

Enhancing Collaborative Rationality between Humans and Machines through Data-Driven Decision Evaluation

Nada Elgandy¹

¹ M3S, Faculty of Information Technology and Electrical Engineering, University of Oulu, Oulu, Finland

Abstract

Organizational decisions have become more data-driven and collaborative, with the increasing utilization of artificial intelligence, machine learning, and analytics to support decision making. While humans and machines are each bounded in their own rationalities, their collaboration has enabled a new, collaborative rationality by augmenting the intelligence and capabilities of each. New research is required to explore the degree and mode of human-machine collaboration, with the aim of enhancing collaborative rationality, and its effect on decision making. Furthermore, the resulting decisions must be evaluated to enable learning, rationalization, and sensemaking from the decision outcomes. However, data-driven decisions are complex in nature, and current theoretical developments fall short in accommodating for their multi-faceted nature and changing context, and there is lack of theoretical support on how, when, and why to evaluate such decisions. Accordingly, we follow a design science research methodology to develop and evaluate a model for data-driven decision evaluation. This model depicts the relationship and role of the multiple elements involved in data-driven decision making, and provides a feedback loop which inputs the results of evaluating decision outcomes back into the process/system, thus enabling learning from the past through experience, and ultimately enhancing collaborative rationality and decision making.

Keywords¹

Data-driven decision making, decision evaluation, collaborative rationality, human-machine collaboration, design science research

1. Introduction

For centuries, decisions have inevitably defined the future of organizations and societies, with decision makers constantly striving to understand and deliberate the complicated process of how to decide. In recent decades, the coexistence of artificial intelligence (AI), or machine learning (ML), systems with human decision makers has ignited an interest in machines augmenting human intelligence and capabilities, which has led to different, and unprecedented, dimensions of “intelligent” data analysis in order to support and enhance decision-making and learning [1–4]. This complex interaction between humans and machines in decision making leads to the creation of metahuman systems, or sociotechnical systems where machines that learn join human learning, consequently complementing and amplifying human capabilities [5]. Accordingly, decision making involves combining the intuition and experience of human decision makers with the analytics and processing capabilities of machines with access to vast and various amounts data, thus breaking beyond the boundaries of each’s limited rationality, and providing more rational choices which can lead to better results [6–8]. Consequently, enhanced collaborative rationality ensues, where humans and machines participate in bringing their various capabilities together to jointly solve problems and make decisions [9].

^BIR 2022 Workshops and Doctoral Consortium, 21st International Conference on Perspectives in Business Informatics Research (BIR 2022), September 20-23, 2022, Rostock, Germany

EMAIL: nada.sanad@oulu.fi

ORCID: 0000-0001-6765-017X



© 2022 Copyright for this paper by its authors.
Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)

There are three main modes of collaboration between humans and machines in decision-making. The first category is full human to AI delegation, where the machine makes the algorithmic decisions, and the role of the human is simply supervisory. The second category is of a hybrid, sequential form. Either the machine provides suitable alternatives to the human, who then selects the most appropriate, or the human decision maker selects a set of alternatives and then passes them to the machine for evaluation. Here, the role of the machine is mainly to provide recommendations or insights. Finally, the third category is aggregated human-machine decision making where different aspects, or elements, of the decision are allocated to each of humans and machines based on their respective strengths, and then aggregated into a collective decision [10,11]. The focus of this research is mainly on the latter two modes of human-machine collaboration, rather than on purely machine decisions.

We utilize the term data-driven decision making to define this collaboration and encompass the relationship between the human decision maker, machine, data, decision-making process, and decision outcome [9]. Data-driven decision making involves the analysis and interpretation of data, typically through human-machine collaboration and the utilization of analytics, algorithms, or ML methods and techniques, to support a collaborative rationality and quality decisions [1,9,12].

However, to argue that such decisions are indeed better, evaluation is necessary. Evaluation clarifies options, reduces uncertainties, and generates information and knowledge about the results within contextual boundaries to make more informed decisions in the future [13]. Moreover, evaluation of past decisions can help establish whether assumptions and analytical methods are reliable or require adjustments and corrective action [14].

In recent years, research has started perceiving the relationship between organizational learning and machine learning [3,10,15,16], showing that evaluation can potentially enhance machine learning as well, since training data can be updated with the results of the evaluation. Mutual learning between humans and machines over time and in the appropriate contexts is necessary for developing a systematic, continuous process of improvement in decision making and collaboration between both. However, there is little agreement in the literature on the role of evaluation, or what and how to evaluate [17,18].

Due to the complexity of data-driven decisions and the interrelationship between the various elements involved, the evaluation and monitoring of the resulting decisions requires additional research of its own for building experiences [5,9,19]. Despite the vast amount of literature in various streams and disciplines (e.g. decision research, information systems (IS), behavioral sciences, AI, ML, information technology (IT), etc.), comprehensive, or holistic, solutions accommodating for the multifaceted nature of collaborative data-driven decisions are not found. The interaction between humans and machines and their roles in decision making is still not clear, and further research is necessary to evaluate the resulting decisions and determine the benefit, impact, and learning that consequently occur from this collaboration. Hence, we need new ways to measure and evaluate the impact of AI-enabled decisions from different perspectives in order to measure the benefits of human-machine collaboration and its role among various factors in data-driven decision making [5,9,19,20].

