# DERIVING DECISION MINING SYSTEM CAPABILITIES: A RESEARCH AGENDA

KOEN SMIT, SAM LEEWIS, MATTHIJS BERKHOUT,
JOHN VAN MEERTEN, CHAIM DE GELDER,
SUSAN BRUGGELING, HANNE DE DECKERE,
ANNEMAE VAN DE HOEF

HU University of Applied Sciences Utrecht, Digital Ethics, Ultrecht, Netherland
koen.smit@hu.nl, sam.leewis@hu.nl, matthijs.berkhout@hu.nl,
john.vanmeerten@hu.nl, chaim.degelder@hu.nl, susan.bruggeling@hu.nl,
hanne.dedeckere@hu.nl, annemae.vandehoef@hu.nl

Decision Mining (DM) is increasingly gaining attention from academia and slowly progressing towards instrumental application in practice by leveraging decision logs to automatically discover, check for conformance and improve derivation patterns for operational decision-making. This study aims to further operationalize DM by identifying capabilities in the form of functional and non-functional requirements that are posed in the current body of knowledge. By identifying and analysing DM contributions with a focus on derivation patterns we were able to point out the aspects of DM getting attention as well as which did not, e.g., a strong focus on input data and algorithms regarding the discovery phase while the output (data) of the improvement phase seems to be detailed insignificantly. Based on this we formulated a research agenda in which five key points for future research studies are presented.

**Keywords:** decision mining, capabilities, systems, requirements, research agenda, literature

University of Maribor Press

DOI https://doi.org/10.18690/um.fov.6.2023.32
ISBN 978-961-286-804-8

## 1 Introduction

Proper operational decision-making is one of the most important capabilities of an organization (Mircea et al., 2012). Especially when organizations designed products and services that focus on high volumes of data processing for operational decision-making (Chalvatzis et al., 2019; Rula et al., 2016), such as governmental services focusing on calculation and application of benefits or financial services for opening accounts at a bank. Adequately managing these operational decisions is becoming increasingly difficult as the digitization of products and services as well as the transition towards fully automated operational decision-making becomes more prevalent. One way of reducing complexity in IT is to separate concerns such as 'data', 'user interfacing', and 'processes' from each other (Dijkstra, 1974; Ossher & Tarr, 2001), which resulted in separate systems to do so as well. A whole research field now focuses on separating decisions, sometimes also referred to as 'the logic', so that its management can be simplified and made more explicit (Bajec & Krisper, 2005; Boyer & Mili, 2011; De Smedt et al., 2017; Graham, 2006; Schlosser et al., 2014; Smit, 2018).

The next step in properly managing decisions is to optimally use data to improve decisions that are explicitly managed, which is similar to what happened in the field of business process management, referred to as Process Mining (van der Aalst, 2011). In the context of separating the concern of 'decisions' and 'logic', a similar approach is referred to as Decision Mining (DM), which is defined as: "*the method of extracting and analyzing decision logs with the aim to extract information from such decision logs for the creation of business rules, to check compliance to business rules and regulations, and to present performance information*" (Leewis, Smit, et al., 2020). DM comprises 1) Discovery, 2) Conformance Checking, and 3) Improvement of decisions and underlying logic (Leewis, Smit, et al., 2020). These phases are very similar to Process Mining. The goal of process mining is to discover process models, check their conformance to theoretical process models and improve process models based on the outcome of the first two phases.

The present state of the research field and the knowledge base on Data Mining (DM) is currently deemed inadequate by scholars (Leewis, Smit, et al., 2020; Vanthienen, 2021). However, there is an increasing interest in DM, particularly in academic circles (Goossens et al., 2023; Scheibel & Rinderle-Ma, 2022), where there is a significant emphasis on quality metrics such as the accuracy of the mining algorithm. However, the comprehensibility of the application of DM algorithms to stakeholders like analysts or end-users is often ignored (Vanthienen, 2021). Therefore, it would be worthwhile to investigate the capabilities of a

DM system, as this would help to advance the research field by putting into practice the theoretically proposed stages that are currently presented in a high level of abstraction in the existing body of knowledge. Within this context, a capability is defined as "*an ability that an organization, person, or system possesses*" (The Open Group, 2011).

