

Design principles for artificial intelligence-augmented decision making: An action design research study

Journal Article**Author(s):**

Herath Pathirannehelage, Savindu ; Shrestha, Yash Raj; von Krogh, Georg

Publication date:

2024

Permanent link:

<https://doi.org/10.3929/ethz-b-000666645>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

European Journal of Information Systems, <https://doi.org/10.1080/0960085X.2024.2330402>

Funding acknowledgement:

197763 - Sustaining Knowledge Creation in Online Communities: Enabling, Creating, and Maintaining (SNF)

Design principles for artificial intelligence-augmented decision making: An action design research study

Savindu Herath Pathirannehelage^a, Yash Raj Shrestha^b and Georg von Krogh^a

^aDepartment of Management, Technology, and Economics, ETH Zurich, Zurich, Switzerland; ^bFaculty of Business and Economics (HEC), University of Lausanne, Lausanne, Switzerland

ABSTRACT

Artificial intelligence (AI) applications have proliferated, garnering significant interest among information systems (IS) scholars. AI-powered analytics, promising effective and low-cost decision augmentation, has become a ubiquitous aspect of contemporary organisations. Unlike traditional decision support systems (DSS) designed to support decisionmakers with fixed decision rules and models that often generate stable outcomes and rely on human agentic primacy, AI systems learn, adapt, and act autonomously, demanding recognition of IS agency within AI-augmented decision making (AIADM) systems. Given this fundamental shift in DSS; its influence on autonomy, responsibility, and accountability in decision making within organisations; the increasing regulatory and ethical concerns about AI use; and the corresponding risks of stochastic outputs, the extrapolation of prescriptive design knowledge from conventional DSS to AIADM is problematic. Hence, novel design principles incorporating contextual idiosyncrasies and practice-based domain knowledge are needed to overcome unprecedented challenges when adopting AIADM. To this end, we conduct an action design research (ADR) study within an e-commerce company specialising in producing and selling clothing. We develop an AIADM system to support marketing, consumer engagement, and product design decisions. Our work contributes to theory and practice with a set of actionable design principles to guide AIADM system design and deployment.

ARTICLE HISTORY

Received 11 August 2022
Accepted 6 March 2024

KEYWORDS

AI-augmented decision making; artificial intelligence; decision support systems; design principles; action design research

1. Introduction


The past decade has witnessed rapid growth in the deployment of AI-based systems within organisations, creating human-AI ensembles (Choudhary et al., 2023, Rai et al., 2019, Van den Broek et al., 2021). A key application of AI-based systems in organisations is for augmenting decision making with AI-generated predictions and insights – a phenomenon referred to as AI-augmented decision making (AIADM) (Keding & Meissner, 2021, Raisch & Krakowski, 2021). Interest in AIADM is fuelled by AI's capacity to mine complex patterns from large volumes of data and generate accurate predictions, which when combined with human judgement and decision making can generate value for organisations (Shrestha et al., 2019). Consequently, AIADM systems are increasingly emerging as a prominent subclass of decision support systems (DSS) (Arnott & Pervan, 2014, Hevner & Storey, 2023, Rai, 2016). With recent advancements in deep learning architectures (Shrestha et al., 2021) and large language model-powered generative AI (Dasborough, 2023), this uptrend shows no signs of stagnation.

DSS have been instrumental in organisations as IS artefacts, delivering significant benefits by supporting communication, data processing and knowledge

management, and the construction of decision models, thereby aiding decisionmakers in problem identification, process execution, and decision making (Arnott & Pervan, 2014, Power, 2001). With the ever-increasing amount of data accessible to organisations, AI-powered DSS promise significant organisational value by boosting human productivity, reducing coordination costs and enhancing decision-making speed and accuracy (Brynjolfsson & McElheran, 2016, Lebovitz et al., 2022, Shrestha et al., 2019, Tinguely et al., 2020).

Given these developments, IS researchers endeavour to design effective DSS that empower humans and AI to jointly make decisions, creating superior business value for organisations while also producing societal impact (Fang et al., 2021, Gregor & Hevner, 2013, Padmanabhan et al., 2022, Samtani et al., 2021). Furthering this objective, design science research (DSR) attempts to create novel artefacts that solve previously unresolved issues or improve existing solutions (Hevner et al., 2004). Action design research (ADR) adopts a DSR approach wherein an IS artefact is constructed within a specific client context to glean prescriptive design knowledge to address a class of problems (Iivari, 2015, Maedche et al., 2021,

CONTACT Savindu Herath Pathirannehelage  sherath@ethz.ch

 Supplemental data for this article can be accessed online at <https://doi.org/10.1080/0960085X.2024.2330402>.

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

Mandviwalla, 2015). Recent calls for design research on human-AI systems underscore its practical utility in addressing practitioner problems and its potential to enhance our comprehension of AI technology (Padmanabhan et al., 2022, Rai et al., 2019).

The challenge is that DSR has thus far primarily concentrated on the formulation of design theories (Miah et al., 2019) based on conventional (i.e., non-AI-based) DSS artefacts (Golovianko et al., 2022, Pan et al., 2021), overlooking increasingly important AIADM systems. Recent research shows the extrapolation of prevailing prescriptive design knowledge from conventional DSS to AIADM systems faces four key design challenges (Hevner & Storey, 2023). First, while conventional DSS employs fixed, predefined decision rules and models, mostly leading to deterministic outcomes, AI algorithms learn from data and adapt over time, resulting in stochastic outputs (Padmanabhan et al., 2022). Such outputs influence the decisions within organisations markedly differently compared to conventional DSS, demanding oversight and careful evaluation. For instance, biases and errors in AIADM systems are comparatively difficult to expose and handle (Shrestha et al., 2021). Second, ambiguity in agency within AIADM systems gives rise to uncertainty around decision making authority, autonomy, responsibility, and accountability in organisations (Abdul et al., 2018, von Krogh et al., 2021). As a result, AIADM systems may face significant organisational resistance driven by the perceived loss of managerial control and costly organisation-wide transformations (Feuerriegel et al., 2022). Third, AIADM systems are increasingly raising regulatory and ethical concerns that extend far beyond those associated with conventional DSS (Berente et al., 2021, Mikalef et al., 2022). Finally, conventional DSS exhibits limited configurability and contextual sensitivity (Arnott & Pervan, 2008, Miah et al., 2019), raising concerns about the transferability of design knowledge across decision contexts. Practitioners often find extant DSS research irrelevant due to a lack of configurability and contextual adaptability, demanding the incorporation of contextual idiosyncrasies and domain knowledge in DSS design (Arnott, 2006, Arnott & Pervan, 2014).

The markedly distinct characteristics of AIADM compared to traditional DSS (see Appendix A) influence various facets of organising (Abbasi et al., 2016, Bailey et al., 2022, Baird & Maruping, 2021) and underscore the need for new prescriptive design knowledge incorporating contextual idiosyncrasies (Miah et al., 2019) and practice-based domain knowledge (Padmanabhan et al., 2022). Therefore, *how to design AIADM systems* while navigating unique organisational challenges remains an important open question for design theory and practice (Abbasi et al., 2016, Hevner & Storey, 2023, Padmanabhan

et al., 2022). This dearth of design knowledge on AIADM systems further aggravates the already polarised academic debate around the bright and dark side of AI in decision making (Mikalef et al., 2022). Conflicting expectations and assumptions about the use of AI may partly stem from the lack of in-depth examination of AI artefact design or limited first-hand understanding of how decision-making processes unfold in human-AI ensembles. Further, lack of prescriptive knowledge also limits practitioners' capacity to apply AI in decision making (Padmanabhan et al., 2022). Popular accounts show that despite the technological superiority of AI algorithms, challenges in adopting them could result in AIADM being dismissed altogether (Joshi et al., 2021, Ransbotham et al., 2019).

Against this background, the current study examines the following research questions:

- (1) What are the challenges involved in designing and deploying AIADM systems in organisations?
- (2) What are the principles for designing AIADM systems in organisations?

To this end, we design, deploy, and evaluate an AIADM system consisting of three AI use cases in a young online fashion retailing company (TBô Clothing) serving a global customer base (Iivari's (2015) strategy 2). TBô envisaged AI design, development, and deployment and effectively and efficiently using data-driven insights, as core components of its strategy to operationalise a business model of co-creating products as a community-led brand. The AIADM system augmented decisions across three pivotal domains: a) customer segmentation and targeting, b) customer retention, and c) redesign of the product and service portfolio.

We use ADR (Sein et al., 2011) to design an AIADM system at TBô. ADR applies an iterative build-intervene-evaluate process covering key IS design and deployment stages within organisational contexts (Peffer et al., 2018). We critically examine the three AI use cases, drawing on a rich set of primary data, including interviews with the firm's executive members and data scientists; archival data related to sales and customer feedback, the website, and corporate presentations; field notes from weekly meetings; and experiments in TBô's customer portals and customer surveys. The AIADM system deployment is followed by evaluation, reflection and learning and systematising the research-based design knowledge as a set of design principles (Chandra et al., 2015), advancing the IS literature on AIADM systems design and deployment. Our design principles guide IS practitioners to successfully overcome the challenges of transforming to AIADM.

2. Background

2.1. Decision making in organisations

Within the decision theory literature, scholars have explored decision-making processes from different vantage points (e.g., individual and collective). The two primary approaches to decision making include: a) following the logics of preferences and expectations (March, 1994, Schoemaker, 1982) and b) following appropriateness, obligation, rules, and routines (March & Olsen, 1989, March & Simon, 1993). Scholars adopting the former approach have traditionally assumed that choices are innately rational given perfect information espoused by neo-classical economic theory (March, 1994). However, later scholars, following the Carnegie school, questioned the information completeness and perfect rationality assumptions by introducing the concept of bounded rationality – limitations of human information processing resulting in satisfactory, rather than optimal, decisions (Simon, 1960). Besides bounded rationality, Simon's (1960) major conceptual contribution – the decision-making phase theorem – identified three iterative and recursive phases in managerial decision making: intelligence, design, and choice. Building on Simon's (1947, 1960) seminal works, a subsequent body of research has examined how organisations process information – information processing view of the firm (Galbraith, 1974) – and integrate it into decision-making processes (Joseph & Gaba, 2020, Tushman & Nadler, 1978).

Drawing on the concept of bounded rationality, Tversky and Kahneman (1974) advanced decision theory towards a post-Simon orthodoxy by developing an array of empirically validated theories about the cognitive processes involved in decision making. They put forth theoretical explanations for systematic failures of human decision making, emphasising human biases and judgement heuristics (the prospect theory). This theory suggests that due to humans' limited information processing capabilities, they apply simplifying heuristics to complex decisions typified by uncertainty and opt for satisficing actions that deviate from the rational optimal alternative. This also applies to decisions within organisations, which are made under uncertainty and are fraught with biases and heuristics (Remus & Kottemann, 1986). Acknowledging these systematic decision-making failures, the follow-up work focused on designing corrective actions to improve decision making using statistical methods (Grove & Meehl, 1996, Grove et al., 2000) and technology (Huber, 1990, Molloy & Schwenk, 1995). With the advent of mainframes and computers and their mainstream adoption, DSS emerged as a prominent sub-field of IS scholarship aimed at facilitating and improving decision making in organisations (Arnott & Pervan, 2005).

2.2. Decision support systems

Prior work on behavioural decision making laid the theoretical foundations for DSS development and research. The seminal works in the Carnegie school (Cyert & March, 1963, Galbraith, 1974, Simon, 1947, 1956, 1960, Tushman & Nadler, 1978) provided the key theoretical constructs to examine the influence of DSS on organisational decision-making processes, decision outcomes, and decision performance (Huber, 1990, Leidner & Elam, 1995, Molloy & Schwenk, 1995, Sharma et al., 2014). Arnott and Pervan (2014) found strong evidence for a shift in decision-making orthodoxy in DSS research from a behavioural view characterised by bounded rationality and satisficing to prospect theory characterised by human biases and judgement heuristics. Following this shift in decision theory, DSS research became increasingly grounded on Tversky and Kahneman's (1974) theory (e.g., Chen & Koufaris, 2015).

The DSS literature has advanced both in terms of general theory and design theory. With respect to general theory, Huber (1990) developed a theory on the effects of computer-assisted communication and DSS on organisational design, intelligence, and decision making. A large body of research examined DSS use (Abouzahra et al., 2022, Kamis et al., 2008), development (Lynch & Gregor, 2004), and the impact on decision-making behaviour, capturing both advantages (Barkhi, 2002, Lilien et al., 2004, Todd & Benbasat, 1999) and challenges (Chen & Koufaris, 2015, Giermindl et al., 2022, Mikalef et al., 2022, Rinta-Kahila et al., 2022). In terms of design theory, Keen's (1980) adaptive design framework for DSS development has been highly influential as a kernel theory for subsequent design studies (Miah et al., 2019). According to Keen (1980), not all uses of DSS can be stipulated during the design phase, but a design evolving through use can overcome the foibles of the seminal Gorry and Morton (1971) framework that assumed a static and technical view on DSS design. Following the prospect theory (Tversky & Kahneman, 1974), DSS design aimed to mitigate humans' cognitive limitations by focusing on enhancing both primary (action selection) and secondary (protocol selection) decisions (Arnott, 2006, Remus & Kottemann, 1986). Recent work in DSS literature has shown a significant rise in DSR (Arnott & Pervan, 2014), developing design theories (Miah et al., 2019) and contextual DSS artefacts often featuring in European IS (Collins et al., 2010, Golovianko et al., 2022, Klör et al., 2018, Pan et al., 2021, Seidel et al., 2018).