Accordingly, the main research question we aim to study is:

RQ 1: *“How can we support data-driven decision making and enhance collaborative rationality between humans and machines through decision evaluation?”*

We intend not just to study collaborative rationality and data-driven decision evaluation, but also to design a solution which is theoretically sound and practically feasible. Accordingly, we adopt a design science research (DSR) methodology to develop and evaluate a model (artifact) which supports evaluating data-driven decision outcomes, enabling collaborative rationality, and creating a feedback loop to enhance human and machine learning from past decisions. Hence, we attempt to research the following sub-questions:

RQ 1.1: *“Why do organizations need to evaluate collaborative rationality-based decisions?”* and *“How can collaborative rationality-based decisions be evaluated?”*

Here we research the theoretical underpinnings behind data-driven decision making and evaluation by reviewing the literature, in order to pave the road for developing more comprehensive evaluation solutions for human-machine collaborative rationality-based decisions.

RQ 1.2: *“What are the requirements and design objectives for ex-post evaluation of data-driven decisions in organizations?”*

Here we define the requirements and design objectives (DOs) for a data-driven decision evaluation solution, contributing to the first two stages of the DSR process (cf. [21]). First, we determine the relevant ex-post evaluation concepts from the literature. These concepts are then exemplified through an industrial case to foresee how ex-post evaluation of data-driven decisions could be done in practice, and accordingly outline the initial requirements for a design solution.

RQ 1.3: *“How can data-driven decisions be evaluated to enable feedback and learning loops, as well as enhance the quality of human-machine collaborative rationality-based decisions?”*

Here we develop and test a model, based on theory and practical case studies, as a design science artifact for evaluating data-driven decisions, thus completing stages 3-6 (design and development, demonstration, evaluation, and communication) of the DSR process (cf. [21]). This model accommodates for the multifaceted and changing nature of these decisions across contextual levels and provides a holistic perspective, which is currently lacking in literature. Moreover, the model depicts the relationship between the data-driven decision making elements, as well as the feedback and learning loops which ensue from ex-ante and ex-post decision evaluation. Accordingly, it provides a modular structure for the practical implementation of a data-driven decision evaluation solution, through which parts of the model can be adapted by organizations within their desired contexts.

This paper is structured as follows. In Section 2, we elaborate the research methodology followed and explain what has been, or will be, done in each of the DSR stages. Finally, in Section 3 we describe the expected contribution of this research to knowledge and practice.

2. Research methodology

The research methodology applied is design science, which consists of a rigorous process to design new artifacts intended to solve observed problems, make research contributions, evaluate the designs of the artifacts, and communicate results to appropriate audiences [21]. This process is followed throughout the research to develop and evaluate a theoretically sound and practically feasible model, or an abstraction that uses constructs to represent a real-world situation, depicting the relationship between the constructs and elements of data-driven decision making, including collaborative rationality between human and machine decision makers, as well as the feedback and learning loops resulting from evaluating data-driven decisions and their outcomes. We chose design science in particular, since our aim is not merely to explain or predict human or organizational behavior as is the core of the behavioral science paradigm, but we also seek to extend the boundaries of human and organizational capabilities by creating new and innovative artifacts which can be implemented to solve a problem [21,22].

In addition to perusing the knowledge base and the available literature, we work closely with decision makers and practitioners in multiple organizations to gather the business needs from the environment, as well as develop, demonstrate, and evaluate the artifact, hence achieving research relevance and rigor [22]. Hence, an organizational case was used to define the DOs of a solution, and multiple organizational cases are applied to demonstrate the model in different contexts.

Peffer's et al.'s [21] DSR process is followed, of which a simplification is shown in Figure 1. The process is iterative, allowing moving back and forth between the stages, with iterative evaluations according to Sonnenberg and vom Brocke's [23] design science research evaluation process in Figure 2. The stages of our research are further elaborated in the following subsections.

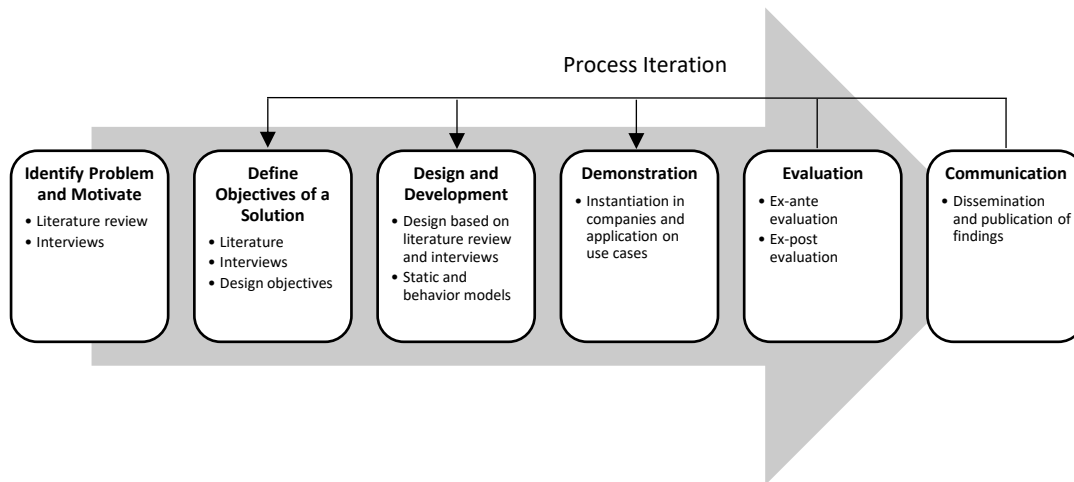


Figure 1: Design science research process (adapted from [20])