In this study, a first step is made towards operationalizing capabilities for discovery, conformance checking, and improvement as part of DM. This paper comprises the exploration of capabilities derived from the current body of knowledge on DM. In a follow-up study we aim to derive capabilities using an empirical approach so that both inputs can be compared and a final set of validated capabilities can be presented so that organizations are able to more swiftly experiment with DM. Therefore, in this paper, we answer the following research question: '*What capabilities can be identified from DM literature focused on the decision viewpoint for the development of DM systems and how should future research into DM systems be conducted?*'

The remainder of the paper is structured as follows. In the following section, section two, we discuss the phases of DM in more detail and explore the current state-of-the-art on DM. Then, in section three, we present the research method used in this study to derive capabilities for the DM system, which we use as a basis to formulate a research agenda. This is followed by the data collection and data analysis in section four. Then, in section five, the identified DM system capabilities are presented. The limitations of the study and its results are presented in section six. Lastly, the paper is concluded in section seven by presenting a research agenda.

## 2 Background

DM is also referred to as decision point analysis, which "*aims at the detection of data dependencies that affect the routing of a case*"(Rozinat & van der Aalst, 2006). The difference between the two is that decision point analysis mines 'sequencing patterns' from a process viewpoint, while DM, in line with the definition presented in the previous section, mines 'derivation patterns' from a decision viewpoint (Leewis, Smit, et al., 2020). Previous studies indicate that the focus towards DM and mining on derivation patterns is necessary to take steps forward in DM (De Smedt et al., 2017; Leewis, Smit, et al., 2020; Sarno et al., 2013).

DM consists of three phases (Leewis, Smit, et al., 2020): discovery of decisions, conformance checking of decisions, and the improvement of decisions, see Figure 1. These phases comprise extracting information from decision logs (discovery), checking this information for compliance with business rules and regulations (conformance), and presenting possible performance information (improvement) (Leewis, Berkhout, et al., 2020).

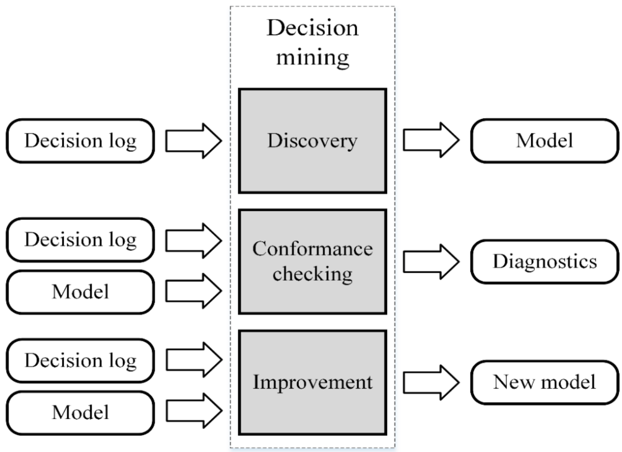


**Figure 1: DM phases with corresponding inputs and outputs**

As depicted in Figure 1, the input data necessary to leverage DM is the 'decision log' (De Smedt et al., 2017; Leewis, Smit, et al., 2020). The structure of this specific log file is essential for the success of DM. De Smedt et. al. (2017) introduced the necessity of decision logs in order to properly mine decisions and Leewis et. al. (2022)described the importance of a decision log for mining decisions with a decision viewpoint in mind. Additionally, DM aims at mining decisions from structured data. Structured data has a predefined data model compared to unstructured data. Therefore, a structure needs to be defined in the form of a decision log.

The structure of a decision log consists of a minimum of four attributes. The first attribute is a unique identifier, for example, a Case ID. Moreover, the condition(s) (second attribute) and conclusion(s) (third attribute) connected to that Case ID are part of the decision log, for example, all the data that is necessary for acquiring a loan and the decisions made with that data. The last required attribute is a timestamp. For example, the time certain input data was received or a decision was made by an actor. An example of a decision log is shown in Figure 2.

| | | Input values | | | Output value |
|---|---|---|---|---|---|
| | | *Conditions* | | | *Conclusion* |
| **ID** | **Timestamp** | **Gender** | **Age** | **Temperature** | **Treatment** |
| 1 | 2021-22-10 | Male | 30 | 37,4 | Treatment A |
| 2 | 2021-26-10 | Female | 25 | 39,1 | Treatment B |
| 5 | 2021-30-10 | Female | 50 | 38,7 | Treatment A |
| 7 | 2021-21-10 | Male | 19 | 37,2 | Treatment A |
| 8 | 2021-22-10 | Male | 18 | 36,8 | Treatment A |
| 10 | 2021-26-10 | Female | 24 | 37,1 | Treatment B |
| 11 | 2021-28-10 | Female | 20 | 37,6 | Treatment B |