With the advent of big data and machine learning (ML), the focus in IS artefacts – including DSS (Arnott & Pervan, 2014) – has shifted from systems, functions, features/requirements, and technology towards deriving knowledge and insights from data (i.e., information

and analytics) (Abbasi et al., 2016). Data-driven decision-support solutions – business analytics, business intelligence, and big data analytics – have been on the rise, generating significant rejuvenation of DSS research and practice (Arnott & Pervan, 2014, Rai, 2016). ML-based AI technologies can extract business relevant patterns from large volumes of data (Ågerfalk, 2020, Berente et al., 2021). This is arguably the most significant movement within the history of DSS, resulting in a substantial footprint of AI across technological, organisational, economical, and social domains simultaneously (Berente et al., 2021, Jain et al., 2021, Rai et al., 2019, von Krogh, 2018). Thus, AI-based DSS are gaining prominence as a sub-class of DSS. The resulting human-AI hybrids have the potential to shape decision making on a spectrum ranging from automation (AI substitutes humans) to augmentation (AI and humans complement each other) (Rai et al., 2019, Raisch & Krakowski, 2021). On this spectrum, we chose decision augmentation and positioned our study within the class of AIADM systems where AI insights enhance managerial decisions.

2.3. AI-augmented decision making

AI-based IS artefacts learn, adapt, and act with limited or no human intervention (Baird & Maruping, 2021). These aspects of AI challenge the primacy of human agency in organisations while shifting the focus towards recognising IS agency (Ågerfalk, 2020). Hence, AI is not merely a technology that harnesses knowledge and insights from data, but it also spurs paradigmatic shifts in relationships between humans and machines (Ågerfalk, 2020, Lyytinen et al., 2021) and how they relate and co-organise to process information to make decisions (Bailey et al., 2022, von Krogh, 2018). A new literature stream has emerged to study different facets of these human-AI ensembles and resulting augmented intelligence. Lyytinen et al. (2021) proposed the concept of “metahuman systems” as a hybrid of humans and machines learning jointly while mutually reinforcing each other’s strengths. Murray et al. (2020) identified four forms of conjoined agency between humans and technologies and the impact of these agency forms on the evolution of organisational routines. They identified ML as an augmenting technology that (1) increases the degree of a routine’s change, (2) decreases the predictability of a routine’s change, and (3) decreases routine responsiveness. These findings point towards the far-reaching organisational implications of AI and the recognition of AI agency. While the extant IS artefact (including DSS) literature is eloquent on human agency, it gives scant attention to IS agency (Ågerfalk, 2020, Baird & Maruping, 2021). Agentic primacy is ambiguous and fluid in AIADM systems (Baird & Maruping, 2021), but there is limited clarity

about who has responsibility and accountability (Abdul et al., 2018) in decision-making protocol development and ultimate action selection (Murray et al., 2020). For instance, Lebovitz et al. (2022) question the locus of accountability when AIADM systems diagnose patients.

Other striking distinctions between conventional DSS and AIADM systems warrant assessing it as a separate class within DSS. First, in conventional DSS, decision rules are programmed to produce an output based on an input. Such designs involve neither training nor learning, as the decision rules often lead to a definitive output. These systems with predefined rules and models do not learn from data or adapt over time. Due to this, conventional DSS models and outputs are often more interpretable than their AI counterparts, making it easy to understand the reasoning behind specific decisions. AI algorithms, on the other hand, are not programmed to perform a fixed task, but to learn to perform the task from data and adapt over time (Padmanabhan et al., 2022). Second, stemming from the first distinction, AIADM systems are more dynamic, stochastic, unpredictable, and less explainable with respect to operations and outcomes (Shrestha et al., 2021). Therefore, AIADM systems can lead to unintended results, causing significant risks and damages to the organisation. As an antidote, the human-in-the-loop literature endorses the presence of humans in ML workflows to identify instances where systems might fail, assess associated risks, and develop contingency plans to mitigate risks (Grønsund & Aanestad, 2020, Xin et al., 2018). Third, as opposed to traditional DSS, AIADM systems may face intensive organisational resistance due to the perceived loss of managerial authority, their opaque and complex algorithms that transcend managerial intuition, their unquantifiable economic benefits, and the fact that they trigger swift, organisation-wide changes (Feuerriegel et al., 2022).

Moreover, extant DSS research and artefacts suffer from a lack of configurability and contextual sensitivity (Arnott & Pervan, 2008, Miah et al., 2019). This casts doubt on whether design knowledge can be applied across different decision contexts, problem domains, and underlying technologies. There have been repeated claims that practitioners find DSS research and artefacts irrelevant as they fail to meet the practitioners’ needs (Arnott, 2006, Arnott & Pervan, 2014). The lack of configurability and contextual dynamism across different application domains, domain-specific languages, and different enabling technologies impedes the wider adoption of DSS knowledge contributions (Miah et al., 2019). Therefore, in the current study, capturing contextual idiosyncrasies and practice-based domain knowledge in AIADM system design is crucial to ensure practitioner relevance and acceptance.

For these reasons, we argue that extrapolating the extant prescriptive design knowledge from conventional DSS to novel and context-specific AIADM systems is contentious. AIADM systems are an emerging phenomenon, distinct from traditional DSS (Abbasi et al., 2016, Baird & Maruping, 2021), and their potential to shape multiple aspects of organising simultaneously (Bailey et al., 2022) calls for novel prescriptive knowledge contributions to the design of AIADM systems while remaining alert to contextual idiosyncrasies (Miah et al., 2019) and contemporary problems in practice (Padmanabhan et al., 2022). Little can be designed a priori, but instead these systems need to be rapidly adjusted to the specific client context (Iivari, 2015). In this study, we contribute to the design knowledge on AIADM in organisations by describing the design and deployment of an AIADM system in a specific context.

3. Research context and methodology

3.1. Case selection

We selected TBô Clothing (<https://tbo.clothing/ch-en/>), a globally operating online fashion retail company headquartered in Switzerland, as our research context. Established in 2019, TBô is a young company with its employees spread across Europe, North America, and Asia. TBô caters to a diverse customer base across three continents, relying exclusively on digital platforms and online stores. It also positions itself as a community-led brand, with its entire product range being co-created using customer input. To do so, TBô has created and maintained an online community where customers can participate as co-creators. Curating ideas on product design and development, TBô routinely (usually weekly) circulates online questionnaires where co-creators can engage and contribute by answering questions on personal information, user experience, product ideas, personal preferences, and personal aspirations.

TBô is a suitable research setting to investigate our research questions. First, TBô envisaged AIADM design and deployment as a core component of its strategy. Second, TBô is a “clean slate” in which we can observe AIADM system design and deployment without much interference from legacy systems, prior routines, decision-making processes, and experience with similar projects. Third, the evolution of a project could be tracked from ideation to deployment from both managerial and operational perspectives. Fourth, the unique technological, organisational, operational, and market conditions of TBô engender certain contextual characteristics and peculiar challenges for AIADM deployment that are congruent with our problem concept. Hence, the design principles we develop can be used to guide AIADM initiatives in similar settings by overcoming the challenges identified.

We also had several practical considerations in choosing TBô, such as the alignment between our research interests and the company’s strategy and vision to leverage AI adoption. TBô’s young age and limited resources, notably in terms of human talent, render it transparent and receptive to collaboration. Consequently, the ADR intervention can be conveniently implemented and meticulously examined.

3.2. AIADM use cases

The strategic data roadmap of TBô officially proposed by the co-founders to all employees and investors highlighted three high-priority AI use cases that centred on three key decision-making areas: (1) customer segmentation and targeting, (2) customer retention, and (3) redesigning the product and service portfolio. While (1) focused on increasing co-creation participation and the efficiency of co-creation campaigns, (2) and (3) aimed to maximise customer lifetime value (CLV) and subsequently sales revenue. Given TBô’s digital business model, rich accumulated customer data enabled the development of AIADM to identify relationships between purchasing and co-creation participation to purposefully nudge the customer community to purchase as well as co-create. TBô anticipated that AIADM would guide the decisionmakers in designing and implementing meaningful interventions for enhancing both sales and co-creation.

We conducted the end-to-end process of AIADM system design and deployment, from when the co-founders first conceived of the idea to adopt AIADM to the final implementation and the company’s post-hoc evaluation of the system. At TBô, AIADM was built on online purchase (order), survey (co-creation), and advertising and promotional campaign data they collected to inform managerial decisions. Before adopting AIADM, TBô relied on manual reading and coding of textual data collected via online surveys to identify and evaluate prominent, attractive, and lucrative ideas and integrate a subset of them into products. Subsequently, based on analysis of the number of pre-orders they received, decisions were made on whether to promote products in the core collection or as limited editions. In essence, the AIADM system at TBô aimed to augment managerial decision making in operationalising its co-creation business model.

3.3. Action design research

We adopt the ADR methodology for three reasons. First, AIADM system design and deployment in an organisation comprises the “inseparable and inherently interwoven activities of building the information technology (IT) artefact, intervening in the organisation,

and evaluating it concurrently” (Sein et al., 2011, p. 37), which aligns with the research process conceptualised in ADR. Extant literature attests that the ADR approach not only fosters richer insights into the interactions of technology and organisation (Altendeitering et al., 2021, Ebel et al., 2016, Sun et al., 2019), but also performs a dual mission of contributing to theory and providing practical insights (Sein et al., 2011). Second, ADR builds on the premise that IS artefacts are *ensembles*: a collection of software/hardware tools, shaped by the organisational and technological context during development and use (Sein et al., 2011, Sun et al., 2019). Relatedly, AIADM represents multiple software/hardware systems and is embedded within the organisational context (Shrestha et al., 2019). Finally, ADR facilitates a dynamic and flexible research process, cycling between building the IS artefact and evaluating its utility (Sein et al., 2011). The artefact emerges through the contemporaneous interaction between (1) design and use and (2) organisational and technological context (Orlikowski & Iacono, 2001), facilitating discovery of both intended and unintended organisational consequences of a specific artefact design and accompanying organisational challenges and mitigating strategies.

Following Sein et al. (2011), we conducted ADR in four stages: (1) problem formulation; (2) building, intervention, and evaluation (BIE); (3) reflection and learning; and (4) formalisation of learning. In each stage, we gathered data from multiple sources (e.g., interviews and field notes) to capture an unbiased, holistic view of AIADM systems design and deployment while navigating through various challenges and trade-offs (see Table 1).

4. Artefact development and evaluation

4.1. Artefact formulation

TBô was driven by an immediate need for automated analysis as the volume of data being collected exponentially increased over time, making their traditional manual analysis impractical. The CEO (I1) expressed the tediousness and slowness of traditional manual efforts of survey analysis:

We used to read every customer survey response by ourselves to find the product ideas. This is impossible with the increasing number of customers and their responses.

Within its co-creation model, TBô had a unique opportunity to accumulate diverse and complementary data about customer interactions through multiple channels (order, co-creation/survey, and campaign data; see Appendix D). Given purely online interactions with customers, digital trace data facilitated an opportunity to identify seasonality and trends in

customer demand and preferences. The online store’s sales data could be exactly and routinely monitored. Data richness and complementarities among the three datasets encouraged TBô to find ways to leverage AI-driven insights to augment decisions such as segmenting and targeting customers, enhancing marketing efforts, and redesigning the product and service portfolio.

Manual data analysis was restrictive in building models that could predict customer journeys, enhance customer engagement, and increase customer repurchasing. Furthermore, manual methods required dedicated organisational roles and employees, increasing labour costs. According to I1, AI-based decision models were critical when competitors such as Zalando and Zara began applying them at scale.

The TBô executive team initially experimented with third-party tools such as Google Analytics, Facebook Business Manager, and Shopify Analytics. Experience and early success with these tools, as well as the CEO’s firm belief that using AI tools could improve the firm’s decision making, became the catalyst to transition to an in-house designed and developed AIADM system. The CEO (I1) succinctly summarised this as follows:

The main advantage of moving from manual to AI tools was the quick summaries and making a nice dataset for us to analyse, making it fast and accurate. Now we also see a big advantage in developing our own tools to further bring co-creation into the community to make it more dynamic.

During artefact formulation, three key challenges, related to a) lack of experience and scepticism, b) managing multiple objectives, and c) competing interests, emerged. First, TBô lacked specialised knowledge of the underlying algorithmic mechanics resulting in initial scepticism of the management team about the possibility to design an AI system to enhance decision making and subsequently create value. Second, we encountered challenges in aligning business problem formulation with AIADM system design. AI algorithms necessitate specific problem formulation, which can be impractical in real-world scenarios, leading to difficulties in aligning AI problem formulation with actual business objectives and metrics. For instance, TBô struggled between its dual business objectives of sales and customer co-creation and in formulating this for AI. TBô operated with a co-creation business model in which all its products were designed and developed based on customer insights. The importance of listening to their customers is exemplified on TBô’s website:

TBô Bodywear is the world’s first DirectByConsumer brand. It’s TBô’s customers—the 400,000-strong Tribe—who decide the brand’s direction and which products get made.

