. W. (2020). Sourcing Decisions under Conditions of Risk and Resilience: A Behavioral Study. *Decision Sciences*, 51(4), 985–1014. <https://doi.org/10.1111/deci.12403>
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). *Playing Atari with Deep Reinforcement Learning*. 1–9. <http://arxiv.org/abs/1312.5602>
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
- NABU. (2019). *Müllverbrennung in Deutschland*. <https://www.nabu.de/umwelt-und-ressourcen/abfall-und-recycling/verbrennung/index.html>
- Nam, K., Dutt, C. S., Chathoth, P., Daghfous, A., & Khan, M. S. (2021). The adoption of artificial intelligence and robotics in the hotel industry: prospects and challenges. *Electronic Markets*, 31(3), 553–574. <https://doi.org/10.1007/s12525-020-00442-3>
- Neumann, O., Guirguis, K., & Steiner, R. (2022). Exploring artificial intelligence adoption in public organizations: a comparative case study. *Public Management Review*, 00(00), 1–27. <https://doi.org/10.1080/14719037.2022.2048685>
- NewVantage Partners. (2019). How big data and AI are accelerating business transformation. In *Big Data and AI Executive Survey 2019*.
- Nian, R., Liu, J., & Huang, B. (2020). A review On reinforcement learning: Introduction and applications in industrial process control. *Computers and Chemical Engineering*, 139, 106886. <https://doi.org/10.1016/j.compchemeng.2020.106886>
- Norrman, A., & Jansson, U. (2004). Ericsson’s proactive supply chain risk management approach after a serious sub-supplier accident. *International Journal of Physical Distribution & Logistics Management*, 34(5), 434–456. <https://doi.org/10.1108/09600030410545463>
- Nortje, M. A., & Grobbelaar, S. S. (2020). A Framework for the Implementation of Artificial Intelligence in Business Enterprises: A Readiness Model. *Proceedings - 2020 IEEE International Conference on Engineering, Technology and Innovation, ICE/ITMC 2020*. <https://doi.org/10.1109/ICE/ITMC49519.2020.9198436>
- OECD. (2021). *Strengthening Economic Resilience Following the COVID-19 Crisis*. OECD. <https://doi.org/10.1787/2a7081d8-en>
- Öksüz, N., Bouschery, S., Schlappa, M., Sporkmann, J., & Unterberg, M. (2021). Towards an Artificial Intelligence based Approach for Manufacturing Resilience. In V. D. I. W. GmbH (Ed.), *Automation 2021: Navigating towards resilient Production* (1st ed., pp. 649–664). VDI Verlag. <https://doi.org/10.51202/9783181023921-649>
- Oliveira, T., & Martins, M. F. (2011). Literature Review of Information Technology Adoption Models at Firm Level. *The Electronic Journal Information Systems Evaluation*, 14(1), 110–121.
- Omidvar, M., Mazloumi, A., Mohammad Fam, I., & Nirumand, F. (2017). Development of a framework for resilience measurement: Suggestion of fuzzy Resilience Grade (RG) and fuzzy Resilience Early Warning Grade (REWG). *Work*, 56(3), 463–474. <https://doi.org/10.3233/WOR-172512>
- Ota, K., Jha, D. K., Romeres, Di., Van Baar, J., Smith, K. A., Semitsu, T., Oiki, T., Sullivan, A., Nikovski, D., &

- Tenenbaum, J. B. (2021). Data-Efficient Learning for Complex and Real-Time Physical Problem Solving Using Augmented Simulation. *IEEE Robotics and Automation Letters*, 6(2), 4241–4248. <https://doi.org/10.1109/LRA.2021.3068887>
- Owen, G. (2020). What makes climate change adaptation effective? A systematic review of the literature. *Global Environmental Change*, 62, 102071. <https://doi.org/10.1016/j.gloenvcha.2020.102071>
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *The BMJ*, 372. <https://doi.org/10.1136/bmj.n71>
- Panter-Brick, C., Hadfield, K., Dajani, R., Eggerman, M., Ager, A., & Ungar, M. (2018). Resilience in Context: A Brief and Culturally Grounded Measure for Syrian Refugee and Jordanian Host-Community Adolescents. *Child Development*, 89(5), 1803–1820. <https://doi.org/10.1111/cdev.12868>
- Pasquale, F. (2015). *The black box society*. Harvard University Press.