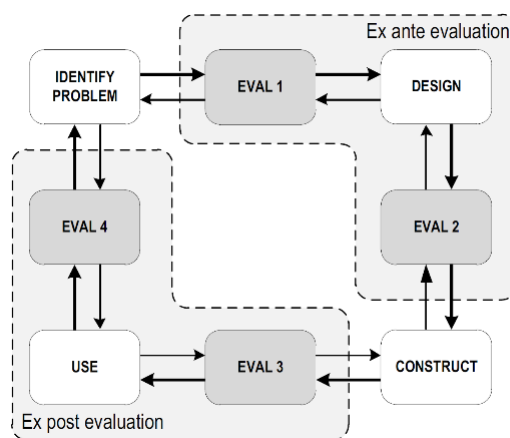


Figure 2: Design science research evaluation process [26]

2.1. Problem identification and motivation

We start with a problem-centered initiation by identifying the problem and motivation. This is mainly done by analyzing the literature and related research. The defined problem and motivation for research were evaluated and supported in practice through the expert interviews conducted and explained in the following subsection. While human-machine collaboration can potentially enhance decision making [6,7], more research is needed. Haphazard implementations without clear guidelines, criteria, and theory-backed methods are usually short-lived and destined to fail. Several failures have already been seen in organizations and agencies which rushed to automate their data-driven decision processes [9].

Moreover, it is perceived throughout research and practice that data-driven decision making can lead to the optimizing of decisions and result in more informed, quality decisions which could have otherwise been unattainable. Controversially, decision-making can also become more difficult when different combinations of data or analytics show different patterns, sometimes conflicting with the preferred choice of the decision-maker, and hence the decision-maker does not know how to proceed with the results [6].

While classical decision-making research focuses on the decision-making process, the decision maker, and the decision itself, the emergence of big data analytics has led to an evolution of modern data-driven decision making. Consequently, two new elements, the (big) data and the analytics need to be incorporated and integrated with the classical decision-making elements [9], as shown in Figure 3.

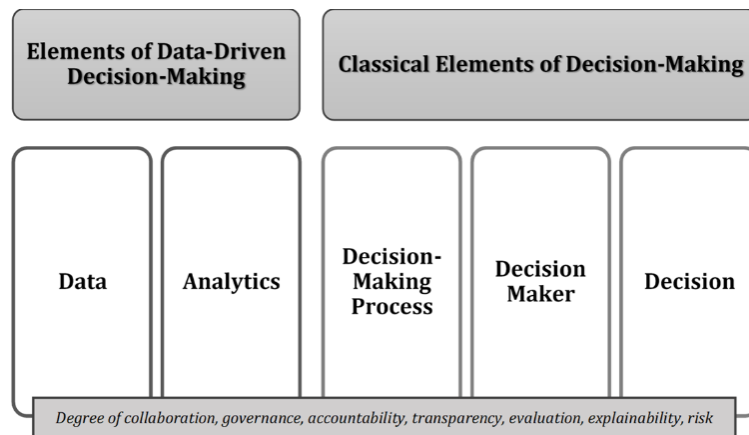


Figure 3: The elements of data-driven decision making [6]

However, a great deal of effort and research is required in order to meaningfully integrate these elements. Human and machine decision makers can coexist and collaborate to reach a higher level of rationality and maximize collaborative intelligence, unattainable by either one individually, which necessitates innovative models of decision making. Thus, the degree of collaboration between the human and the machine, the selection of quality (big) data, the appropriate use of analytics methods and tools, the definition of the decision-making process, and how to integrate all of these elements together, along with the resulting information and knowledge, are all imperative aspects to study for optimizing data-driven decisions and enhancing collaborative rationality between humans and machines.

Several recent studies have elaborated different modes of collaboration between humans and machines in organizational decision making for combining human and machine intelligence and capabilities, according to the context and types of the decision [4,10,11]. However, to organize these metahuman systems, new organizational functions are needed to delegate human vs. machine decisions, monitor how these decisions are taken, cultivate criteria to evaluate such decisions, and reflect through double-loop learning for continuous development [5].

This leads to collective intelligence, where future research on cognitive computing has been highlighted, with the goal of building a rational, combined, and collective mechanism motivated by the capability of the human mind and strengths of AI systems [24]. Accordingly, there is a synergy between the unique strengths of humans and machines, augmenting the intelligence of one another, however the level of collaboration differs according to the tasks and types of decisions on hand, which still requires future work [4].

Nevertheless, no comprehensive, or holistic, solutions accommodating for the multifaceted nature of collaborative data-driven decisions were found in the literature. The interaction between humans and machines and their roles in decision making is still not clear, the concept of collaborative rationality between humans and machines and its effect on decision making requires thorough elaboration, and further research is necessary to evaluate the resulting decisions and determine the benefit, impact, and learning that consequently occur from this collaboration [19]. Such an evaluation would provide information about the results within contextual boundaries and support more informed decisions in the future [13]. Furthermore, it can enable organizational and experiential learning [25,26], rationalization [27], and sensemaking [28] from the decision outcomes and consequences, as well as allow for analysis, benchmarking, and comparison of the results [29]. Iterations can be performed based on the knowledge gained from the impact of the decision, consequently creating a feedback loop between actions and outcomes.