Table 1. Summary of the ADR process at TBó Clothing.

ADR stage	ADR principles	Activities	Actors involved	Data collection and analysis	Artefact
Problem formulation	<p>Practice-inspired research (P1): The research was driven by (1) a vision to improve decision making on CLV and co-creation participation by leveraging AI-driven insights and (2) a lack of prescriptive knowledge about AIADM system design, and deployment.</p> <p>Theory-ingrained artefact (P2): We chose the decision augmentation thesis (over automation) of the human-AI ensemble theories (Fügener et al., 2022, Murray et al., 2020).</p> <p>Reciprocal shaping (P3): The technological context and organisational context mutually influenced the AIADM system design and deployment, as the challenges encountered were technical and/or organisational.</p> <p>Mutually influential roles (P4): The ADR team was comprised of researchers and practitioners to capture theoretical, technical, and practical insights. The lead designer (a PhD student) worked full time on developing AIADM tools with TBó.</p> <p>Authentic and concurrent evaluation (P5): The ADR team concurrently evaluated whether the AIADM system delivered business value through field experiments and by interviewing the executive users.</p> <p>Guided emergence (P6): The AIADM system was developed and implemented in the firm so that the deployment process would not only reflect the technical design (P2) but also undergo ongoing shaping by organisational use (P3), participant perspectives (P4), and concurrent evaluation (P5).</p>	<ol style="list-style-type: none"> 1. Understanding the co-creation business model, types and sources of data, the processes, and the systems at TBó 2. Reviewing the literature on AIADM system design and deployment 3. Formulating research questions 1. Developing data-driven decision models to augment customer intervention decisions at TBó 2. Implementing recommendations of the developed decision models in a real-world business context 3. Evaluating whether the AIADM system delivered improvements in terms of sales and co-creation input 	<p>Researchers, CEO, COO, data engineer</p>	<ol style="list-style-type: none"> 1. Interviews with CEO (11; 60 min; see Appendix B) 2. 16 Weekly meetings with CEO/COO and data engineer 3. Exploratory data analysis on archival data 4. Company data workshop presentation (21 slides) 5. Company website 6. Literature review (see Section 2) 	<p>Recognition: 1. Understanding the problem and solution domain as “AIADM system design and deployment” 2. Shortcomings of traditional decision-making approaches in business practice. Identifying the factors driving AIADM adoption 3. Absence of prior prescriptive knowledge and empirical research on design and deployment of AIADM systems</p> <p>Alpha version: Three AIADM models for identified use cases (see Section 3.2) as proof of concept (POC)</p> <p>Beta version: Upon POC, recommendations from AIADM models were deployed to evaluate proof of value (POV), proof of use (POU), and design activities</p>
Reflection and learning	<p>Authentic and concurrent evaluation (P5): The ADR team concurrently evaluated whether the AIADM system delivered business value through field experiments and by interviewing the executive users.</p> <p>Guided emergence (P6): The AIADM system was developed and implemented in the firm so that the deployment process would not only reflect the technical design (P2) but also undergo ongoing shaping by organisational use (P3), participant perspectives (P4), and concurrent evaluation (P5).</p>	<ol style="list-style-type: none"> 1. Recognising the challenges encountered in AIADM system design and deployment 2. Realisation of design activities while navigating through the challenges identified in design and deployment <p>Proposing a set of design principles for AIADM system design and deployment, positioning TBó as an instance/case</p>	<p>Researchers, CEO, COO</p>	<ol style="list-style-type: none"> 1. Meetings with CEO and COO. Literature review to rationalise the design principles³. Design principle reusability evaluation (livari et al., 2021) 	<p>Emerging version and realisation: The revised version of the AIADM system, specific challenges in each phase emerging from organisational and technological context, and a reflection on how these challenges were addressed through design activities</p>
Formalisation of learning	<p>Generalised outcomes (P7): The resulting ensemble from ADR was a solution to a problem. We conceptualised the problem (P1) and our solution (P6) as an instance of the class to generalise our findings. In our study, this class is “AIADM system design and deployment”.</p>	<p>Proposing a set of design principles for AIADM system design and deployment, positioning TBó as an instance/case</p>	<p>Researchers, CEO, COO</p>	<ol style="list-style-type: none"> 1. Meetings with CEO and COO. Literature review to rationalise the design principles³. Design principle reusability evaluation (livari et al., 2021) 	<p>Ensemble version: Ensemble embodying an AIADM system for customer interventions, challenges encountered, and design principles addressing these challenges</p>

Although TBô benefited from design ideas from its customers and sales growth, it remained unclear how to weight these two related but distinct objectives in the concrete objective function that AI requires. To circumvent this challenge, we initiated AIADM with two separate objectives (instead of aggregated objectives) stemming from AIADM use cases (see Section 3.2): to (1) increase co-creation participation and (2) improve the CLV and subsequently sales revenue. We designed for synergy in input (data, domain knowledge), AI model, and the output (predictions, visualisations) iteratively to derive useful insights for above objectives.

Third, to schedule resource utilisation, at the inception of operations, TBô concentrated on investments that were likely to immediately strengthen their business model (e.g., building website infrastructure, setting up co-creation channels, marketing and promotions) and AIADM transformation was considered a second step.

The CEO took charge of championing the change. This required a transformation in decision-making structures, reporting hierarchy, and data management practices, which induced uncertainty in the organisation. To garner support and prevent strong risk aversion amongst employees, the CEO formulated a concrete data roadmap. The roadmap outlined short-, medium-, and long-term goals of AI design and deployment and thus formed a concrete and actionable object for curating organisational support and trust, resulting in enhanced coordination (see Figure 1).

The CEO (I1) highlighted the benefits of a clear roadmap as follows:

The data roadmap that we used in the workshop in mid-September with an overview of our business provided confidence [to employees and investors] in our data-driven approach going forward.

4.2. Artefact development

Following the second principle of developing a theory-ingrained artefact (P2), we drew on the extant human-AI ensemble literature, which is germane to our class of systems. This literature examines how to integrate decisions involving humans and AI while recognising the agency of AI artefacts (Baird & Maruping, 2021, Murray et al., 2020, Shrestha et al., 2019). Research addressing this fundamental question of human-AI ensemble decision making coalesces into two major conceptualisations: decision automation and augmentation (Raisch & Krakowski, 2021). Raisch and Krakowski (2021) defined automation as machines substituting humans, whereas augmentation refers to humans collaborating with AI in making decisions. Recent work has evidenced the superiority of the augmentation theory, citing improved decision-making performance. Fügner et al. (2022) found that humans and AI working collaboratively can outperform the AI that outperforms humans when they work independently. However, the combined performance improves only when the AI delegates work to humans – not when humans delegate work to the AI. Bouschery et al. (2023) found that AI can augment human innovation teams by fostering divergent processes to explore wider problem and solution spaces in new product development. These crucial discoveries align with the augmentation theory and our class of

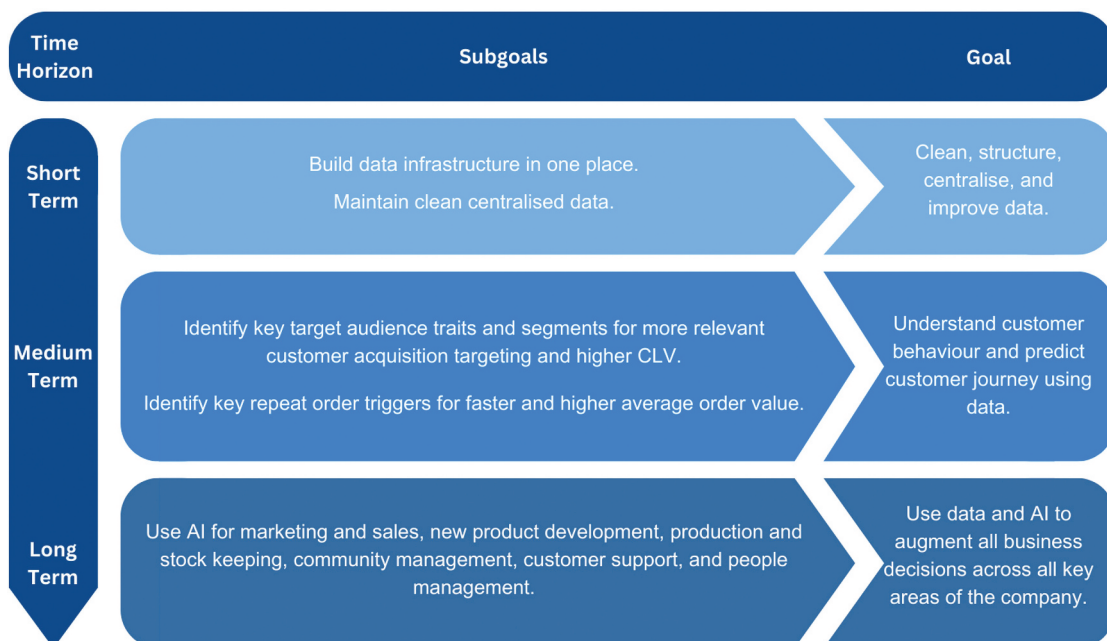


Figure 1. Goals of TBô as extracted from the company's data roadmap.

AIADM systems, where AI parses large amounts of data, detects patterns therein, and provides recommendations, while humans assume responsibility over decision and action selection. Therefore, we rely on augmentation theory.

Augmentation triggers a partial shift from human-driven to AI-supported decision making, in which AI systems provide recommendations (output of AIADM systems) for humans to act. This ensures the involvement of humans without losing the characteristics of decision making, such as responsibility, accountability, context specificity, and utility expectations (value seeking). Thus, we formulate the initial design principles (DPs), prescriptive statements to constitute the basis of the design actions (Chandra Kruse et al., 2016), as context specificity (Miah et al., 2019), utility (Sein et al., 2011), and responsibility (Mikalef et al., 2022), while keeping human involvement (Van den Broek et al., 2021) as the primary and fundamental design principle.

Following the fourth principle of mutually influential roles (P4), we emphasised learning and cross-fertilisation between the research team members and the TBô executives by enabling a combination of academic insights with domain knowledge from industry and practice (Sein et al., 2011). The lead designer (the first author) worked full time on developing the AIADM tools with TBô and interacted regularly with TBô staff in weekly meetings (see Table 1 and Appendix D). The co-authors had multiple roles, including facilitating the technical development; managing the research partnership; and undertaking the organisational and theoretical introspection, synthesis, reflection, and learning. The CEO and data engineer facilitated the ADR procedures by contributing their practical experience (see Figure 2).

The artefact development consisted of business and data understanding, followed by AI modelling and validation.

4.2.1. Business and data understanding

AIADM is the confluence of insight from data (exploration/induction) and the domain expertise of decision-makers (Agrawal et al., 2018, Tarafdar et al., 2019). Managers' experience and their understanding of consumer behaviour and products were necessary for the AIADM system design process. Transferring adequate domain expertise to data scientists to work on the problem(s) was crucial. This domain knowledge

transfer helped the data scientists to better understand what the business problems/tasks are and to formulate those into an objective function that an ML algorithm can comprehend (see Section 3.2 for AI use cases). The ADR team exchanged domain knowledge (e.g., about the co-creation business model and its performance metrics) with the data scientists in several collaborative sessions and meetings (see Table 1) which helped in formulating evaluation criteria for the effectiveness of AIADM (see Section 4.3). Although at the beginning of the process, there were misunderstandings (e.g., data scientists lacked experience with the co-creation model and its performance metrics), after several discussions, the team members converged on a common language and found a way to further collaborate.

The data science team took significant steps in describing, exploring, and verifying the quality of the data. This included descriptive statistics, visualisation, assessing data quality, and discussing potential use cases with the domain experts. We found that data understanding and business understanding benefitted from many iterations between the domain experts and data science team.

4.2.2. AI modelling

First, data was prepared for AI modelling following standard steps such as removing redundant features, feature engineering, and treatment of missing values. Pre-processing mechanisms such as feature selection and reweighting were used to debias training datasets before feeding them into learning algorithms (Kordzadeh & Ghasemaghaei, 2022). We describe data preparation and subsequent training and testing of predictive ML models in Appendix E. To understand the relationship between purchasing and co-creation behaviour in TBô's business model, three decision models (DMs) were developed (see Figure 3):

4.2.2.1. DM1 for customer segmentation and targeting. DM1 aimed at predicting co-creation behaviour based on purchasing behaviour. Observing purchasing behaviour, this ML model guided segmentation of the customer base and subsequently targeting promising segments with interventions aimed at increasing co-creation participation.

4.2.2.2. DM2 for customer retention. DM2 aimed at identifying the difference in purchasing behaviour

ADR team			
Domain experts		Data science team	
Researchers	Practitioners	Researchers	Practitioners
Researcher 1	CEO	Lead designer	Data engineer
Researcher 2			

Figure 2. ADR team.