- Pereira, C. R., Christopher, M., & Lago Da Silva, A. (2014). Achieving supply chain resilience: the role of procurement. *Supply Chain Management*, 19(November), 626–642. <https://doi.org/10.1108/SCM-09-2013-0346>
- Perera, A. T. D., & Kamalaruban, P. (2021). Applications of reinforcement learning in energy systems. *Renewable and Sustainable Energy Reviews*, 137(November 2020), 110618. <https://doi.org/10.1016/j.rser.2020.110618>
- Pettit, T. J., Croxton, K. L., & Fiksel, J. (2013). Ensuring supply chain resilience: Development and implementation of an assessment tool. *Journal of Business Logistics*, 34(1), 46–76. <https://doi.org/10.1111/jbl.12009>
- Pettit, T. J., Fiksel, J., & Croxton, K. L. (2010). Ensuring Supply Chain Resilience: Development of a Conceptual Framework. *Journal of Business Logistics*, 31(1), 1–21. <https://doi.org/10.1002/j.2158-1592.2010.tb00125.x>
- Pires Ribeiro, J., & Barbosa-Povoa, A. (2018). Supply Chain Resilience: Definitions and quantitative modelling approaches – A literature review. *Computers and Industrial Engineering*, 115(May 2017), 109–122. <https://doi.org/10.1016/j.cie.2017.11.006>
- Ponis, S. T., & Koronis, E. (2012). Supply Chain Resilience: Definition Of Concept And Its Formative Elements. *Journal of Applied Business Research (JABR)*, 28(5), 921. <https://doi.org/10.19030/jabr.v28i5.7234>
- Ponomarov, S. Y., & Holcomb, M. C. (2009). Understanding the concept of supply chain resilience. In *The International Journal of Logistics Management* (Vol. 20, Issue 1). <https://doi.org/10.1108/09574090910954873>
- Pumplun, L., Tauchert, C., & Heidt, M. (2019). A new organizational chassis for artificial intelligence – exploring organizational readiness factors. *ECIS 2019 Proceedings*. https://aisel.aisnet.org/ecis2019_rp/106
- PWC. (2022). *Maschinenbau-Barometer*. <https://www.pwc.de/maschinenbau-barometer>
- Qin, S. J., & Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7), 733–764. [https://doi.org/10.1016/S0967-0661\(02\)00186-7](https://doi.org/10.1016/S0967-0661(02)00186-7)
- Queiroz, M. M., Fosso Wamba, S., & Branski, R. M. (2022). Supply chain resilience during the COVID-19: empirical evidence from an emerging economy. *Benchmarking: An International Journal*, 29(6), 1999–2018. <https://doi.org/10.1108/BIJ-08-2021-0454>
- Rajesh, R. (2019). A fuzzy approach to analyzing the level of resilience in manufacturing supply chains. *Sustainable Production and Consumption*, 18, 224–236. <https://doi.org/10.1016/j.spc.2019.02.005>
- Rashid, Y., Rashid, A., Warraich, M. A., Sabir, S. S., & Waseem, A. (2019). Case Study Method: A Step-by-Step Guide for Business Researchers. *International Journal of Qualitative Methods*, 18, 1–13. <https://doi.org/10.1177/1609406919862424>
- Roberts, K. H., & Rousseau, D. M. (1989). Research in Nearly Failure-Free, High-Reliability Organizations: Having the Bubble. *IEEE Transactions on Engineering Management*, 36(2), 132–139. <https://doi.org/10.1109/17.18830>
- Robertson, I. T., Cooper, C. L., Sarkar, M., & Curran, T. (2015). Resilience training in the workplace from 2003 to 2014: A systematic review. *Journal of Occupational and Organizational Psychology*, 88(3), 533–562. <https://doi.org/10.1111/joop.12120>
- Rogers, E. M. (1995). *Diffusion of innovations* (4th ed.). Free Press.
- Rousseau, D. M., Manning, J., & Denyer, D. (2008). 11 Evidence in Management and Organizational Science: Assembling the Field's Full Weight of Scientific Knowledge Through Syntheses. *Academy of Management Annals*, 2(1), 475–515. <https://doi.org/10.5465/19416520802211651>
- Sambowo, A. L., & Hidayatno, A. (2021). Resilience Index Development for the Manufacturing Industry based

- on Robustness, Resourcefulness, Redundancy, and Rapidity. *International Journal of Technology*, 12(6), 1177. <https://doi.org/10.14716/ijtech.v12i6.5229>
- Sapeciay, Z., Wilkinson, S., & Costello, S. B. (2017). Building organisational resilience for the construction industry. *International Journal of Disaster Resilience in the Built Environment*, 8(1), 98–108. <https://doi.org/10.1108/IJDRBE-05-2016-0020>
- Sarkar, M., & Fletcher, D. (2014). Psychological resilience in sport performers: a review of stressors and protective factors. *Journal of Sports Sciences*, 1–16. <https://doi.org/10.1080/02640414.2014.901551>
- Sasaki, H., Hirabayashi, T., Kawabata, K., Onuki, Y., & Matsubara, T. (2020). Bayesian policy optimization for waste crane with garbage inhomogeneity. *IEEE Robotics and Automation Letters*, 5(3), 4533–4540. <https://doi.org/10.1109/LRA.2020.3002204>
- Schlappa, M., Hegemann, J., Mathur, T., & Spinler, S. (2022). *Overcoming Barriers to AI Adoption: A Case Study from the Waste Incineration Industry [Unpublished working paper]*.