**Figure 2: Example of a decision log with the required attributes**

Despite the current focus on DM with structured data, more and more research is being conducted on extracting and discovering decision logic and decision dependencies in (semi) unstructured data such as laws and regulations or protocols and guidelines (Etikala et al., 2020; Goossens et al., 2022; Vanthienen, 2021). While mining decision logic from text using text mining algorithms and remodelling them to a decision table is a complex task in itself, finding relations between the individual decision tables by mining dependencies, among the other elements on the requirements level, is even more complex and is currently being explored (Vanthienen, 2021).

Recent studies also focus on the conformance-checking phase of DM where anomalies in individual decision tables, as well as decision requirements diagrams, are found (Batoulis & Weske, 2018; Corea & Delfmann, 2018; Smit et al., 2017). The techniques described in these studies deal with the completeness and consistency of the individual rules, the decision table as a whole, and the dependencies between decisions. All models can be completely checked and audited, and anomalies can be found. Additionally, they can even be (semi) automatically altered (Corea & Delfmann, 2018). We argue that the current body of knowledge, although with a limited focus on the decision viewpoint, contains useful pointers for the operationalization towards capabilities.

One way to define capabilities is by formulating functionalities that a system supporting the DM phases must possess. In software engineering, requirements are used to express functionalities (Kotonya & Sommerville, 1998). Different types of requirements exist, for example, functional requirements, non-functional requirements, business requirements, user requirements, and constraints

(Sommerville & Sawyer, 1997). In this paper, we solely focus on functional and non-functional requirements as these are recognizable archetypes (Kotonya & Sommerville, 1998). Furthermore, a functional requirement emphasizes *what* is required and a non-functional requirement emphasizes the general properties of a system. This does not define the *how*, which is in line with the notion of a capability, that also focuses on *what* (value) an organization can deliver, but not *how* the value is delivered. In this study, we use thematic coding to give meaning to a wide range of possible functional requirements and are domain-specific. For non-functional requirements we use the ISO25010 standard comprising system and software quality clustering, which is a thoroughly validated framework (ISO, 2014). Examples of clusters are Reliablity (e.g., availability, fault tolerance) and operability (e.g., learnability, technical accessibility).

## 3 Research method

In order to identify functional and non-functional requirements for the development of DM systems we start, in this study, with a thorough analysis of the scientific literature. As described earlier, the focus lies on requirements which are defined for solutions specified from a decision viewpoint. Therefore, these sources need to be identified and coded with regard to functional and non-functional requirements. Because the field of DM is in a nascent state (Leewis, Smit, et al., 2020) the body of knowledge is limited. Even more so when limiting contributions to the decision viewpoint. Therefore, the study is considered an explorative one.

## 4 Data Collection & Analysis

The data collection of relevant contributions focussed on DM from a decision viewpoint focusing on deriving derivation patterns took place during November 2022. For this search, two criteria were adhered to exclude non-relevant contributions; 1) contributions referring to the utilization of event logs only to derive sequence data, and 2) papers referring to decision-point analysis as a form of DM. The first criterion was used because it excludes contributions that focus (predominantly) on process mining, which is not the focus of this study. The second criterion was used because decision-point analysis is considered a form of DM, which however, aims at deriving sequencing patterns in the context of business processes and business process management, i.e., how a sequence of a business process is routed based on one or more conditions. Google Scholar has been used as our search engine. The exclusion criteria were as follows: 1) the contribution focuses on decision-point analysis (thus using a process instead of a decision perspective), 2) the contribution is accessible for the research team, and 3) the paper is written in English. The search resulted in eight relevant contributions, presented in Table 1.