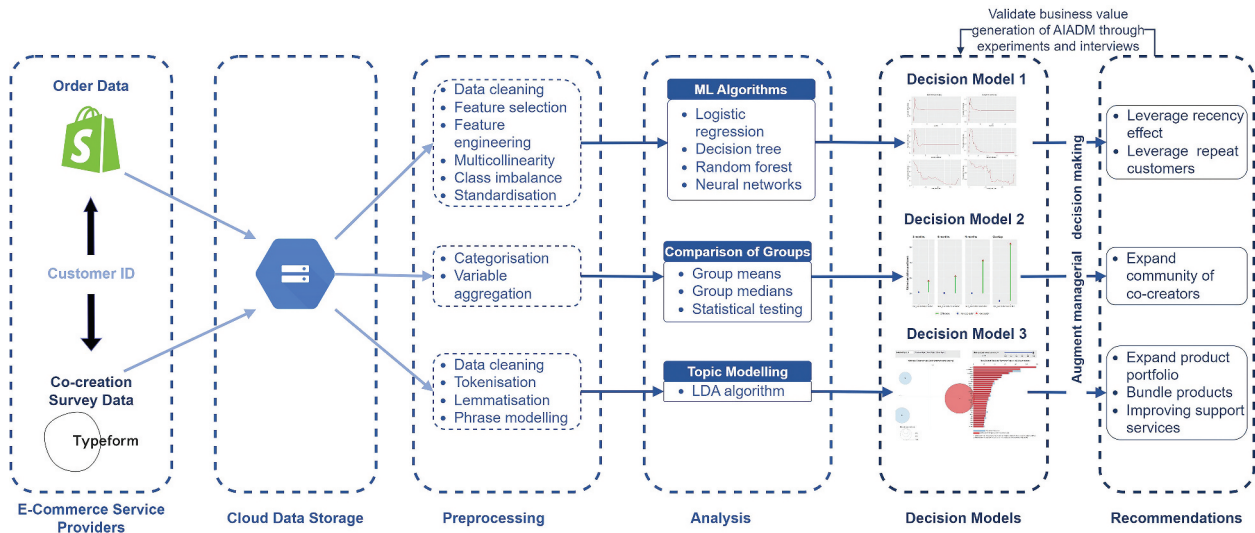


Figure 3. Architecture of the AIADM system at TBô.

between co-creators and non-co-creators by comparing their purchases. This model guided retention of lucrative customer segments.

4.2.2.3 DM3 for redesigning product portfolio and services. DM3, via topic modelling, focused on identifying salient product and service issues that customers raise as reasons not to place repeat orders. It guided TBô to redesign their product and service portfolio.

4.2.3 AI model validation

Model validation appraised the predictive performance of the models built. For DM1, a collection of three metrics was used to validate the predictive performance of the models, namely accuracy, mean log-loss score, and area under the receiver operating characteristic curve (AUC). The best predictive performance appeared with the random forest model. For DM2, we used standard statistical testing. For DM3, the coherence score was used as the performance metric to choose the best hyperparameter combination in our grid search. The coherence score is a measure of semantic similarity between words within a topic, and it measures the quality of the generated topics.

We observed that model validation also required collective decisions from managers and data scientists, such as deciding on accuracy metrics, model selection, and interpreting the topics in topic modelling. By integrating data science knowledge with business expertise, the ADR team was able to significantly improve the predictive performance of the models (see Appendix E) by demonstrating the effectiveness of our artefact (POC) (Venable et al., 2012, Nunamaker et al., 2015).

Three key challenges related to a) resource constraints, b) data constraints, and c) technological

constraints were identified during artefact development. First, TBô faced constraints in talent, capital, and time while developing AIADM. It struggled to fulfil its vacancies requiring specialist technical and business domain expertise. The following quote from I1 highlights the CEO's earnest search for experience and expertise in AI-related technology:

The main bottleneck [in adopting AIADM] is the lack of engineers and data scientists to develop the tech and algorithms.

This statement was corroborated by the job advertisements posted on the company website, which failed to attract suitable applicants for more than nine months (see Figure 4).

Moreover, IT infrastructure to store and analyse data is costly and time consuming to install. The two co-founders (CEO and COO) of TBô also found it difficult to dedicate time and managerial attention to implement AIADM while managing their day-to-day business operations.

Second, TBô faced several data challenges in the adoption of AIADM. To pursue the AIADM journey, the firm needed to capture, clean, and store data, because data is the centrepiece of AIADM (Kuguoglu et al., 2021). AIADM initiatives that fail to capture and store clean and relevant data are destined to be error prone and hence unsuccessful (Kuguoglu et al., 2021). By “relevant”, we mean that the data can give meaningful insights to solve the problem(s) under consideration. Data management is challenging, especially the crucial tasks of collecting, cleaning, and storing data. TBô used third-party service providers to run email and social media campaigns to collect co-creation data (see Appendix D). To accumulate order data, TBô leveraged a third-party proprietary e-commerce platform for online stores and retail point-of-sale systems. These external service providers offer

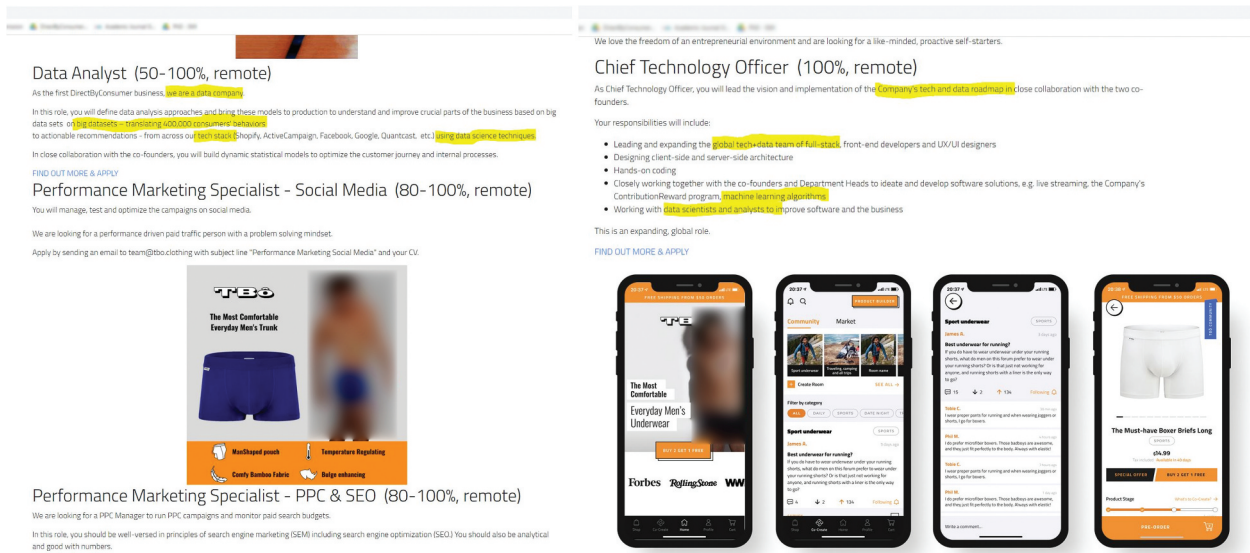


Figure 4. AI-related job advertisements on TBô website.

interaction platforms, methods to collect data in real time, data analysis tools, trouble-free integration with the firm’s internal systems, and storing this data in (cloud) data storage systems via dedicated application programming interfaces. The company experienced initial challenges in curating complementary datasets, such as finding behavioural data for customers, since order data only contained limited behavioural information. TBô thus relied on surveys to collect behavioural data. Combining these three datasets – order, co-creation, and campaign data – was challenging, as the process demanded a unique customer identifier (e.g., email address) across datasets. If customers used different emails, joining the datasets became inefficient, creating multiple copies of the same customers in a fragmented form.

Third, systems and technologies, collectively called the “technology” or “solution stack”, are leveraged for diverse business tasks in organisations. Even though TBô tried to integrate these systems running on heterogeneous technologies through standard interfaces, they often found the interlinking difficult, while also failing to achieve desired results. The CMD highlighted this issue (I3):

Another issue is interlinking all the different software seamlessly and having it all in one software or dashboard—that is, combining email, SMS, social media, the website, and other outlets.

4.3. Artefact deployment and evaluation

The decision recommendations derived from the developed decision models were deployed. It is important to mention that diverse recommendations were identified. In Table 2, our intention is not to provide a comprehensive list of all the recommendations of the decision models, but to elucidate the deployment with a few examples. By doing so, we demonstrate how an organisation envisaging transformation to AIADM may replicate a similar approach.

Best-validated ML models based on suitable performance metrics such as accuracy, mean log-loss score, and AUC were deployed (see Section 4.2, model validation). However, it is likely that the model’s accuracy might not align with the additional value being generated by AIADM. In artefact evaluation, we thus examined the expected gains from our AIADM system (P5).

Table 2. Recommendations stemming from decision models.

Decision model 1	Decision model 2	Decision model 3
<p>R1. Leverage the recency effect of sales on co-creation Insight: The probability of co-creation drops considerably as the time from the last purchase increases. Hence, segment the customers depending on the recency of the purchase and target the recent customers with co-creation surveys.</p> <p>R2. Leverage repeat customers to optimise co-creation initiatives Insight: Co-creation initiatives can be made more efficient (cost per response, etc.) if they target repeat customers.</p>	<p>1. Expand community of co-creators Insight: Co-creators account for higher CLV than non-co-creators. Managers should integrate co-creation initiatives into their business models and incentivise co-creation.</p>	<p>From topic modelling, we identified three latent topics that hamper repeat orders: (T1) no need to buy, (T2) high cost, and (T3) dissatisfaction with service. We present one recommendation for each topic identified.</p> <p>T1: Expand product portfolio to cater to different product needs T2: Product bundling T3: Improve support services and customer inquiry handling</p>

We conducted field experiments to evaluate the model's actual benefits (POV) and interviewed the responsible stakeholders of the organisation to identify both desirable and undesirable consequences of its use (POU) at TBô. We adopted the DSR evaluation approach proposed by Venable et al. (2012) and Nunamaker et al. (2015) and applied by Tuunanen and Peffers (2018), Nguyen et al. (2021), and Golovianko et al. (2022).

As discussed above, we created a set of recommendations from each decision model. We executed the two recommendations from DM1 (R1 and R2 in Table 2), which were then evaluated in the field. Specifically, we created the treatment groups with the new customer segments suggested by our AIADM system while the control groups were predefined by TBô. The resulting field experiment returned co-creation survey response rates – the performance metric reflecting co-creation participation – of 1% and 4.4% for treatment groups compared to 0.1% and 0.2% for control groups. Interviews I2 (CEO) and I3 (CMD) confirmed that treatment groups' survey response rates are significantly superior compared to what was previously observed in the company and the industry in general. In conclusion, the results of the experiment confirm that the selected recommendations derived from DM1 (R1 and R2) are effective in delivering significant gains in co-creation participation.

After this experiment, the ADR team conducted meetings with TBô to evaluate the effectiveness of the topic model (DM3). The CEO confirmed that these topics are highly relevant, and they find great value in such a topic model to extract insights hidden in large amounts of textual data they gather from various channels (I2). Current approaches in practice, including manual reading and coding of textual data, were also discussed. Manual text processing had already identified several issues that had some commonality with the topics we found. The CEO expressed the firm's interest in implementing an AI-based automated text analysis tool, especially to analyse customer conversations in the newly implemented social space on their website. Via successful implementations and deployments, and practitioners' intention to extend the use cases, we demonstrated the POC, POV, and POU of our artefact.

We observed several challenges in artefact evaluation, related to a) challenges in experiment design, b) consumer/user engagement and fairness, and c) data shifts. First, technical challenges emanated from implementing experiment design. Experimental settings are widely leveraged to validate the effectiveness of AIADM (Senoner et al., 2022), and they mainly compare the change in the performance indicator between the treatment scenario (AIADM case) against the control scenario (conventional decision-making

methods). We created two groups for each recommendation: a control group and a treatment group. The treatment of our experiment design was an intervention in the customer's purchasing behaviour, that is, making a purchase for R1 and making repeat purchases for R2. However, we could not force the customers to make (repeat) purchases. To overcome this challenge, we categorised existing customers into treatment and control groups using thresholds for days from the last purchase (purchase recency) and the number of orders (purchase frequency).

The second challenge was obtaining enough engagement in the experiment. This is evident from the low response rates of our experimental groups (the maximum response rate was 4.4%). When the participants' engagement was sporadic and sparse, the internal and external validity of the experiment's results suffered. Grouping participants into treatment and control groups also raises ethical and fairness concerns.

Third, covariate shifts, that is, changes in data distribution, pose risks to AIADM system performance. ML algorithms assume stable data-generating processes, and any changes in the underlying process could lead to performance deterioration. For instance, the COVID-19 pandemic unfolded during the study, and it was difficult to disentangle the effects of COVID-19 (e.g., increased remote working, higher e-commerce, higher savings, etc.) on consumer purchasing behaviour and the resulting evaluation of the AIADM system.

4.4. Artefact sustenance

The decision to sustain AIADM at TBô relied on three key aspects: (1) the adoption of AIADM delivered measurable performance gains when juxtaposed with conventional decision-making methods, (2) the benefits exceeded the recurrent costs of the AIADM system, and (3) the general expectation that AIADM improves as the models learn. After observing the performance gains of AIADM over traditional decision-making methods and the benefits that outweigh the costs of AIADM (demonstrated POC, POV, and POU), TBô expressed its eagerness to apply AIADM not only to the business cases considered in this study, but also to other future use cases (see Figure 1, long-term goals).