- Schlappa, M., Spinler, S., & Hegemann, J. (2022). Optimizing Control of Waste Incineration Plants Using Reinforcement Learning and Digital Twins. *IEEE Transactions on Engineering Management*. <https://doi.org/10.1109/TEM.2022.3201434>
- Schlappa, M., Sporkmann, J., Unterberg, M., & Bouschery, S. (2022). *Benchmarking Organizational Resilience in the German Manufacturing Industry [Unpublished working paper]*.
- Schmidhuber, J. (2015). Deep Learning in neural networks: An overview. *Neural Networks*, 61, 85–117. <https://doi.org/10.1016/j.neunet.2014.09.003>
- Schoukens, J., & Ljung, L. (2019). Nonlinear System Identification: A User-Oriented Road Map. *IEEE Control Systems*, 39(6), 28–99. <https://doi.org/10.1109/MCS.2019.2938121>
- Schubnel, B., Carrillo, R. E., Alet, P.-J., & Hutter, A. (2020). A Hybrid Learning Method for System Identification and Optimal Control. *IEEE Transactions on Neural Networks and Learning Systems*, 32(9), 4096–4110. <https://doi.org/10.1109/tnnls.2020.3016906>
- Seiler, C., & Wohlrabe, K. (2013). Das ifo Geschäftsklima und die deutsche Konjunktur. *Ifo Schnelldienst*, 66(18).
- Sharifi, A. (2016). A critical review of selected tools for assessing community resilience. *Ecological Indicators*, 69, 629–647. <https://doi.org/10.1016/j.ecolind.2016.05.023>
- Sheffi, Y., & Rice, J. B. (2005). A Supply Chain View of the Resilient Enterprise. *MIT Sloan Management Review*, 47(1), 41–48.
- Shet, A., Carr, K., Danovaro-Holliday, M. C., Sodha, S. V., Prosperi, C., Wunderlich, J., Wonodi, C., Reynolds, H. W., Mirza, I., Gacic-Dobo, M., O'Brien, K. L., & Lindstrand, A. (2022). Impact of the SARS-CoV-2 pandemic on routine immunisation services: evidence of disruption and recovery from 170 countries and territories. *The Lancet Global Health*, 10(2), e186–e194. [https://doi.org/10.1016/S2214-109X\(21\)00512-X](https://doi.org/10.1016/S2214-109X(21)00512-X)
- Shin, S., Lee, Y., Kim, M., Park, J., Lee, S., & Min, K. (2020). Deep neural network model with Bayesian hyperparameter optimization for prediction of NOx at transient conditions in a diesel engine. *Engineering Applications of Artificial Intelligence*, 94(June), 103761. <https://doi.org/10.1016/j.engappai.2020.103761>
- Shirali, G. A., Mohammadfam, I., & Ebrahimipour, V. (2013). A new method for quantitative assessment of resilience engineering by PCA and NT approach: A case study in a process industry. *Reliability Engineering & System Safety*, 119, 88–94. <https://doi.org/10.1016/j.res.2013.05.003>
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489. <https://doi.org/10.1038/nature16961>
- Singh, C. S., Soni, G., & Badhotiya, G. K. (2019). Performance indicators for supply chain resilience: review and conceptual framework. *Journal of Industrial Engineering International*, 15(s1), 105–117. <https://doi.org/10.1007/s40092-019-00322-2>
- Sinn, H.-W., & Abberger, K. (2006). Zur Prognosekraft des ifo Indikators. *Ifo Schnelldienst*, 59(04), 35–36. <https://econpapers.repec.org/RePEc:ces:ifosdt:v:59:y:2006:i:04:p:35-36>
- Sjödén, D., Parida, V., & Visnjic, I. (2022). How Can Large Manufacturers Digitalize Their Business Models? A Framework for Orchestrating Industrial Ecosystems. *California Management Review*, 64(3), 49–77. <https://doi.org/10.1177/00081256211059140>
- Spearman, C. (1987). The Proof and Measurement of Association between Two Things. *The American Journal of Psychology*, 100(3/4), 441. <https://doi.org/10.2307/1422689>
- Spieske, A., & Birkel, H. (2021). Improving supply chain resilience through industry 4.0: A systematic literature review under the impressions of the COVID-19 pandemic. *Computers and Industrial Engineering*, 158(June), 107452. <https://doi.org/10.1016/j.cie.2021.107452>
- Steege, F.-F., Hartmann, A., Schaffernicht, E., & Gross, H.-M. (2010). Reinforcement Learning Based Neural Controllers for Dynamic Processes without Exploration. *Artificial Neural Networks - ICANN 2010, 20th*

- International Conference, Thessaloniki, Greece, September 15-18, 2010, Proceedings, Part II*, 223(4), 222–227. https://doi.org/10.1007/978-3-642-15822-3_29
- Stephan, V., Debes, K., Gross, H.-M., Wintrich, F., & Wintrich, H. (2001). A new control scheme for combustion processes using reinforcement learning based on neural networks. *International Journal of Computational Intelligence and Applications*, 01(02), 121–136. <https://doi.org/10.1142/S1469026801000172>
- Stephan, V., Wintrich, F., König, A., & Debes, K. (2004). Application of Action Dependent Heuristic Dynamic Programming to Control an Industrial Waste Incineration Plant. *SOAVE 2004*, 743, 262–270.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. A Bradford Book.