The analysis of the identified contributions started by identifying functional and non-functional requirements through thematic coding as specified by Gibbs (Gibbs, 2007). We employed four individual coders that used a pre-defined coding scheme, based on the definitions of functional and non-functional requirements from (Sommerville & Sawyer, 1997): "*functional requirements describe what the system should do and non-functional requirements place constraints on how these functional requirements are implemented.*"

**Table 1: Identified contributions focusing on the decision viewpoint**

| ID: | Title: | Reference: |
|---|---|---|
| 1 | Utilizing Algorithms for Decision Mining Discovery | (Berkhout & Smit, 2022) |
| 2 | Decision Mining versus Process Mining: a Comparison of Mining Methods | (de Jong et al., 2021) |
| 3 | Decision Mining in a Broader Context: An Overview of the Current Landscape and Future Directions | (De Smedt et al., 2017) |
| 4 | Business Rules Management and Decision Mining- Filling in the Gaps | (Leewis et al., 2022) |
| 5 | Deep Learning for the Identification of Decision Modelling Components from Text | (Goossens et al., 2021) |
| 6 | Putting Decision Mining into Context: A Literature Study | (Leewis, Smit, et al., 2020) |
| 7 | Extracting Decision Model and Notation models from text using deep learning techniques | (Goossens et al., 2023) |
| 8 | Future challenges in decision mining at governmental institutions | (Leewis, Berkhout, et al., 2020) |

Three coding rounds were conducted in order to reliably identify functional (F) and non-functional (NF) requirements from the identified contributions, depicted in Figure 3. During the first round of coding, all coders coded all contributions with regard to functional and non-functional requirements separately. An example of a functional requirement code is: "*all main functionalities: text classification, decision dependency extraction and decision logic extraction.*" An example of a non-functional requirement code is: "*The output of an algorithm must be explainable and comprehensible by Subject Matter Experts.*"

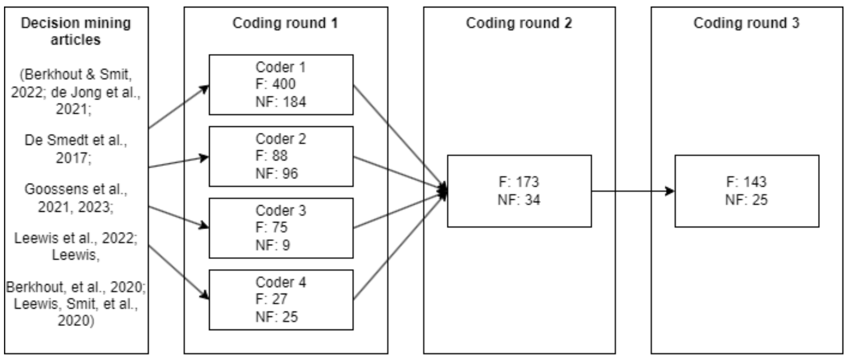


**Figure 3: Coding process and results per step**

A total of 590 functional requirements and 157 non-functional requirements were coded during coding round one. Table 2 shows the coding results of each coder, for each contribution, regarding both functional and non-functional requirements.

The second round of coding is used as a consolidation round. During this round, all coders discussed the individual coding results from the first round. Consensus or disagreement was focused on keeping or removing individual codes. The second coding round resulted in 173 functional requirements (417 were identical or removed) and 34 non-functional requirements (123 were identical or removed).

**Table 2: Coding totals**

| | Coder 1: | | Coder 2: | | Coder 3: | | Coder 4: | |
|---|---|---|---|---|---|---|---|---|
| Article ID: | F | NF | F | NF | F | NF | F | NF |
| 1 | 10 | 1 | 8 | 2 | 6 | 1 | 0 | 0 |
| 2 | 60 | 2 | 34 | 24 | 25 | 0 | 14 | 5 |
| 3 | 31 | 4 | 2 | 9 | 6 | 0 | 6 | 4 |
| 4 | 13 | 7 | 12 | 2 | 15 | 0 | 1 | 0 |
| 5 | 62 | 0 | 13 | 10 | 5 | 0 | 1 | 0 |
| 6 | 25 | 0 | 1 | 11 | 5 | 0 | 3 | 0 |
| 7 | 173 | 4 | 8 | 15 | 6 | 0 | 1 | 0 |
| 8 | 26 | 9 | 10 | 23 | 7 | 8 | 1 | 16 |