We identified three pressing challenges in sustaining AIADM, related to a) trust and confidence in AIADM, b) economics of AIADM, and c) managerial over-optimism. First, during I2, I3, and company meetings, the co-founders and team members emphasised the importance of the reliability of the decision models. In other words, AIADM should be reliable for at least a partial delegation of decision making. We observed that

some results obtained by the decision models (e.g., the total purchase value of a customer negatively affected co-creation probability) were thoroughly scrutinised by employees and we observed their scepticism in accepting some of the insights the AIADM system presented. The CMD expressed his concerns about trusting AI (I3):

Some tools give shallow analysis and insights and require more labour or other software to extract the insights we need.

Second, we increasingly recognised that the AIADM system could not be sustained without demonstrating sufficient economic value. During our study, we identified significant recurring labour costs (both for employees and outsourced work) and costs of maintaining AI infrastructure (data storage, computational power, etc.). To this end, the company had to ensure that the benefits of the AIADM system outweighed the costs in the long run.

Third, we found that learning from failures was an integral part of the continuous improvement at TBô. The AIADM initiative should not be viewed as sequential but as cyclical. In essence, AI models are not oracles that provide predictions, but instead cross-learning agents that evolve over time through multiple interactions. Several iterations can pave the way for gradual system improvements over time. Promising AIADM projects can be scrapped when they are audited against utopian managerial expectations. Hence, it was important to set clear objectives with proper business understanding and define thresholds to audit the performance gains of AIADM, keeping in mind that failures can lead to success in subsequent iterations.

5. Prescriptive learning

5.1. Reflection and learning

We noted two important practices that facilitated the guided emergence (P6) of the AIADM system within the organisational and business context: pursuing AIADM system deployment in a real-world uncontrolled corporate context and leveraging a variety of data collection procedures to collect a diverse yet rich dataset to reflect on and learn from. This enabled the identification of an eclectic set of challenges that organisations face from formulation of to sustaining AIADM systems. We call these challenges the “unanticipated outcomes” of our IS artefact. Table 3 provides an overview of the challenges we identified for each phase and our design activities addressing these challenges. As we reflected and learned about anticipated (e.g., strategic data roadmap) and unanticipated outcomes (identified challenges) demanding ongoing changes to the preliminary artefact design, we developed a set of design activities on how the identified challenges can be addressed in our class of systems. Based on our reflections on the design activities in Table 3, we refined our initial design principles and formalised our learning into an expanded set of final design principles.

5.2. Design principles of AIADM systems

The final stage relates to formalisation of the learning. ADR suggests that generalisation of outcomes (P7) has three levels: problem instance, solution instance, and derivation of design principles (Sein et al., 2011). While we deployed an AIADM system within a single organisation, we aimed to extract insights that may extend beyond a single business problem.

Table 3. Overview of challenges and design activities.

	Challenges	Design activities
Formulation	<ol style="list-style-type: none"> 1. Technological ambiguity 2. Managing multiple objectives 3. Organisational and cultural inertia 4. Competing interests 	<ol style="list-style-type: none"> 1. Formulate a strategic data roadmap and communicate it to all the project stakeholders. 2. Identify measurable AIADM use cases. 3. Develop a prototype to prove the business value creation and encourage user acceptance. 4. Make the AI outcomes understandable (e.g., visualisations—partial dependence plots [PDPs] and feature importance).
Development	<ol style="list-style-type: none"> 1. Resource constraints 2. Data constraints—storage, integration, and ethical use 3. Technological constraints 	<ol style="list-style-type: none"> 1. Leverage open and free AI resources (code, libraries, etc.). 2. Foster partnerships (e.g., industry-academia, solution providers). 3. Combine multiple datasets using common fields. 4. Adopt regulatory guidelines (e.g., European Commission, 2021, OECD, 2021). 5. Develop AI guiding principles in-house. 6. Audit AIADM to ensure ethical outcomes. 7. Implement interactive user interfaces with customisable parameters for human inputs.
Deployment and Evaluation	<ol style="list-style-type: none"> 1. Challenges in experiment design 2. Consumer/user engagement and fairness 3. Data shifts—dynamic environment 	<ol style="list-style-type: none"> 1. Identify measurable AIADM use cases. 2. Use flexible quasi-experimental designs. 3. Adhere to AI regulations and the firm’s code of ethics. 4. Maintain domain expert involvement.
Sustenance	<ol style="list-style-type: none"> 1. Trust and confidence in AIADM 2. Economics of AIADM 3. Managerial over-optimism: expectation vs reality 	<ol style="list-style-type: none"> 1. Make the AI outcome understandable and hence trustable, e.g., explainability libraries. 2. Involve humans to correct AIADM errors. 3. Follow a comprehensive evaluation of the AIADM system (POC, POV, and POU). 4. Accept failures and adopt a continuous improvement mindset.

We investigated the general problem of “AIADM system design and deployment in organisations”. The final design principles distilled from the design and deployment of AIADM system at TBô are as follows.

5.2.1. DP1: Design for alignment between the business model and organisational resources and capabilities

Formulation of an actionable strategic roadmap aligning with a company’s business model is crucial for AI-based decision augmentation in firms. Such a roadmap should include and explicitly illustrate (1) measurable and easily interpretable AI use cases; (2) availability of domain expertise associated with identified business cases; (3) technical feasibility of the (proposed) AI tech stack; and (4) clear and concrete goals, sub-goals, timeline, and likely challenges in implementing AIADM. This principle serves three purposes. First, it makes the AIADM system specific to the decision-making and business context, thus increasing practitioner relevance and acceptance (Miah et al., 2019). Second, it coordinates the project communication in line with corporate vision and mission for organisational support. Third, the strategic roadmap helps steer the project in overcoming technological ambiguity and managing diverse use cases with multiple objectives. The TBô data roadmap included all these aspects and thus demonstrates this design principle. TBô’s leadership championed the proposed project with an actionable strategic roadmap, remained accountable for the AI project outcomes, and proactively led the process of gathering employee support.

5.2.2. DP2: Design for synergy in input, model, and output to ensure business value

Once the AIADM use cases were defined, we employed an iterative design approach to synergise input elements (data, domain knowledge), the AI model (ML and natural language processing), and output (predictions, visualisations) of the system to derive meaningful insights for the use cases. To overcome data constraints, we merged multiple datasets. Leveraging these comprehensive datasets alongside existing domain expertise, we evaluated several AI models to identify the best performing models on the chosen ML performance metric (see Appendix E). As an integral part of the output, we included visualisations (PDPs, feature importance) and explanations (SHapley Additive exPlanations [SHAP], Local Interpretable Model-agnostic Explanations [LIME]). Finally, we followed a comprehensive three-pronged evaluation of the AIADM system to demonstrate POC, POV, and POU. This principle serves two purposes. First, establishing synergy in the key components of AIADM ensures accurate and comprehensible recommendations and prevents massive failures. Second, demonstrating

value and usability is key to convincing stakeholders and gaining organisational commitment for additional resources over other competing interests. At TBô, this ensured that the organisation pursued the value-led solid AI use cases rather than blindly following AI-hype-led implementations.

5.2.3. DP3: Design for ethical AI governance frameworks

Accompanying the great promises and possibilities of AI is a host of intricate thorny issues related to security and privacy, fairness, deskilling, surveillance, and accountability (Berente et al., 2021, Mikalef et al., 2022). Organisations are highly susceptible to these perils due to the rudimentary state of the guidelines, inadequate expertise in these guidelines, low institutional support, and the dire need to scale up rapidly (Bessen et al., 2022; Singh et al., 1986). These perils are circumvented by an *AI governance framework*—a set of normative declarations on how AI is developed, deployed, and governed, adhering to legal, ethical, social, and organisational values. Through our ADR study, we offer three pathways to a responsible AI design: (1) Adopt extant regulatory guidelines (e.g., European Commission, 2019, 2021, OECD, 2021); (2) develop own AI guiding principles consistent with customer and user expectations (Bessen et al., 2022; Google, 2022); and (3) establish an AI auditing and governance framework (Grønsund & Aanestad, 2020). AI auditing should evaluate not only potential business value but also potential business risks. A responsible design should be transparent in its operation and should not compromise ethical values for business value. This principle helps gain customer and user trust, fostering fairness and engagement.

5.2.4. DP4: Design for human involvement and engagement

Several crucial advantages arise from human involvement in AIADM systems. First, human domain expertise is an essential input for AIADM systems in organisational contexts. Humans possessing tacit knowledge about decision contexts can comprehend intangible information that may elude AIADM systems. The integration of this tacit knowledge into the AIADM system, whenever possible, improves the system performance. Second, AI algorithms are prone to errors and might yield unintended results (Shrestha et al., 2021, Xin et al., 2018). Such errors may lead to detrimental consequences and incur many types of risks for organisations. For instance, exogenous shocks such as pandemics and climate disasters could result in drastic changes in the quality of data for making predictions. Humans can identify and rectify such errors, contributing to user acceptance and trust in the systems.

To attain these benefits, we facilitated human involvement in two ways. First, by closely involving the domain experts in every phase of design, development and evaluation/auditing, we could integrate tacit knowledge components into our AIADM system, prevent unintended outcomes, and preserve responsibility and accountability. Human decisions also act as benchmarks for AI decisions in evaluation, as we showed. We devised interactive user interfaces with customisable parameters, enabling domain experts to seamlessly integrate tacit domain knowledge into the system during operation. Moreover, by grounding our work in decision augmentation over automation, we leave the responsibility and accountability of action selection with humans and illustrate the significance of having the human in the loop (Feuerriegel et al., 2022, Grønsund & Aanestad, 2020). Second, to benefit from aggregation and interaction, AI systems should be designed to connect different users (customers and employees) who interact with them. Outputs from the AI systems (dashboards, reports, plots, user interfaces, etc.) should facilitate engagement and interpretability/explainability. In our AIADM system, we demonstrated the explainability of AI outcomes using two concepts of explainable AI: feature importance and feature attribution (SHAP, LIME). We further enhanced explainability to wider audiences by using visualisations such as PDPs, variable importance, and topic modelling visualisations using pyLDAvis (Mabey, 2018). In the absence of explainability, gaining trust and confidence in AI is particularly challenging (Burkart & Huber, 2021). User feedback should always be used in updating models. A human-centred design fosters stakeholder and user trust and aegis (Bauer et al., 2023), overcoming algorithmic aversion (Dietvorst et al., 2018). Such a design principle is particularly useful for customer-centric business models such as that of TBò.

While embracing the benefits and possibilities of integrating AI into decision making, organisations should also recognise that lack of human involvement and over-reliance on AI could lead to decisionmakers losing their domain knowledge and autonomy and deskilling of the workforce (Xue et al., 2022). One way of mitigating that is introducing decision-making designs in which humans are involved (e.g., Choudhary et al., 2023, Te'eni et al., 2023). In such designs, employees enhance their proficiency in working efficiently with the system, preserving their skills and knowledge.

5.2.5. DP5: Design for continuous learning and adaptation

Our ADR study revealed that, given the emerging characteristics of technology that AI represents, its design and deployment cannot be fixed from the outset (Bailey et al., 2022). The organisation should

embrace failures and adopt a continuous improvement mindset to overcome various challenges and uncertainties that may arise in different stages of AI development and deployment. E.g., the prototype models of AI might not be highly effective and accurate in their predictions due to lack of training data. The development of AI is a staged process, and as data is accumulated over time, AI models and corresponding use cases need to be adapted. An AI model's effectiveness increases when various users engage with it and the system improves over time. If use is restricted, opportunities to update become limited. Our AIADM system design and deployment was characterised by many iterations and adaptations. Hence, we learned that AIADM system development should be guided by adaptive and iterative enhancements to minimise the risk of failure and accumulate the learning effects over time.

Equally important is keeping the design and expectations around AIADM realistic, as AI is not a technological panacea for all business ills (Berente et al., 2021). AI implementation is filled with various trade-offs in different stages, as we demonstrated in this study. Significant trade-offs derive from high costs for recruiting talent and amassing resources and changes in organisational structures and decision-making processes. This often induces significant risks (Mikalef et al., 2022). It is thus crucial to curb managerial over-expectations and subsequent disappointments. Managers should view AI neither as a magic bullet nor a quick fix.

5.2.6. DP6: Design for open knowledge and resource utilisation

Given the massive costs of full internal development (Tarafdar et al., 2019) and the necessity of a multidisciplinary approach (Lyytinen et al., 2021, von Krogh, 2018), AI projects should follow an open and collaborative design. By "open", we mean utilisation of community-developed open-source code, AI/ML libraries, platforms, datasets, tools, etc. Developing modern AI models (e.g., large language models) requires huge initial investments and many AI experts, which is beyond reach of most organisations. The core reason for adopting open resources – datasets, source codes, and models – to build corporate AI is the benefit of attracting external knowledge to supplement the internal knowledge while alleviating exorbitant development costs (Shrestha et al., 2023, von Krogh & Haefliger, 2010). In our ADR study, technology reuse alleviated talent and resource constraints.