Yet, with the various number of elements involved in complex, data-driven decisions, evaluation considering individual metrics or perspectives is insufficient. Accordingly, we can neither rely solely on the machine's evaluation of its models, nor the human's evaluation of their choices or judgements. Conversely, evaluation should consider the decision as a whole. However, many solutions focus on evaluating machine output rather than the entire decision, or evaluate with individual metrics, despite studies showing that even decisions with high accuracy are not necessarily correct and suffer in other

dimensions [30,31]. Moreover, limited evaluation metrics may be insufficient in considering multiple decision factors and are sometimes conflicting [32]. Hence, we need new ways to assess and evaluate the impact and benefits of data-driven decisions and human-machine collaboration from different perspectives [5,9,19,20]. Thus, novel theorizing in this area is crucial to provide a systematic understanding through an integrated conceptualization [20,33].

2.2. Defining the objectives of a solution

The objectives of a solution were defined both from theory and practice in [34]. The knowledge base was first perused to extract the theoretical concepts and factors relevant to data-driven decision making and evaluation from the appropriate streams of literature. These concepts include: embedded contexts (environmental, organizational, and decision contexts) in which the decision is made and which affects the decision and its evaluation, time (when and how often evaluation is conducted since the outcomes of the decision may vary across time), data driven-decision elements (decision maker, decision-making process, data, analytics/machine, and the decision outcome which needs to be evaluated), impact and consequences of the decision, conformance of the decision to certain criteria, metrics used for evaluating the decision, and errors and biases (both human and machine) which may affect the decision.

The concepts were then exemplified through an industrial example of a chemical production plant to foresee how ex-post evaluation of data-driven decisions could be done in practice and outline the initial requirements for a design solution. Two expert interviews were conducted to gather data on the plant's data-driven decisions and discuss the needs and requirements for a data-driven decision solution, from the viewpoints of both the production planner and process operator roles.

From the production planner's perspective, the data-driven decision involves determining and planning the production targets and capacities for a specific time interval and to schedule production. The purpose is to optimize the production rate and product portfolio to meet market demand. While these operational decisions are made short-term, they have a long-term impact. The human determines the objectives and constraints, and the ML tool supports the decision by simulating scenarios and suggesting alternative schedules. However, the human selects the best schedule to meet the designated criteria and makes the final decision (mode of collaboration: AI-recommends, human decides). Commonly, the human overlooks the output of the ML tool and decides not to use it, due to lack of trust in the reliability of the results.

From the process operator's perspective, the decisions include choices on set points for the process (such as feed rates and temperature), steering the process, and avoiding/overcoming fault situations. The purpose is to optimize the process efficiency (energy, material), avoid faults, and solve possible problems. Decisions are continuous and may also be triggered whenever there is a new target for the process, or a disturbance occurs. The ML tool provides outputs, insights, and predictions based on the data and process parameters. However, the inference and final decision are made by the human, who may use their own knowledge and expertise, along with additional monitoring methods (mode of collaboration: AI recommends, human decides and AI generates insights, which human uses in decision process).

The need for an ex-post evaluation solution was mainly to assess the reliability of the ML tool and increase trust, enhance both human and machine learning from evaluation feedback, evaluate decisions at different time intervals to indicate if the reliability of the tool increases, evaluate the monitoring methods, and evaluate uncertainties in the measurement data and their effect on the decision.

By comparing the current and desired approaches for decision evaluation stated in the interviews, the results were summarized into a set of evaluation requirements for each of the theoretical evaluation concepts previously defined from the literature. According to their functional similarities, the requirements were further thematized and mapped to more abstract and implementable DOs for an evaluation solution.

Accordingly, this study resulted in four DOs for a data-driven decision evaluation solution, shown in Table 1. The first DO for an implementable data-driven decision evaluation method is that it should be comprehensive and incorporate multi-faceted criteria ranging across different contextual levels. The second DO is that a processual evaluation is performed across different stages in time to accommodate changing contexts and consequences of the decision. The third DO suggests incorporating into the

evaluation the mode of collaboration between humans and machines and the consequent effect on decision making, the decision outcomes, and achieving collaborative rationality. Finally, the fourth DO would be to enable a potentially automated feedback loop which ensues from the evaluation. Decision makers then learn a new set of lessons from experiences and from evaluating the outcomes of their decision (link between actions and outcomes), which then leads to learning within the organization, as well as amongst other organizations [26]. This could potentially enhance machine learning as well since training data can be updated with the results of evaluation.

Table 1
Design objectives for ex-post evaluation of data-driven decisions [34]

No.	Design objective
1	Incorporate multi-faceted (potentially conflicting) evaluation criteria across contextual levels.
2	Perform processual evaluation across time.
3	Define the applicable mode of collaboration between humans and machines and evaluate its effect on decision-making, decision outcomes, and collaborative rationality.
4	Enable a (potentially automated) feedback loop for learning from the (discrete or continuous) evaluation of past decisions.

2.3. Artifact design and development (current stage)

The next stage in the research process is to design and develop an artifact based on the previous DOs. Our artifact, in this research, is a model providing a holistic, multifaceted, and multidisciplinary view for evaluating collaborative human-machine decisions. The essential constructs and attributes which are relevant to the model, as well as the underlying assumptions, were defined from the literature and previous theories.

An additional expert interview was conducted in another organization, focusing on data-driven decisions for predicting and preventing customer churn. This interview helped externally validate the design objectives and add new perspectives from a different case and decision context. Utilizing the theoretical concepts, and the knowledge gained from the three expert interviews, an initial version of a structural model for evaluating data-driven, or human-AI centric decisions, shown in Figure 4, was developed.