Then, a third coding round was used to further validate the results of the first two coding rounds. This was done by one of the original coders that worked on coding during the first two rounds but was accompanied by two senior researchers with more experience in the field of study. This resulted in a total of 143 functional requirements (30 were removed) and 25 non-functional requirements (9 were removed). In this last round, we used an additional coding scheme presented in Table 3. The coding scheme is based on the three DM phases on the x-axis, while the y-axis comprises three attention areas concerning DM that we chose to further explore concerning the research agenda for further operationalisation of DM. The attention areas are selected based on the fact that, for each of the DM phases on the x-axis, the algorithm needs specific 1) *input data*, characterized by requirements (e.g., a minimum of one condition and one conclusion). Furthermore, there are requirements for what the 2) *algorithm* itself should be able to do with the input data (e.g., which transformative steps have to be taken). The transformation of data then results in certain 3) *output data*, which is characterized by requirements as well (e.g., which types of data need to be presented in what manner).

**Table 3: Round 3 functional requirements coding scheme**

| | Discovery | Conformance Checking | Improvement |
|---|---|---|---|
| Input data | - Code 1<br>- Code 2<br>- Code n | ...<br>...<br>... | ...<br>...<br>... |
| Algorithm | ... | ... | ... |
| Output data | ... | ... | ... |

## 5 Results

Based on the coding of our data we describe the functional requirements in this section using the three phases of DM. Additionally, we describe the non-functional requirements according to the ISO25010 categories. Due to space limitations, we summarize the results for each category.

### 5.1 Discovery phase

**Input data – 26 requirements**

As the second-largest category by requirements identified, many assumptions and requirements were posed in the literature about how the input data should (ideally) be structured so that it can be used for the discovery of decisions and underlying logic. All contributions seem to focus on the need for a decision log and refer to the same composition of a decision log used to generate output, including implicit or explicit conditions, conclusions, timestamps, and dependencies. An example of a coded fragment referring to this category is: *"The DMN model extraction tool takes as input a decision description and automatically classifies sentences into irrelevant(for the decision model), dependency or logic."*

**Algorithm – 51 requirements**

The predominant category of requirements applies to the algorithmic technique used for decision discovery. The found literature revealed specific specifications that decision discovery algorithms must satisfy. These specifications apply to both structured data, which utilizes event or decision logs, and (semi) unstructured data, such as laws and regulations or textual descriptions. While there is significant overlap in the requirements presented, minor differences are noticed. For instance, some algorithms must create a decision requirements diagram (DRD) with supporting elements prior to discovering the decision logic, while others first discover the decision logic and then create a DRD based on the decision logic. An example of a coded fragment referring to this category is: *"The construction of a decision model from text requires a sequence of steps, each with their own challenges, regardless of whether a human or a machine is performing it: coreference resolution (where all expressions that refer to the same entity are resolved); preprocessing (preparing the data for analysis); text classification (identifying the* *relevant sentences for a modelling problem) and decision dependency and logic extraction (identifying the relevant elements needed for the construction of the model).*"

**Output data – 13 requirements**

The requirements identified in this category entail the transformation of data into various representations, guided by the construction of a decision model. The resulting output comprises a DRD, decision tables, and/or business rules. Additionally, a crucial aspect of the output data is that the algorithms should indicate the under or overfitting of data. An example of a coded fragment referring to this category is: *"The output of the Discovery phase is a business decision architecture, e.g., a DRD, as well as decision tables, and business rules."*

## 5.2 Conformance Checking phase

**Input data – 8 requirements**

The requirements identified in the literature highlighted the relevant components necessary for conformance checking. DM utilizes decision logs and models, such as DMN models, during the conformance-checking phase. Therefore, the input for the conformance-checking activity consists of the log file and the discovered model, which usually emerges from the output of the discovery activity. An example of a coded fragment referring to this category is: *"Decision conformance checking has the same purpose and consists of the same input components as described for Process mining (log file and discovered model).*"

**Algorithm – 10 requirements**

Requirements in this category emphasize the diagnostic aspect of conformance checking. The algorithms or techniques used must possess the capability to identify and quantify discrepancies between the model and associated log files. Conformance checking should function to identify, locate, and detect deviations. Furthermore, by analyzing the conditions present in the decision log, the model can be utilized to detect bottlenecks and incorrect dependencies on both the DRD level as well as decision logic level. An example of a coded fragment referring to this category is:

*"Conformance checking analyzes whether the reality, as recorded in a logbook, corresponds to the model and vice versa. The aim is to measure their severity and to detect abnormalities.*"