Furthermore, AI deployment in organisations is a complex process that requires expertise in multiple disciplines. As we observed, both data science competence and business domain competence are needed to address these challenges. The data

scientists bring extensive knowledge in areas such as natural language processing, ML algorithms, statistical inference, data analysis, and knowledge representation and reasoning. Business domain experts bring deep hands-on knowledge about the tasks, workflows, and business models, and they reckon the logic of deriving business value from AI deployment. AI technologies are evolving rapidly, and it is logical to set up industry-academia collaborations and external expert partnerships and to engage in open innovation initiatives to be at the forefront of the AI frontier (Berente et al., 2021). TBô successfully led the industry-academia collaboration by proactively engaging with the researchers and subsequently conducting an ADR study within their firm. This strategy will help uphold high standards for both operational and scientific excellence.

Table 4 stipulates design goals, as well as the mechanisms to achieve these goals for each design principle (Gregor et al., 2020).

6. Discussion

6.1. Implications for research

We provide a twofold contribution to the IS design literature. First, we investigate the organisational challenges facing AIADM implementation by clearly documenting our ADR approach and illustrating potential trade-offs and challenges that managers might face in AIADM system design and deployment. Second, based on the challenges we identified, we propose a set of six design principles to guide organisations in designing and deploying AIADM systems. Unlike the traditional view of DSS design as a passive tool with static decision rules, lacking the ability to learn, adapt, initiate decision-making processes and accept decision rights and responsibilities for achieving optimal outcomes under uncertainty, our design principles take into consideration the stochastic, adaptive, and agentic nature of AI systems (Baird & Maruping, 2021). This approach, in turn, captures

Table 4. Design principles, design goals, and mechanisms to achieve them.

Design principle	Design goals/sub-goals	Mechanisms to achieve goals
DP1: Design for alignment between the business model and organisational resources and capabilities	<ol style="list-style-type: none"> 1. Align the AIADM system with the business model 2. Align the AIADM system with the organisational resources and capabilities 	Formulate an actionable strategic roadmap illustrating: <ol style="list-style-type: none"> 1a. measurable and interpretable AIADM use cases 1b. domain expertise needed for the identified business cases 2a. technical feasibility of the (proposed) AI tech stack 2b. clear and concrete goals, sub-goals, timeline, and challenges in implementing AIADM
DP2: Design for synergy in input, model, and output to ensure business value	<ol style="list-style-type: none"> 1. Overcome data constraints 2. Make the AI outcomes accurate and understandable 3. Demonstrate value for business and stakeholders to help navigate through competing interests 	To demonstrate this principle, a firm should: <ol style="list-style-type: none"> 1. Find relevant data for business problems or combine multiple datasets using common fields 2a. Pick the right ML performance metric to optimise 2b. Try different models and prioritise according to the chosen performance metric 2c. Make visualisations, e.g., PDPs and feature importance 3. Develop the prototype system and follow a comprehensive evaluation (POC, POV, and POU)
DP3: Design for ethical AI governance frameworks	<ol style="list-style-type: none"> 1. Ensure ethical management and use of data 2. Establish trust and engagement with customers and users 	<ol style="list-style-type: none"> 1a. Adopt extant regulatory guidelines (e.g., European Commission, 2019, 2021; OECD, 2021) 1b. Establish an AI auditing and governance framework (Grønsund & Aanestad, 2020) 2. Develop own AI guiding principles consistent with customer and user expectations (Bessen, Impink, & Seamans, 2022; Google, 2022)
DP4: Design for human involvement and engagement	<ol style="list-style-type: none"> 1. Integrate domain knowledge into system design, operation, and evaluation 2. Enhance employees' proficiency in working efficiently with the system 	<ol style="list-style-type: none"> 1. Keep the domain experts in the loop to actively integrate tacit domain knowledge into system design and, when in operation, identify instances in which AI errs. e.g., implement interactive user interfaces with customisable parameters 2. Facilitate user engagement and interpretability/explainability of AI outcomes through elements of explainable AI and visualisations, e.g., explainability libraries, SHAP values, or LIME
DP5: Design for continuous learning and adaptation	<ol style="list-style-type: none"> 1. Continuous improvement over new advancements in AI technology (algorithms, hardware, data, etc.) 2. Adaptability to changes in environment and decision-making contexts 	<ol style="list-style-type: none"> 1. Monitor model decay (e.g., through ML performance metrics) 2. Embrace an iterative process to overcome various challenges and uncertainties that may arise in different stages of AI design and deployment. As data is accumulated over time, AI models and corresponding use cases need to be updated and improved iteratively
DP6: Design for open knowledge and resource utilisation	<ol style="list-style-type: none"> 1. Alleviate knowledge and resource constraints 2. Alleviate talent and domain expertise constraints 	<ol style="list-style-type: none"> 1. Utilise community-developed open-source code (Github, Stackoverflow) for collaborative problem solving; AI/ML libraries (Scikit-learn, PyTorch, Keras, Carat, Rpart, etc.) for pre-built functionalities; and platforms (TensorFlow, OpenML, MLflow), datasets (Kaggle, Datahub.io), and tools (ClearML, CNTK) 2. Foster partnerships (e.g., industry-academia, solution providers, outsourcing/offshoring)

contextual idiosyncrasies and practice-based domain knowledge, ensuring practitioner relevance and acceptance (Miah et al., 2019).

Our design principles speak to several important aspects of the AIADM phenomenon. DP1 addresses the management of AIADM by listing the essential elements of an AI adoption strategy. Managing AI is different to typical IT management due to AI's agentic nature, superior learning capabilities, and incomprehensibility compared to that of IT artefacts in the past (Baird & Maruping, 2021, Berente et al., 2021). Pursuing the unprecedented opportunities realised through AI, managers of AI initiatives also grapple with myriad new challenges. Therefore, in strategic alignment, we considered five elements to guide managers in their AI adoption strategy formulation: use cases, domain expertise, goals, feasibility, and likely challenges.

DP2 recommends synergy in the input elements, AI modelling, and output elements as the pathway to ascertaining business value. This delves into the polarised academic debate on whether AI fosters intended performance improvements. Scholars express scepticism that much of the traction on AIADM is merely hype and that it may fail to produce measurable performance gains (Aaen et al., 2022, Ermakova et al., 2021, Rana et al., 2022). Attempts to integrate AI into the decision-making process and value chains often fail (Joshi et al., 2021, Ransbotham et al., 2019). Therefore, we emphasise that organisations should make the system components congruent and rigorously scrutinise their AI systems for value creation and usability.

DP3 underscores the notion of responsible AI – principles that involve ethical, fair, secure, and accountable design and deployment of AI (Golovianko et al., 2022, Mikalef et al., 2022). Regardless of the elementary AI guidelines put forth by regulators and the continually evolving AI technology stack, we emphasised three requisites: the adoption of extant regulations, developing AI policies in-house to cater to customer and user needs, and establishing an AI auditing and governance framework.

DP4 endorses two closely related concepts of human-AI ensembles: human-in-the-loop frameworks (Grønsund & Aanestad, 2020, Xin et al., 2018) and explainable AI (Bauer et al., 2023, Senoner et al., 2022). Our study demonstrates how human-in-the-loop frameworks unfold in practice to successfully integrate tacit domain knowledge into AI system design and audit AI outcomes to prevent error propagation. By opting for decision augmentation over automation, we held the human accountable and responsible for action selection while the protocol development was vested in AI (Murray et al., 2020). By doing so, we successfully combined the benefits of an efficient AI system with the unique abilities of

humans in decision processes. For the humans in the loop to function effectively and efficiently, the results of AI systems should be sufficiently interpretable and explainable. Our study showcased how AI outcomes can be explained in practice. Thereby, we mitigate problems associated with the black-box nature of many contemporary AI systems while garnering wider user and stakeholder acceptance (Bauer et al., 2023).

DP5 extends the understanding of data-driven value propositions (Günther et al., 2022, Wiener et al., 2020). Recent work in this domain attests that “the process of creating data-driven value propositions is emergent, consisting of iterative resourcing cycles” (Günther et al., 2022, p. 1). Realising value from data relies on reconstruction and repurposing of both data and algorithms, as it is an interconnected process of trial and error (e.g., Chapman et al., 2000). Moreover, unlike conventional IS, AI systems improve over time and use as they learn from the accumulated deeper pools of data. Therefore, the extant literature concurs about the need for a flexible and iterative design process for data projects including AI from a purely technical standpoint (Abbasi et al., 2016). Our study broadens the scope of examination by including organisational processes surrounding the technical AI development process. Our study uncovers the organisational challenges encountered in the process and observes how these challenges create iterative back-and-forth workflows.

DP6 endorses open innovation (Chesbrough, 2003) in AI development (Shrestha et al., 2023). During AI development, organisations can leverage free and open resources – data, code, models, and developer communities – to overcome resource constraints. Subsequently, firms can decide to open their innovations to encourage community-driven improvements to the system at minimal marginal costs.

6.2. Implications for practice

Data generation, access, and collection is a hallmark of contemporary organisations. With the rapid scaling of data, AI becomes indispensable for value creation and value capture in firms (Iansiti & Lakhani, 2020). Within organisations, the rapid scaling of data can become an obstacle unless new systems can be designed and deployed to aid managers in making timely and effective decisions (Agrawal et al., 2018). Yet, the value from data can only be captured effectively when the quality of the data is matched with well-designed and deployed AI systems (Bessen et al., 2022). Our study demonstrates that the design and deployment of an AIADM system needs to consider technical, organisational, human, and social factors equally. Our design principles are intended to offer guidance to managers in developing and adopting

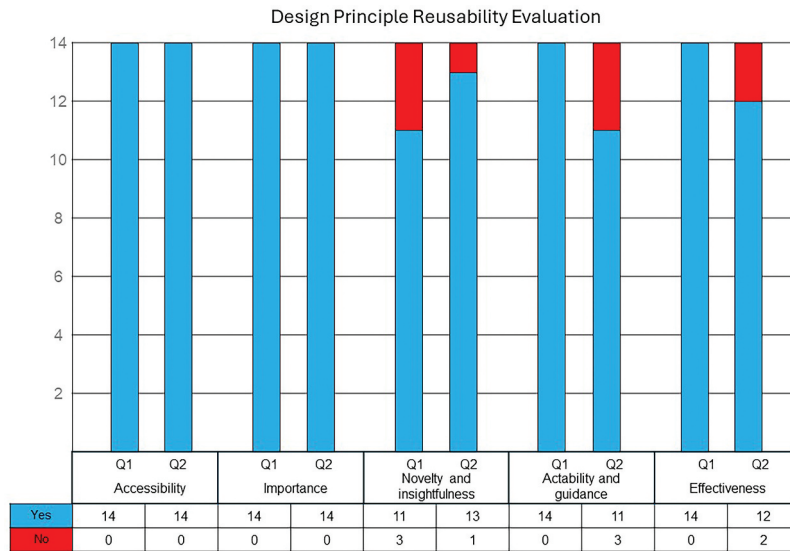


Figure 5. Design principle reusability evaluation.

powerful AIADM systems in their organisations while remaining aware of this wide range of factors.

A key insight from our research is that although AI algorithms are designed to automate or augment managerial decision making, the process of designing and deploying AI is itself filled with trade-offs and challenges that require critical managerial judgements. Our study shows some of the challenges and trade-offs managers may encounter when pushing AI within their organisations. We found that the organisation requires best practices (e.g., a strategic roadmap for AI adoption), which we outline as our six design principles to mitigate technical, social, and organisational barriers.

We adopted the Design Principle Reusability Evaluation Framework by Iivari et al. (2021) to assess the transferability of our design principles to other contexts (external validity). This encompasses five key criteria: (1) accessibility, (2) importance, (3) novelty and insightfulness, (4) actability and guidance, and (5) effectiveness. The evaluation involved engagement with 14 managers and developers of AIADM systems – the target audience of the design principles – with extensive experience in IT and digital transformation. Substantial evidence (see Figure 5 and Appendix F) indicates that our design principles are helpful in practical applications.

6.3. Limitations and future research

Our paper also has some limitations which provide opportunities for future research. First, we derive our design principles from a single case sample, characterised by a relatively small company that predominantly works within an e-commerce set-up. As a result, some challenges of AI deployment that we identified may be specific to this company. Organisations in other data-heavy industries like

finance, pharmaceuticals, healthcare, fast-moving consumer goods, and manufacturing and business functions like hiring, marketing, distribution, and quality assurance are likely to bring interesting AI use cases that are different from the case we studied. Therefore, for the sake of the generalisability of the insights derived, future research needs to be conducted in a larger set of organisations with a cross-comparison of identified mechanisms, challenges, and design principles. Second, the timeline of this research work coincided with the COVID-19 outbreak in Switzerland; as a result, most of the interviews and meetings took place online via Zoom or offline with protective measures in place. This set-up might have reduced capture of some social cues that are usually available in face-to-face interviews and observations. Based on our work, we see promising opportunities for design science research to generate prescriptive knowledge pertinent to AI use in practice.

7. Conclusion

This study advances the argument that AIADM systems represent a novel class of information systems that exhibit unique socio-organisational dynamics and technological complexity when juxtaposed with conventional DSS. AIADM systems present unprecedented opportunities and challenges for contemporary organisations. Accordingly, striking a balance between the potential gains and risks associated with AIADM systems requires identifying specific design principles to guide practice. To this end, we offer actionable design principles following a comprehensive design science-based investigation of the design and deployment of AIADM systems.