For structuring the model with a conceptual modeling grammar [35] we utilized the unified modeling language (UML) due to its demonstrated application for agent-based systems [36]. However, as controversial perspectives exist on the consistency, vagueness, and comprehensibility issues which plague UML notation and semantic representation [37], we have implemented our own minor adaptations as we saw fit to ensure simplicity and comprehension. Class diagrams are a type of structural diagram which describe the collection of declarative model elements and classes (in this case, our constructs), and their contents (their attributes) and relationships (association represented with solid lines, and feedback results represented with dashed lines) [36].

The model thus accommodates for the multifaceted and changing nature of data-driven decisions across contextual levels and provides a holistic perspective by depicting the relationship between the different concepts, such as the decision maker, machine agent, data, collaborative rationality, decision, criteria to which the decision should conform, ex-ante evaluation of the decision, ex-post evaluation of the decision, and the resulting feedback loop and types of learning enabled.

We assume recurring data-driven decisions, taking place within an organizational context, further embedded within an external, environmental context. A human decision maker and artificial agent (machine, analytics, etc.) are informed by a set of available data, of a certain quality. Both are characterized by a level of rationality, have a particular role in the decision-making process, have perceived models of the environment under which they operate, and a set of reference points which serve as baseline criteria, or governing variables, to which their decisions must conform. They are also prone to errors and biases. The human is further driven by emotions, which is a crucial attribute which the machine lacks.

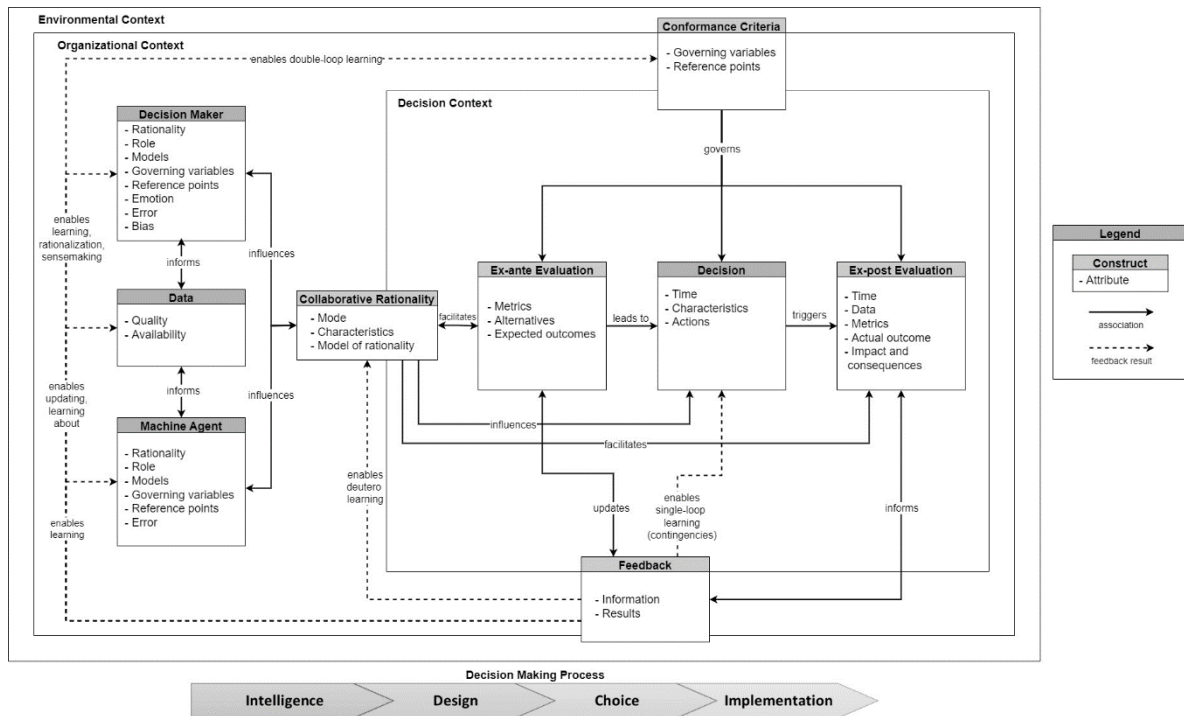


Figure 4: A model for data-driven decision evaluation

Together, the decision maker and machine, within an organizational context, engage in collaborative rationality, which can be utilized within the context of the decision, and in other contexts as well. The collaborative rationality is defined by the mode of collaboration, its characteristics, and its supported models of rationality. This collaborative rationality facilitates, and is facilitated by, the ex-ante evaluation, as new insights are generated through the evaluation of choices and alternatives, and it further supports ex-post evaluation, from which it is supported through the feedback loop. The ex-ante evaluation, utilizing its own set of metrics to evaluate alternatives, and the perceived outcomes leads to a decision. This is defined by the time of the decision, its characteristics, and the actions involved for implementation.

Subsequently, certain events resulting from the decision trigger an ex-post evaluation, utilizing another set of metrics and available data regarding the decision, at a specific time, to evaluate the actual outcomes and their impact and consequences. The ex-ante evaluation, the decision outcome, and the ex-post evaluation are within a particular decision context; however, they are governed by a set of conformance criteria spanning across multiple contexts. These criteria include reference values to which the entire decision should conform (e.g. benchmarks, KPIs, objectives, etc.).