- Takaghaj, S. M., MacNab, C. J. B., Westwick, D., & Boiko, I. (2014). Neural-adaptive control of waste-to-energy steam generators. *IEEE Transactions on Control Systems Technology*, 22(5), 1920–1926. <https://doi.org/10.1109/TCST.2013.2292818>
- Tang, O., & Nurmaya Musa, S. (2011). Identifying risk issues and research advancements in supply chain risk management. *International Journal of Production Economics*, 133(1), 25–34. <https://doi.org/10.1016/j.ijpe.2010.06.013>
- Tao, C., Longya, R., & Xian, G. (2019). AI innovation for advancing public service: The case of China's first Administrative Approval Bureau. *ACM International Conference Proceeding Series, August*, 100–108. <https://doi.org/10.1145/3325112.3325243>
- Tesauro, G. (1994). TD-Gammon, a Self-Teaching Backgammon Program, Achieves Master-Level Play. *Neural Computation*, 6(2), 215–219. <https://doi.org/10.1162/neco.1994.6.2.215>
- Tornatzky, L. G., Fleischer, M., & Chakrabarti, A. K. (1990). *The processes of technological innovation*. Lexington Books.
- Tranfield, D., Denyer, D., & Smart, P. (2003). Towards a Methodology for Developing Evidence-Informed Management Knowledge by Means of Systematic Review. *British Journal of Management*, 14(3), 207–222. <https://doi.org/10.1111/1467-8551.00375>
- Tukamuhabwa, B. R., Stevenson, M., Busby, J., & Zorzini, M. (2015). Supply chain resilience: Definition, review and theoretical foundations for further study. *International Journal of Production Research*, 53(18), 5592–5623. <https://doi.org/10.1080/00207543.2015.1037934>
- Usher, L. E., Yusuf, J. W., & Covi, M. (2019). Assessing tourism business resilience in Virginia Beach. *International Journal of Tourism Cities*, 6(2), 397–414. <https://doi.org/10.1108/IJTC-02-2019-0019>
- Utami, I. D., Santosa, I., & Vidya Leila, M. R. (2021). Priority resilience strategy for micro, small, and medium enterprises for dealing with natural disasters. *International Journal of Disaster Risk Reduction*, 55(November 2019), 102074. <https://doi.org/10.1016/j.ijdrr.2021.102074>
- Venkatesh, Morris, Davis, & Davis. (2003). User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly*, 27(3), 425. <https://doi.org/10.2307/30036540>
- Vuong, Q. H., Ho, M. T., Vuong, T. T., La, V. P., Ho, M. T., Nghiem, K. C. P., Tran, B. X., Giang, H. H., Giang, T. V., Latkin, C., Nguyen, H. K. T., Ho, C. S. H., & Ho, R. C. M. (2019). Artificial intelligence vs. Natural stupidity: Evaluating ai readiness for the vietnamese medical information system. *Journal of Clinical Medicine*, 8(2). <https://doi.org/10.3390/jcm8020168>
- Wagner, S. M., & Neshat, N. (2012). A comparison of supply chain vulnerability indices for different categories of firms. *International Journal of Production Research*, 50(11), 2877–2891. <https://doi.org/10.1080/00207543.2011.561540>
- Waters, J., & Adger, W. N. (2017). Spatial, network and temporal dimensions of the determinants of adaptive capacity in poor urban areas. *Global Environmental Change*, 46, 42–49. <https://doi.org/10.1016/j.gloenvcha.2017.06.011>
- Weber, K., Quicker, P., Hanewinkel, J., & Flamme, S. (2020). Status of waste-to-energy in Germany, Part I – Waste treatment facilities. *Waste Management and Research*, 38(1_suppl), 23–44. <https://doi.org/10.1177/0734242X19894632>
- Weiner, B. J. (2009). A theory of organizational readiness for change. *Implementation Science*, 4(1), 1–9. <https://doi.org/10.1186/1748-5908-4-67>