Acknowledgements

We thank Roy Bernheim, Allan Perrottet, Matthew Soroka, and Nina Geilinger for their invaluable support in project coordination. We acknowledge the exceptional assistance of Noah Hampf, Esther Le Mair, Parinitha Mundra, and Dmitry Plekhanov in preparing the manuscript.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the Swiss National Science Foundation under Grant [197763].

References

- Aaen, J., Nielsen, J. A., & Carugati, A. (2022). The dark side of data ecosystems: A longitudinal study of the DAMD project. *European Journal of Information Systems*, 31(3), 288–312. <https://doi.org/10.1080/0960085X.2021.1947753>
- Abbasi, A., Sarker, S., & Chiang, R. H. (2016). Big data research in information systems: Toward an inclusive research agenda. *Journal of the Association for Information Systems*, 17(2), 3. <https://doi.org/10.17705/1jais.00423>
- Abdul, A., Vermeulen, J., Wang, D., Lim, B. Y., & Kankanhalli, M. (2018). Trends and trajectories for explainable, accountable and intelligible systems: An HCI research agenda. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, Montreal, Canada (pp. 1–18).
- Abouzahra, M., Guenter, D., & Tan, J. (2022). Exploring physicians' continuous use of clinical decision support systems. *European Journal of Information Systems*, 33(2), 1–22. <https://doi.org/10.1080/0960085X.2022.2119172>
- Ågerfalk, P. J. (2020). Artificial intelligence as digital agency. *European Journal of Information Systems*, 29(1), 1–8. <https://doi.org/10.1080/0960085X.2020.1721947>
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Press.
- Altendeitering, M., Fraunhofer, I. S. S. T., & Guggenberger, T. (2021). Designing data quality tools: Findings from an action design research project at Boehringer Ingelheim. In *Proceedings of the 29th European Conference on Information Systems*, Marrakech, Morocco (pp. 17).
- Arnott, D. (2006). Cognitive biases and decision support systems development: A design science approach. *Information Systems Journal*, 16(1), 55–78. <https://doi.org/10.1111/j.1365-2575.2006.00208.x>
- Arnott, D., & Pervan, G. (2005). A critical analysis of decision support systems research. *Journal of Information Technology*, 20(2), 67–87. <https://doi.org/10.1057/palgrave.jit.2000035>
- Arnott, D., & Pervan, G. (2008). Eight key issues for the decision support systems discipline. *Decision Support Systems*, 44(3), 657–672. <https://doi.org/10.1016/j.dss.2007.09.003>
- Arnott, D., & Pervan, G. (2014). A critical analysis of decision support systems research revisited: The rise of design science. *Journal of Information Technology*, 29(4), 269–293. <https://doi.org/10.1057/jit.2014.16>
- Bailey, D., Faraj, S., Hinds, P., von Krogh, G., & Leonardi, P. (2022). Special issue of *organization Science*: Emerging technologies and organizing. *Organization Science*, 30(3), 642–646. <https://doi.org/10.1287/orsc.2019.1299>
- Baird, A., & Maruping, L. M. (2021). The next generation of research on is use: A theoretical framework of delegation to and from agentic is artifacts. *MIS Quarterly*, 45(1), 315–341. <https://doi.org/10.25300/MISQ/2021/15882>
- Barkhi, R. (2002). The effects of decision guidance and problem modeling on group decision-making. *Journal of Management Information Systems*, 18(3), 259–282.
- Bauer, K., von Zahn, M., & Hinz, O. (2023). Expl (AI) ned: The impact of explainable artificial intelligence on users' information processing. *Information Systems Research*, 34(4), 1582–1602. <https://doi.org/10.1287/isre.2023.1199>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450.
- Bessen, J., Impink, S. M., Reichensperger, L., & Seamans, R. (2022). The role of data for AI startup growth. *Research Policy*, 51(5), 104513. <https://doi.org/10.1016/j.respol.2022.104513>
- Bessen, J., Impink, S. M., & Seamans, R. (2022). The cost of ethical AI development for AI startups. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, Oxford, United Kingdom (pp. 92–106).
- Bouschery, S. G., Blazevic, V., & Piller, F. T. (2023). Augmenting human innovation teams with artificial intelligence: Exploring transformer-based language models. *Journal of Product Innovation Management*, 40(2), 139–153. <https://doi.org/10.1111/jpim.12656>
- Brynjolfsson, E., & McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, 106(5), 133–39. <https://doi.org/10.1257/aer.p20161016>
- Burkart, N., & Huber, M. F. (2021). A survey on the explainability of supervised machine learning. *Journal of Artificial Intelligence Research*, 70, 245–317. <https://doi.org/10.1613/jair.1.12228>
- Chandra Kruse, L., Seidel, S., & Purao, S. (2016). Making use of design principles. In *Tackling Society's Grand Challenges with Design Science: 11th International Conference, DESRIST 2016, St John's, NL, Canada, May 23–25, 2016, 11* (pp. 37–51). Springer International Publishing.
- Chandra, L., Seidel, S., & Gregor, S. (2015). Prescriptive knowledge in IS research: Conceptualizing design principles in terms of materiality, action, and boundary conditions. In *2015 48th Hawaii International Conference on System Sciences*, Hawaii, USA (pp. 4039–4048). IEEE.
- Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., & Wirth, R. (2000). CRISP-DM 1.0: Step-by-step data mining guide. *SPSS Inc*, 9(13), 1–77. <https://mineraoedados.files.wordpress.com/2012/12/crisp-dm-1-0.pdf>
- Chen, C. W., & Koufaris, M. (2015). The impact of decision support system features on user overconfidence and risky behavior. *European Journal of Information Systems*, 24(6), 607–623. <https://doi.org/10.1057/ejis.2014.30>
- Chesbrough, H. W. (2003). *Open innovation: The new imperative for creating and profiting from technology*. Harvard Business Press.

- Choudhary, V., Marchetti, A., Shrestha, Y. R., & Puranam, P. (2023). Human-AI ensembles: When can they work? *Journal of Management*. <https://doi.org/10.1177/01492063231194968>
- Collins, J., Ketter, W., & Gini, M. (2010). Flexible decision support in dynamic inter-organisational networks. *European Journal of Information Systems*, 19(4), 436–448. <https://doi.org/10.1057/ejis.2010.24>
- Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm*. Englewood Cliffs, New Jersey: Prentice-Hall.
- Dasborough, M. T. (2023). Awe-inspiring advancements in AI: The impact of ChatGPT on the field of organizational behavior. *Journal of Organizational Behavior*, 44(2), 177–179. <https://doi.org/10.1002/job.2695>
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64(3), 1155–1170. <https://doi.org/10.1287/mnsc.2016.2643>
- Ebel, P., Bretschneider, U., & Leimeister, J. M. (2016). Leveraging virtual business model innovation: A framework for designing business model development tools. *Information Systems Journal*, 26(5), 519–550. <https://doi.org/10.1111/isj.12103>
- Ermakova, T., Blume, J., Fabian, B., Fomenko, E., Berlin, M., & Hauswirth, M. (2021). Beyond the hype: Why do data-driven projects fail? In *Proceedings of the 54th Hawaii International Conference on System Sciences*, Hawaii, USA (pp. 5081).
- European Commission. (2019). Ethics guidelines for trustworthy AI. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
- European Commission. (2021). Proposal for a regulation of the European Parliament and of the council: Laying down harmonised rules on artificial intelligence (artificial intelligence Act) and amending certain union legislative acts. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52021PC0206&from=EN>
- Fang, X., Gao, Y., & Hu, P. J. (2021). A prescriptive analytics method for cost reduction in clinical decision making. *MIS Quarterly*, 45(1), 83–115. <https://doi.org/10.25300/MISQ/2021/14372>
- Feuerriegel, S., Shrestha, Y. R., von Krogh, G., & Zhang, C. (2022). Bringing artificial intelligence to business management. *Nature Machine Intelligence*, 4(7), 611–613. <https://doi.org/10.1038/s42256-022-00512-5>
- Fügener, A., Grahl, J., Gupta, A., & Ketter, W. (2022). Cognitive challenges in human-artificial intelligence collaboration: Investigating the path toward productive delegation. *Information Systems Research*, 33(2), 678–696. <https://doi.org/10.1287/isre.2021.1079>
- Galbraith, J. R. (1974). Organization design: An information processing view. *Interfaces*, 4(3), 28–36. <https://doi.org/10.1287/inte.4.3.28>
- Giermindl, L. M., Strich, F., Christ, O., Leicht-Deobald, U., & Redzepi, A. (2022). The dark sides of people analytics: Reviewing the perils for organisations and employees. *European Journal of Information Systems*, 31(3), 410–435. <https://doi.org/10.1080/0960085X.2021.1927213>
- Golovianko, M., Gryshko, S., Terziyan, V., & Tuunanen, T. (2022). Responsible cognitive digital clones as decision-makers: A design science research study. *European Journal of Information Systems*, 32(5), 1–23. <https://doi.org/10.1080/0960085X.2022.2073278>
- Google. (2022) *Responsible AI Practices*. <https://ai.google/responsibilities/responsible-ai-practices/>
- Gorry, G. A., & Morton, M. S. (1971). A framework for management information systems. *Sloan Management Review*, 13(1), 1–22.
- Gregor, S., & Hevner, A. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337–355. <https://doi.org/10.25300/MISQ/2013/37.2.01>
- Gregor, S., Kruse, L. C., & Seidel, S. (2020). Research perspectives: The anatomy of a design principle. *Journal of the Association for Information Systems*, 21(6), 1622–1652. <https://doi.org/10.17705/1jais.00649>
- Grønsund, T., & Aanestad, M. (2020). Augmenting the algorithm: Emerging human-in-the-loop work configurations. *The Journal of Strategic Information Systems*, 29(2), 101614. <https://doi.org/10.1016/j.jsis.2020.101614>
- Grove, W. M., & Meehl, P. E. (1996). Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: The clinical-statistical controversy. *Psychology, Public Policy, and Law*, 2(2), 293. <https://doi.org/10.1037/1076-8971.2.2.293>
- Grove, W. M., Zald, D. H., Lebow, B. S., Snitz, B. E., & Nelson, C. (2000). Clinical versus mechanical prediction: A meta-analysis. *Psychological Assessment*, 12(1), 19. <https://doi.org/10.1037/1040-3590.12.1.19>
- Günther, W. A., Mehrizi, M. H. R., Huysman, M., Deken, F., & Feldberg, F. (2022). Resourcing with data: Unpacking the process of creating data-driven value propositions. *The Journal of Strategic Information Systems*, 31(4), 101744. <https://doi.org/10.1016/j.jsis.2022.101744>
- Hevner, A., March, S., Park, J., & Ram, S. (2004). Design science in Information Systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
- Hevner, A., & Storey, V. (2023). Research challenges for the design of human-artificial intelligence systems (HAIS). *ACM Transactions on Management Information Systems*, 14(1), 1–18. <https://doi.org/10.1145/3549547>
- Huber, G. P. (1990). A theory of the effects of advanced information technologies on organizational design, intelligence, and decision making. *The Academy of Management Review*, 15(1), 47–71. <https://doi.org/10.2307/258105>
- Iansiti, M., & Lakhani, K. R. (2020). *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world*. Harvard Business Press.
- Iivari, J. (2015). Distinguishing and contrasting two strategies for design science research. *European Journal of Information Systems*, 24(1), 107–115. <https://doi.org/10.1057/ejis.2013.35>
- Iivari, J., Rotvit Perlt Hansen, M., & Haj-Bolouri, A. (2021). A proposal for minimum reusability evaluation of design principles. *European Journal of Information Systems*, 30(3), 286–303. <https://doi.org/10.1080/0960085X.2020.1793697>
- Jain, H., Padmanabhan, B., Pavlou, P. A., & Raghu, T. S. (2021). Editorial for the special section on humans, algorithms, and augmented intelligence: The future of work, organizations, and society. *Information Systems Research*, 32(3), 675–687. <https://doi.org/10.1287/isre.2021.1046>
- Joseph, J., & Gaba, V. (2020). Organizational structure, information processing, and decision-making: A retrospective and road map for research. *Academy of Management Annals*, 14(1), 267–302. <https://doi.org/10.5465/annals.2017.0103>
- Joshi, M. P., Su, N., Austin, R. D., & Sundaram, A. K. (2021). Why so many data science projects fail to deliver. *MIT*