**Output data – 5 requirements**

As the smallest category of identified requirements, the focus of the discovered requirements for output data was diagnostic in nature. For example, the diagnosis, in the form of output data, examines whether decisions are executed as intended, providing diagnostic information that highlights differences and similarities between the model and the input data. An example of a coded fragment referring to this category is: "*The output consists of diagnostic information that shows differences and similarities between the model and the input data (log files).*"

## 5.3 Improvement phase

**Input data – 10 requirements**

The majority of the found assumptions and requirements in this category overlap with conformance checking. The focus is on the utilization of structured data within an algorithm to enhance or improve discovered models in DM activities. A decision log is employed as input but the improvement phase also needs the inclusion of both a decision log and a decision model, which is the same as conformance checking. An example of a coded fragment referring to this category is: *"The enhancement/improvement activity needs an event or decision log and a model as input.*"

**Algorithm – 8 requirements**

As the second smallest category coded, most requirements found in the literature were about providing potential improvements to an existing model ultimately resulting in a new model. The algorithm must support the changing or extending of a decision model, based on a decision log and the theoretical model. An example of a coded fragment referring to this category is: *" enhancement/improvement activity aims at changing or extending the model."*

**Output data – 12 requirements**

The last category coded is output data during the improvement phase. The requirements coded in this category were mainly about a new model that is created using the input of a discovered model and the decision log. This phase must not only identify and present improvements but also outputs a new model with the improvements, which acts as a basis for revision during the improvement phase. An example of a coded fragment referring to this category is: "*In the enhancement/ improvement activity a new process/ decision model is created or an existing process or decision model is adapted.*"

## 5.4 Non-functional requirements

In terms of non-functional requirements, most codes referred to the usability (8 codes) and reliability (12 codes). Regarding usability, we see that contributions mention different stakeholder groups from different domains that should be able to work with DM systems. Therefore, the output of the algorithm should be explainable and understandable. For example, one of the coded fragments referring to this is: "*Further focus seems required to ensure a user-friendly interface where non-experts could use the capabilities of DM and thereby not confronted with algorithms where expert interpretation is needed.*" Another contribution mentioned that offering too much transparency can pose a risk to the accuracy of the DM algorithm used, which should be further explored. Regarding reliability, many generic data science constraints are identified, such as data quality (comprising the decision log), avoidance of data contamination, over and underfitting of DM algorithms, DM algorithm accuracy levels, (sample) data representativeness, and detection of outliers. An example of a coded fragment referring to this is: "*High validity can be ensured by utilizing accurate and reliable DM techniques (ensuring internal validity), and utilizing a data sample representative towards the population when decision mining (ensuring external validity).*"

# 6 Discussion

This study and its results have some limitations that should be discussed. Firstly, this study only aimed to identify capabilities from a theoretical perspective. Although the eight sources have some empirical basis, most of them are conceptual in nature and should be further supplemented and validated using empirical observations in a realistic context, i.e., using case studies. This also helps integrate DM processes into other IT processes as the current contributions on DM with a decision viewpoint mostly examine DM systems as a standalone phenomenon. A second limitation is the low number of contributions in the current body of knowledge with a focus on the decision viewpoint that aims to identify derivation patterns. This severely limits the theoretical richness of capabilities identified in this paper, but also points out a gap in the current body of knowledge that should be covered by future research. Another direction for future research would be to include contributions from the process mining research field that focus on the operationalisation of the three phases and look at systems for process mining specifically, a perspective that has not been included in this study. A third limitation would be the overestimation of depth and richness we expected the current body of knowledge to have in terms of potential for operationalisation of capabilities into requirements. If we look at the results from the first coding round alone, we identified many coded fragments of concepts that are described at a high level of abstraction. This also limits the value of the (patterns of) requirements we unearthed in this study, although we argue that the results are still a useful starting point for the development of tooling to support DM systems as well as that it serves as a research agenda for what needs further consideration in future research. The high level of abstraction adhered to in the identified contributions is presumably also caused by the focus of these studies as they are not meant to operationalize capabilities to the level we pursued to do so in this study. Thus, this observation also calls for future research focused on applying DM capabilities and their operationalization in DM systems in practice.