The ex-post evaluation provides different types of feedback to enable experiential and organizational learning, either pertaining to the context of the decision or providing information and knowledge to other contexts as well. Such feedback serves as input to the other constructs and processes. First of all, it supports single-loop learning by updating the actions to resolve any errors. The feedback also provides knowledge for updating the ex-ante evaluation, and the opportunity to compare predicted outcomes with actual outcomes, thus making sense of the consequences and learning through experience. This could enhance future evaluations and predictions, as well as provide insight for revisiting the evaluation metrics and assumptions. Furthermore, learning from the feedback could entail updating the models themselves with which the alternatives were evaluated, which could lead to double-loop learning. Double-loop learning can occur when the feedback results in human or machine agents updating their reference points or perceived models of the environment, and when it leads to updating the conformance criteria for the decision.

Furthermore, this feedback can provide knowledge of the data required, and help detect discrepancies which may need to be solved. It can support rationalization and result in retrospective sensemaking. Finally, the feedback resulting from ex-post evaluation can enhance collaborative rationality by contributing to deuterio learning. Here the decision makers can reflect on their

collaborative, human-machine learning processes, and create new learning strategies for improving the collaborative rationality between them.

This model was then presented to an expert panel of four practitioners in the second organization and validated. An example instantiation of the model was created for the chemical manufacturing plant case, and an instantiation of the model was created for the customer churn case, together with two of the experts, as a prototypical evaluation. Finally, the model was presented to four experts in a third organization, with which we still aim to implement the model on their data-driven decision making case, thus achieving additional validation and fulfilling the Eval 2 stage in Sonnenberg and vom Brocke's [23] design science research evaluation process.

Since the structural model is static, and data-driven decision making is dynamic and requiring processual evaluation due to constantly changing contexts and consequences, a behavioral model depicting the relationship between the elements and their resulting behavior is still necessary and is the next stage in future work. It is intended to develop the behavioral model through collaboration with the third organization implementing data-driven decision making in order to observe the processual nature of decision making and evaluation, collaborative rationality between humans and machines, and the interaction between the elements and their impact on the decision.

2.4. Artifact demonstration and evaluation (future work)

The models should then be demonstrated through a prototypical instantiation in the context of an organization and evaluated (ex-post) in the Eval 3 and Eval 4 stages of Sonnenberg and vom Brocke's [23] design science research evaluation process. This is necessary to show that the models are applicable and useful in practice in various scenarios, can easily be integrated with the company's data, systems, and processes, their impact on the organization, and how they support or enhance data-driven decision making within the organization.

Future work thus includes practically implementing the model in a different case in the third organization and studying the relationship between the constructs. This will lead to developing the associated behavioral model to represent the dynamic aspects and the activities, sequences, flows, and implementable relationships between the elements and their resulting behavior, in addition to the behavior of collaborative rationality in each of the decision making stages through demonstration in a practical setting. By understanding these process, we can then influence their change in the desired directions [38].

Consequently, methods may be developed as a set of steps for manipulating the constructs so that the solution statement models are realized. Instantiations can further be used to operationalize the model and method, as a realization of the working artifacts in the environment [39], enabling practical implementation (and possible automation) and in-depth evaluation of the model. Furthermore, as we are currently in the theorizing stage, future research could help develop a more elaborate theory on data-driven decision systems and their evaluation, as well as on enhancing the collaboration between humans and machines. The results will be communicated and disseminated through publications.

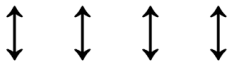
3. Expected contribution

Gregor and Hevner [40] proposed that design science research is neither limited to a particular type of artifact, nor to developing and testing a single artifact; but could rather include several artifacts with different levels of abstraction to demonstrate a contribution to knowledge, as shown in Table 2.

The design objectives contribute by guiding the development of a data-driven decision evaluation solution, based on knowledge gained from theory and practice. The static and behavioral models would serve as level 2 contributions, or nascent design theories. The instantiations of the models in organizations would serve as level 1 contributions. This would allow us to theorize and develop level 3 theories on data-driven decision making and collaborative rationality between humans and machines.

Table 2

Design science research contribution types [40]

	Contribution Types	Example Artifacts
More abstract, complete, and mature knowledge  More specific, limited, and less mature knowledge	Level 3. Well-developed design theory about embedded phenomena	Design theories (mid-range and grand theories)
	Level 2. Nascent design theory—knowledge as operational principles/architecture	Constructs, methods, models, design principles, technological rules.
	Level 1. Situated implementation of artifact	Instantiations (software products or implemented processes)

Hence, the main expected contributions of the research are as follows:

- Suggesting and elaborating the theoretical concept of collaborative rationality, as well as enhancing it through deutero learning which results from evaluating data-driven decisions.
- Highlighting the importance of a holistic, multi-faceted, and processual ex-post evaluation, while considering the different contexts, changing nature of the environment, and constantly challenging the evaluation criteria.
- Creating a feedback loop for enabling learning, rationalization, and sensemaking from data-driven decisions to enhance future decisions.
- Providing a modular structure for the practical implementation of the parts of the models and their relationship, resulting in a data-driven decision evaluation solution.

If the scope of the research allows, future work may also include developing a method for performing data-driven decision evaluation and design principles for potentially automating evaluation.

The contribution and benefit of this research in science and industry is a new approach to evaluating, and possibly automating the evaluation of, data-driven decisions, further supported by rigorous research methods. Accordingly, the impact of this research can help develop good practices in data-driven decision making, and enhance learning from previous data-driven decisions, both for the decision maker and the machine. It can aid decision makers and policy makers in utilizing analytics to extract insights, driving them to make better quality and more informed decisions, and helping them overcome the problems and challenges faced with current practices. Finally, it can provide a better understanding of the collaboration between humans and machines in decision-making across various levels and contexts, as well as enable a collaborative rationality between both.