- Sloan Management Review*, 62(3), 85–89.
- Kamis, A., Koufaris, M., & Stern, T. (2008). Using an attribute-based decision support system for user-customized products online: An experimental investigation. *MIS Quarterly*, 32(1), 159–177. <https://doi.org/10.2307/25148832>
- Keding, C., & Meissner, P. (2021). Managerial overreliance on AI-augmented decision-making processes: How the use of AI-based advisory systems shapes choice behavior in R&D investment decisions. *Technological Forecasting and Social Change*, 171, 120970. <https://doi.org/10.1016/j.techfore.2021.120970>
- Keen, P. G. (1980). Decision support systems: A research perspective. Decision support systems: Issues and challenges, International Institute for Applied systems Analysis (IIASA). *Proceedings Series*, 11, 23–27. https://books.google.ch/books?hl=en&lr=&id=LF0hBQAAQBAJ&oi=fnd&pg=PA23&ots=n3GZiEb nSB&sig=9-r9bb1S-oMNS2UXUeAO9qCRqMA&redir_esc=y#v=onepage&q&f=false
- Klör, B., Monhof, M., Beverungen, D., Bräuer, S., Niehaves, B., Tuunanen, T., & Peffers, K. (2018). Design and evaluation of a model-driven decision support system for repurposing electric vehicle batteries. *European Journal of Information Systems*, 27(2), 171–188. <https://doi.org/10.1057/s41303-017-0044-3>
- Kordzadeh, N., & Ghasemaghaei, M. (2022). Algorithmic bias: Review, synthesis, and future research directions. *European Journal of Information Systems*, 31(3), 388–409. <https://doi.org/10.1080/0960085X.2021.1927212>
- Kuguoglu, B. K., van der Voort, H., & Janssen, M. (2021). The giant leap for smart cities: Scaling up smart city artificial intelligence of things (AIoT) initiatives. *Sustainability*, 13(21), 12295. <https://doi.org/10.3390/su132112295>
- Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science*, 33(1), 126–148. <https://doi.org/10.1287/orsc.2021.1549>
- Leidner, D. E., & Elam, J. J. (1995). The impact of executive information systems on organizational design, intelligence, and decision making. *Organization Science*, 6(6), 645–664. <https://doi.org/10.1287/orsc.6.6.645>
- Lilien, G. L., Rangaswamy, A., Van Bruggen, G. H., & Starke, K. (2004). DSS effectiveness in marketing resource allocation decisions: Reality vs. perception. *Information Systems Research*, 15(3), 216–235. <https://doi.org/10.1287/isre.1040.0026>
- Lynch, T., & Gregor, S. (2004). User participation in decision support systems development: Influencing system outcomes. *European Journal of Information Systems*, 13(4), 286–301. <https://doi.org/10.1057/palgrave.ejis.3000512>
- Lyytinen, K., Nickerson, J. V., & King, J. L. (2021). Metahuman systems= humans+ machines that learn. *Journal of Information Technology*, 36(4), 427–445. <https://doi.org/10.1177/0268396220915917>
- Mabey, B. (2018). pyLdavis: Python library for interactive topic model visualization. *port of the R LDavis package*. <https://github.com/bmabey/pyldavis>
- Maedche, A., Gregor, S., & Parsons, J. (2021). Mapping design contributions in information systems research: The design research activity framework. *Communications of the Association for Information Systems*, 49(1), 12. <https://doi.org/10.17705/1CAIS.04914>
- Mandviwalla, M. (2015). Generating and justifying design theory. *Journal of the Association for Information Systems*, 16(5), 3. <https://doi.org/10.17705/1jais.00397>
- March, J. G. (1994). *Primer on decision making: How decisions happen*. Simon and Schuster.
- March, J. G., & Olsen, J. P. (1989). *Rediscovering institutions*. Free Press.
- March, J. G., & Simon, H. A. (1993). *Organizations*. John Wiley & Sons.
- Miah, S. J., Gammack, J. G., & McKay, J. (2019). A metadesign theory for tailorable decision support. *Journal of the Association for Information Systems*, 20(5), 4. <https://doi.org/10.17705/1jais.00544>
- Mikalef, P., Conboy, K., Lundström, J. E., & Popovič, A. (2022). Thinking responsibly about responsible AI and ‘the dark side’ of AI. *European Journal of Information Systems*, 31(3), 257–268. <https://doi.org/10.1080/0960085X.2022.2026621>
- Molloy, S., & Schwenk, C. R. (1995). The effects of information technology on strategic decision making. *Journal of Management Studies*, 32(3), 283–311. <https://doi.org/10.1111/j.1467-6486.1995.tb00777.x>
- Murray, A., Rhymmer, J., & Sirmon, D. G. (2020). Humans and technology: Forms of conjoined agency in organizations. *Academy of Management Review*, 46(3), 552–571. <https://doi.org/10.5465/amr.2019.0186>
- Nguyen, A., Tuunanen, T., Gardner, L., & Sheridan, D. (2021). Design principles for learning analytics information systems in higher education. *European Journal of Information Systems*, 30(5), 541–568. <https://doi.org/10.1080/0960085X.2020.1816144>
- Nunamaker, J. F., Jr., Briggs, R. O., Derrick, D. C., & Schwabe, G. (2015). The last research mile: Achieving both rigor and relevance in information systems research. *Journal of Management Information Systems*, 32(3), 10–47. <https://doi.org/10.1080/07421222.2015.1094961>
- OECD. (2021). *Artificial Intelligence*. <https://www.oecd.org/going-digital/ai/principles/>
- Orlikowski, W. J., & Iacono, C. S. (2001). Research commentary: Desperately seeking the “IT” in it research. A call to theorizing the it artifact. *Information Systems Research*, 12(2), 121–134. <https://doi.org/10.1287/isre.12.2.121.9700>
- Padmanabhan, B., Sahoo, N., & Burton-Jones, A. (2022). Machine learning in Information Systems research. *Management Information Systems Quarterly*, 46(1), iii–xix.
- Pan, S. L., Li, M., Pee, L. G., & Sandeep, M. S. (2021). Sustainability design principles for a wildlife management analytics system: An action design research. *European Journal of Information Systems*, 30(4), 452–473. <https://doi.org/10.1080/0960085X.2020.1811786>
- Peffers, K., Tuunanen, T., & Niehaves, B. (2018). Design science research genres: Introduction to the special issue on exemplars and criteria for applicable design science research. *European Journal of Information Systems*, 27(2), 129–139. <https://doi.org/10.1080/0960085X.2018.1458066>
- Power, D. J. (2001). Supporting decision-makers: An expanded framework. *Informing Science*, 1(1), 1901–1915. <https://doi.org/10.28945/2384>

- Rai, A. (2016). Editor's comments: Synergies between big data and theory. *MIS Quarterly*, 40(2), iii–ix.
- Rai, A., Constantinides, P., & Sarker, S. (2019). Next generation digital platforms: Toward human-AI hybrids. *MIS Quarterly*, 43(1), iii–ix.
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>
- Rana, N. P., Chatterjee, S., Dwivedi, Y. K., & Akter, S. (2022). Understanding dark side of artificial intelligence (AI) integrated business analytics: Assessing firm's operational inefficiency and competitiveness. *European Journal of Information Systems*, 31(3), 364–387. <https://doi.org/10.1080/0960085X.2021.1955628>
- Ransbotham, S., Khodabandeh, S., Fehling, R., LaFountain, B., & Kiron, D. (2019). *Winning with AI: Pioneers combine strategy, organizational behavior, and technology*. MIT Sloan Management Review and Boston Consulting Group.
- Remus, W. E., & Kottemann, J. E. (1986). Toward intelligent decision support systems: An artificially intelligent statistician. *MIS Quarterly*, 10(4), 403–418. <https://www.jstor.org/stable/249197>
- Rinta-Kahila, T., Someh, I., Gillespie, N., Indulska, M., & Gregor, S. (2022). Algorithmic decision-making and system destructiveness: A case of automatic debt recovery. *European Journal of Information Systems*, 31(3), 313–338. <https://doi.org/10.1080/0960085X.2021.1960905>
- Samtani, S., Chai, Y., & Chen, H. (2021). Linking exploits from the dark web to known vulnerabilities for proactive cyber threat intelligence: An attention-based deep structured semantic model. *MIS Quarterly*, 46(2), 911–946. <https://doi.org/10.25300/MISQ/2022/15392>
- Schoemaker, P. J. (1982). The expected utility model: Its variants, purposes, evidence and limitations. *Journal of Economic Literature*, 20(2), 529–563. <https://www.jstor.org/stable/2724488>
- Seidel, S., Chandra Kruse, L., Székely, N., Gau, M., Stieger, D., Peffers, K., Tuunanen, T., Niehaves, B., & Lyytinen, K. (2018). Design principles for sensemaking support systems in environmental sustainability transformations. *European Journal of Information Systems*, 27(2), 221–247. <https://doi.org/10.1057/s41303-017-0039-0>
- Sein, M. K., Henfridsson, O., Puro, S., Rossi, M., & Lindgren, R. (2011). Action design research. *MIS Quarterly*, 35(1), 37–56. <https://doi.org/10.2307/23043488>
- Senoner, J., Netland, T., & Feuerriegel, S. (2022). Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science*, 68(8), 5704–5723. <https://doi.org/10.1287/mnsc.2021.4190>
- Sharma, R., Mithas, S., & Kankanhalli, A. (2014). Transforming decision-making processes: A research agenda for understanding the impact of business analytics on organisations. *European Journal of Information Systems*, 23(4), 433–441. <https://doi.org/10.1057/ejis.2014.17>
- Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), 66–83. <https://doi.org/10.1177/0008125619862257>
- Shrestha, Y. R., Krishna, V., & von Krogh, G. (2021). Augmenting organizational decision-making with deep learning algorithms: Principles, promises, and challenges. *Journal of Business Research*, 123, 588–603. <https://doi.org/10.1016/j.jbusres.2020.09.068>
- Shrestha, Y. R., von Krogh, G., & Feuerriegel, S. (2023). Building open-source AI. *Nature Computational Science*, 1–4. <https://doi.org/10.2139/ssrn.4614280>
- Simon, H. A. (1947). *Administrative behavior: A study of decision-making processes in administrative organization*. Palgrave Macmillan.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138. <https://doi.org/10.1037/h0042769>
- Simon, H. A. (1960). *The new science of management decision*. Harper.
- Singh, J. V., Tucker, D. J., & House, R. J. (1986). Organizational legitimacy and the liability of newness. *Administrative Science Quarterly*, 31(2), 171–193. <https://doi.org/10.2307/2392787>
- Sun, D., Ying, W., Zhang, X., & Feng, L. (2019). Developing a blockchain-based loyalty programs system to hybridize business and charity: An action design research. *International Conference on Information Systems 2019 Proceedings*, Munich, Germany (pp. 6).
- Tarafdar, M., Beath, C. M., & Ross, J. W. (2019). Using AI to enhance business operations. *MIT Sloan Management Review*, 60(4), 37–44.
- Te'eni, D., Yahav, I., Zagalsky, A., Schwartz, D., Silverman, G., Cohen, D., Mann, Y., & Dafna, L. (2023). *Reciprocal human-machine learning: A theory and an instantiation for the case of message classification*. <https://doi.org/10.1287/mnsc.2022.03518>
- Tinguely, P., Shrestha, Y. R., & von Krogh, G. (2020). How does your labor force react to COVID-19? Employing social media analytics for preemptive decision making. *California Management Review*. <https://cmr.berkeley.edu/2020/08/social-media-analytics/>
- Todd, P., & Benbasat, I. (1999). Evaluating the impact of DSS, cognitive effort, and incentives on strategy selection. *Information Systems Research*, 10(4), 356–374. <https://doi.org/10.1287/isre.10.4.356>
- Tushman, M. L., & Nadler, D. A. (1978). Information processing as an integrating concept in organizational design. *The Academy of Management Review*, 3(3), 613–624. <https://doi.org/10.2307/257550>
- Tuunanen, T., & Peffers, K. (2018). Population targeted requirements acquisition. *European Journal of Information Systems*, 27(6), 686–711. <https://doi.org/10.1080/0960085X.2018.1476015>
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Van den Broek, E., Sergeeva, A., & Huysman, M. (2021). When the machine meets the expert: An ethnography of developing AI for hiring. *MIS Quarterly*, 45(3), 1557–1580. <https://doi.org/10.25300/MISQ/2021/16559>
- Venable, J., Pries-Heje, J., & Baskerville, R. (2012). A comprehensive framework for evaluation in design science research. In *International conference on design science research in information systems*, Las Vegas, USA (pp. 423–438). Springer.
- von Krogh, G. (2018). Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Academy of Management Discoveries*, 4(4), 404–409. <https://doi.org/10.5465/amd.2018.0084>
- von Krogh, G., Ben-Menahem, S. M., & Shrestha, Y. R. (2021). Artificial Intelligence in Strategizing: Prospects and Challenges. *Strategic Management: State of the Field and Its Future* (pp. 625–646). New York: Oxford

- University Press. <https://academic.oup.com/book/39240/chapter/338769107>
- von Krogh, G., & Haefliger, S. (2010). Opening up design science: The challenge of designing for reuse and joint development. *The Journal of Strategic Information Systems*, 19(4), 232–241. <https://doi.org/10.1016/j.jsis.2010.09.008>
- Wiener, M., Saunders, C., & Marabelli, M. (2020). Big-data business models: A critical literature review and multi-perspective research framework. *Journal of Information Technology*, 35(1), 66–91. <https://doi.org/10.1177/0268396219896811>
- Xin, D., Ma, L., Liu, J., Macke, S., Song, S., & Parameswaran, A. (2018). Accelerating human-in-the-loop machine learning: Challenges and opportunities. In *Proceedings of the second workshop on data management for end-to-end machine learning*, Houston, USA (pp. 1–4).
- Xue, M., Cao, X., Feng, X., Gu, B., & Zhang, Y. (2022). Is college education less necessary with AI? Evidence from firm-level labor structure changes. *Journal of Management Information Systems*, 39(3), 865–905. <https://doi.org/10.1080/07421222.2022.2096542>