- Windle, G., Bennett, K. M., & Noyes, J. (2011). A methodological review of resilience measurement scales. *Health and Quality of Life Outcomes*, 9(1), 8. <https://doi.org/10.1186/1477-7525-9-8>
- World Bank. (2022). *Manufacturing, value added (% of GDP) - Germany*. <https://data.worldbank.org/indicator/NV.IND.MANF.ZS?locations=DE>
- Wouters, M., Anderson, J. C., & Kirchberger, M. (2018). New-technology startups seeking pilot customers: Crafting a pair of value propositions. *California Management Review*, 60(4), 101–124. <https://doi.org/10.1177/0008125618778855>
- Yazdani, A., & Jeffrey, P. (2012). Applying Network Theory to Quantify the Redundancy and Structural Robustness of Water Distribution Systems. *Journal of Water Resources Planning and Management*, 138(2), 153–161. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000159](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000159)

- Yazdanparast, R., Tavakkoli-Moghaddam, R., Heidari, R., & Aliabadi, L. (2021). A hybrid Z-number data envelopment analysis and neural network for assessment of supply chain resilience: a case study. *Central European Journal of Operations Research*, 29(2), 611–631. <https://doi.org/10.1007/s10100-018-0596-x>
- Yin, R. K. (2018). *Case Study Research and Applications: Design and Methods* (6th ed.). SAGE Publications, Inc.
- Zeng, B., & Yen, B. P. C. (2017). Rethinking the role of partnerships in global supply chains: A risk-based perspective. *International Journal of Production Economics*, 185(June 2016), 52–62. <https://doi.org/10.1016/j.ijpe.2016.12.004>
- Zhan, X., Xu, H., Zhang, Y., Huo, Y., Zhu, X., Yin, H., & Zheng, Y. (2021). *DeepThermal: Combustion Optimization for Thermal Power Generating Units Using Offline Reinforcement Learning*. x, 1–19. <http://arxiv.org/abs/2102.11492>
- Zhang, L. J., Shan, Z., & Mao, Z. H. (2010). An optimal-control-based decision-making model and consulting methodology for service enterprises. *IEEE Transactions on Engineering Management*, 57(4), 607–619. <https://doi.org/10.1109/TEM.2009.2034270>
- Zheng, G., Zhang, F., Zheng, Z., Xiang, Y., Yuan, N. J., Xie, X., & Li, Z. (2018). DRN: A deep reinforcement learning framework for news recommendation. *The Web Conference 2018 - Proceedings of the World Wide Web Conference, WWW 2018*, 2, 167–176. <https://doi.org/10.1145/3178876.3185994>
- Zheng, Z.-H., Lin, X.-D., Yang, M., He, Z.-M., Bao, E., Zhang, H., & Tian, Z.-Y. (2020). Progress in the Application of Machine Learning in Combustion Studies. *ES Energy & Environment*, 1–14. <https://doi.org/10.30919/ese8c795>
- Zhou, L., Song, Y., Ji, W., & Wei, H. (2022). Machine learning for combustion. *Energy and AI*, 7, 100128. <https://doi.org/10.1016/j.egyai.2021.100128>
- Zhou, Z., Li, X., & Zare, R. N. (2017). Optimizing Chemical Reactions with Deep Reinforcement Learning. *ACS Central Science*, 3(12), 1337–1344. <https://doi.org/10.1021/acscentsci.7b00492>
- Zhuang, Z., Sun, Z., Cheng, Y., Yao, R., & Zhang, W. (2018). Modeling and optimization of paper-making wastewater treatment based on reinforcement learning. *Chinese Control Conference, CCC, 2018-July*, 8342–8346. <https://doi.org/10.23919/ChiCC.2018.8482733>
- Zou, L., Cheng, Y., Zhuang, Z., Sun, Z., & Zhang, W. (2018). An optimization control of thermal power combustion based on reinforcement learning. *Chinese Control Conference, CCC, 2018-July*, 3553–3558. <https://doi.org/10.23919/ChiCC.2018.8482853>