4. Acknowledgements

I would like to acknowledge my supervisors:

Professor Tero Päivärinta, M3S, Faculty of Information Technology and Electrical Engineering, University of Oulu, Finland

Professor Ahmed Elragal, Department of Computer Science, Electrical and Space Engineering, Luleå University of Technology, Sweden

I would also like to acknowledge the ITEA3 project Oxilate (<https://itea3.org/project/oxilate.html>), and the organizations and experts who have willingly collaborated in this research.

5. References

- [1] T. Grønsund, M. Aanestad, Augmenting the algorithm: Emerging human-in-the-loop work configurations, *The Journal of Strategic Information Systems*. 29 (2020) 101614. <https://doi.org/10.1016/j.jsis.2020.101614>.

- [2] V. Grover, A. Lindberg, I. Benbasat, K. Lyytinen, The Perils and Promises of Big Data Research in Information Systems, *Journal of the Association for Information Systems*. 21 (2020). <https://doi.org/10.17705/1jais.00601>.
- [3] T. Sturm, J. Gerlach, L. Pumplun, N. Mesbah, F. Peters, C. Tauchert, N. Nan, P. Buxmann, Coordinating Human and Machine Learning for Effective Organization Learning, *MISQ*. 45 (2021) 1581–1602. <https://doi.org/10.25300/MISQ/2021/16543>.
- [4] A. Trunk, H. Birkel, E. Hartmann, On the current state of combining human and artificial intelligence for strategic organizational decision making, *Bus Res*. 13 (2020) 875–919. <https://doi.org/10.1007/s40685-020-00133-x>.
- [5] K. Lyytinen, J.V. Nickerson, J.L. King, Metahuman systems = humans + machines that learn, *Journal of Information Technology*. 36 (2020) 427–445. <https://doi.org/10.1177/0268396220915917>.
- [6] M. Janssen, H. van der Voort, A. Wahyudi, Factors influencing big data decision-making quality, *Journal of Business Research*. 70 (2017) 338–345. <https://doi.org/10.1016/j.jbusres.2016.08.007>.
- [7] D.J. Power, Data science: supporting decision-making, *Journal of Decision Systems*. 25 (2016) 345–356. <https://doi.org/10.1080/12460125.2016.1171610>.
- [8] F. Provost, T. Fawcett, Data Science and its Relationship to Big Data and Data-Driven Decision Making, *Big Data*. 1 (2013) 51–59. <https://doi.org/10.1089/big.2013.1508>.
- [9] N. Elgendy, A. Elragal, T. Päivärinta, DECAS: a modern data-driven decision theory for big data and analytics, *Journal of Decision Systems*. (2021) 1–37. <https://doi.org/10.1080/12460125.2021.1894674>.
- [10] S. Ransbotham, S. Khodabandeh, D. Kiron, F. Candelon, M. Chu, B. LaFountain, Expanding AI's Impact With Organizational Learning, *MIT Sloan Management Review*. (2020). <https://sloanreview.mit.edu/projects/expanding-ais-impact-with-organizational-learning/>.
- [11] Y.R. Shrestha, S.M. Ben-Menahem, G. von Krogh, Organizational Decision-Making Structures in the Age of Artificial Intelligence, *California Management Review*. 61 (2019) 66–83. <https://doi.org/10.1177/0008125619862257>.
- [12] E.B. Mandinach, A Perfect Time for Data Use: Using Data-Driven Decision Making to Inform Practice, *Educational Psychologist*. 47 (2012) 71–85. <https://doi.org/10.1080/00461520.2012.667064>.
- [13] C. Herrick, D. Sarewitz, Ex Post Evaluation: A More Effective Role for Scientific Assessments in Environmental Policy, *Science, Technology, & Human Values*. 25 (2000) 309–331. <https://doi.org/10.1177/016224390002500303>.
- [14] D.J. Power, *Decision support systems: concepts and resources for managers*, Quorum Books, Westport, Conn, 2002.
- [15] R. Afiouni-Monla, Organizational Learning in the Rise of Machine Learning, in: *ICIS 2019 Proceedings*, 2019: p. 18.
- [16] T. Sturm, T. Koppe, Y. Scholz, P. Buxmann, The Case of Human-Machine Trading as Bilateral Organizational Learning, (2021) 18.
- [17] E. Klecun, T. Cornford, A critical approach to evaluation, *European Journal of Information Systems*. 14 (2005) 229–243. <https://doi.org/10.1057/palgrave.ejis.3000540>.
- [18] R. Stockdale, C. Standing, An interpretive approach to evaluating information systems: A content, context, process framework, *European Journal of Operational Research*. 173 (2006) 1090–1102. <https://doi.org/10.1016/j.ejor.2005.07.006>.
- [19] Y.K. Dwivedi, L. Hughes, E. Ismagilova, G. Aarts, C. Coombs, T. Crick, Y. Duan, R. Dwivedi, J. Edwards, A. Eirug, V. Galanos, P.V. Ilavarasan, M. Janssen, P. Jones, A.K. Kar, H. Kizgin, B. Kronemann, B. Lal, B. Lucini, R. Medaglia, K. Le Meunier-FitzHugh, L.C. Le Meunier-FitzHugh, S. Misra, E. Mogaji, S.K. Sharma, J.B. Singh, V. Raghavan, R. Raman, N.P. Rana, S. Samothrakakis, J. Spencer, K. Tamilmani, A. Tubadji, P. Walton, M.D. Williams, Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy, *International Journal of Information Management*. 57 (2021) 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>.
- [20] Y. Duan, J.S. Edwards, Y.K. Dwivedi, Artificial intelligence for decision making in the era of Big Data – evolution, challenges and research agenda, *International Journal of Information Management*. 48 (2019) 63–71. <https://doi.org/10.1016/j.ijinfomgt.2019.01.021>.

- [21] K. Peffers, T. Tuunanen, M.A. Rothenberger, S. Chatterjee, A Design Science Research Methodology for Information Systems Research, *Journal of Management Information Systems*. 24 (2007) 45–77. <https://doi.org/10.2753/MIS0742-1222240302>.
- [22] A.R. Hevner, S.T. March, J. Park, S. Ram, Design Science in Information Systems Research, *MIS Quarterly*. 28 (2004) 75–105. <https://doi.org/10.2307/25148625>.
- [23] C. Sonnenberg, J. vom Brocke, Evaluations in the science of the artificial --- reconsidering the build-evaluate pattern in design science research, in: *Proceedings of the 7th International Conference on Design Science Research in Information Systems: Advances in Theory and Practice*, Springer-Verlag, Berlin, Heidelberg, 2012: pp. 381–397. https://doi.org/10.1007/978-3-642-29863-9_28.
- [24] P. Gupta, A.W. Woolley, Articulating the Role of Artificial Intelligence in Collective Intelligence: A Transactive Systems Framework, *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. 65 (2021) 670–674. <https://doi.org/10.1177/1071181321651354c>.
- [25] C. Argyris, D.A. Schön, *Organizational Learning: A Theory of Action Perspective*, Addison-Wesley. (1997) 345. <https://doi.org/10.2307/40183951>.
- [26] J.G. March, *Primer on Decision Making: How Decisions Happen*, Simon and Schuster, 1994.
- [27] J.G. March, Bounded Rationality, Ambiguity, and the Engineering of Choice, *The Bell Journal of Economics*. 9 (1978) 587–608. <https://doi.org/10.2307/3003600>.
- [28] K.E. Weick, *Sensemaking in organizations*, Sage Publications, Thousand Oaks, 1995.
- [29] E.M. Masha, The Case for Data Driven Strategic Decision Making, *European Journal of Business and Management*. (2014) 10.
- [30] A. Köchling, S. Riazzy, M.C. Wehner, K. Simbeck, Highly Accurate, But Still Discriminatory: A Fairness Evaluation of Algorithmic Video Analysis in the Recruitment Context, *Bus Inf Syst Eng*. 63 (2021) 39–54. <https://doi.org/10.1007/s12599-020-00673-w>.
- [31] M. Teodorescu, L. Morse, Y. Awwad, G. Kane, Failures of Fairness in Automation Require a Deeper Understanding of Human-ML Augmentation, *MISQ*. 45 (2021) 1483–1500. <https://doi.org/10.25300/MISQ/2021/16535>.
- [32] S. Lebovitz, N. Levina, H. Lifshitz-Assa, Is AI Ground Truth Really True? The Dangers of Training and Evaluating AI Tools Based on Experts’ Know-What, *MISQ*. 45 (2021) 1501–1526. <https://doi.org/10.25300/MISQ/2021/16564>.
- [33] H. Benbya, S. Pachidi, S.L. Jarvenpaa, Special Issue Editorial: Artificial Intelligence in Organizations: Implications for Information Systems Research, *Journal of the Association for Information Systems*. 22 (2021) 281–303. <https://doi.org/10.17705/1jais.00662>.
- [34] N. Elgendy, A. Elragal, M. Ohenoja, T. Päiväranta, Ex-Post Evaluation of Data-Driven Decisions: Conceptualizing Design Objectives, in: Ę. Nazaruka et al. (Ed.), *Perspectives in Business Informatics Research*, Springer Nature Switzerland, 2022. https://doi.org/10.1007/978-3-031-16947-2_2.
- [35] Y. Wand, R. Weber, Research Commentary: Information Systems and Conceptual Modeling—A Research Agenda, *Information Systems Research*. 13 (2002) 363–376. <https://doi.org/10.1287/isre.13.4.363.69>.
- [36] B. Bauer, J. Odell, UML 2.0 and agents: how to build agent-based systems with the new UML standard, *Engineering Applications of Artificial Intelligence*. 18 (2005) 141–157. <https://doi.org/10.1016/j.engappai.2004.11.016>.
- [37] J. Erickson, A Decade and More of UML: An Overview of UML Semantic and Structural Issues and UML Field Use, (2008) 8.
- [38] J. vom Brocke, W. van der Aalst, T. Grisold, W. Kremser, J. Mendling, B. Pentland, J. Recker, M. Roeglinger, M. Rosemann, B. Weber, *Process Science: The Interdisciplinary Study of Continuous Change*, Social Science Research Network, Rochester, NY, 2021. <https://doi.org/10.2139/ssrn.3916817>.
- [39] S.T. March, G.F. Smith, Design and natural science research on information technology, *Decision Support Systems*. 15 (1995) 251–266. [https://doi.org/10.1016/0167-9236\(94\)00041-2](https://doi.org/10.1016/0167-9236(94)00041-2).
- [40] S. Gregor, A.R. Hevner, Positioning and Presenting Design Science Research for Maximum Impact, *MIS Quarterly*. 37 (2013) 337–355.