## 7 Conclusion

To conclude this paper we revisit the research question posed in the introduction section: '*What capabilities can be identified from Decision Mining literature focused on the decision viewpoint for the development of decision mining systems and how should future research into DM systems be conducted?*' Based on a thorough analysis of the current body of knowledge on DM focusing on the decision viewpoint we derived requirements on nine areas on a functional level and used the ISO 25010 software quality standard to identify a predominant focus on usability and reliability non-functional requirements for DM systems. From a theoretical perspective, our results point out directions for future research. Additionally, we demonstrated that our theoretical model of DM phases and focus areas for the operationalisation of DM system capabilities could be a useful approach, which can be used in future studies as suggested further in this section. From a practical perspective, our results help practitioners to 'unbox' the conceptual level of DM that the current body of knowledge on DM comprises. Doing so, DM capabilities can be further exploited in practice as proper tooling requires attention from both researchers and practitioners in the coming years, similar to the development of process mining systems in the last decade.

Based on the results we formulate the following research agenda for the operationalisation of DM capabilities:

1. Analysis of the current body of knowledge on DM reveals that most contributions focus on deriving sequencing patterns and are also referred to as decision-point analyses based on event data from business processes. In general, we argue that (more) research attention should be directed towards DM from a decision point of view focused on discovering, conformance checking and improvement of derivation patterns in data captured in decision logs. This also further helps in maturing the research field of DM and provides a basis for further operationalisation towards proper software systems to support businesses in leveraging the power of their operational decision-making.
2. Although the identified contributions did not primarily focus on presenting requirements for DM systems we argue that many of them describe DM capabilities from a high level of abstraction that should be further explored in the future, e.g., by describing how to process open norms in decision logs, how outliers in decision logs can be detected and managed, or what changing or extending of decision models exactly entails. Doing so helps future research in becoming more practically applicable as the current contributions are limited in their use for practitioners. Also, future studies can focus on validating capabilities that are detailed well enough.
3. Our findings show that most contributions focus on exploring the discovery phase of DM, while the conformance checking and improvement phases seem to be defined in less detail. Future studies should secure that both the conformance checking and improvement phases are considered and explored further. This is important to close the feedback loop of decision-making using DM capabilities as only discovery only gets organisations so far. Furthermore, we see that the identified contributions primarily focus on

detailing the input data, while less attention is focused on the algorithm and especially the output (data) of the conformance checking and improvement phases. This observation also calls for further exploration and definition regarding these aspects.

4. The current body of knowledge offers a narrow glimpse into non-functional requirements specifically relevant to DM systems. From the contributions, we identified a focus on usability and reliability, but how the other software quality aspects come into play is yet to be discovered. Exploration of the non-functional requirements (also referred to as constraints by some) is important as these are contextual boundaries that should be taken into account by design as much as possible. Future studies that explore the use of DM systems in practice could benefit from these findings.
5. We argue that it is important to conceptually ground DM phases and concepts. The identified contributions focus on theoretically proving that DM based on deriving derivation patterns is a technique that should be further explored. However, the current body of knowledge seems to lack a strong empirical basis, which should be taken into account in future studies, e.g., involving subject-matter experts and experts from outside academia as well as setting up case studies in which DM phases and systems are integrally evaluated and lessons learned formulated.

**References**

Bajec, M., & Krisper, M. (2005). A methodology and tool support for managing business rules in organisations. Information Systems, 30(6), 423–443. https://doi.org/10.1016/j.is.2004.05.003

Batoulis, K., & Weske, M. (2018). Disambiguation of DMN Decision Tables. In Lecture Notes in Business Information Processing (Vol. 320, Issue March, pp. 236–249). https://doi.org/10.1007/978-3-319-93931-5_17

Berkhout, M., & Smit, K. (2022). Utilizing Algorithms for Decision Mining Discovery. 35 Th Bled EConference Digital Restructuring and Human (Re)Action, 343–358. https://doi.org/10.18690/um.fov.4.2022.21

Boyer, J., & Mili, H. (2011). Agile business rule development: Process, Architecture and JRules Examples. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-19041-4

Chalvatzis, K. J., Malekpoor, H., Mishra, N., Lettice, F., & Choudhary, S. (2019). Sustainable resource allocation for power generation: The role of big data in enabling interindustry architectural innovation. Technological Forecasting and Social Change, 144(March 2018), 381–393. https://doi.org/10.1016/j.techfore.2018.04.031

Corea, C., & Delfmann, P. (2018). A tool to monitor consistent decision-making in business process execution. CEUR Workshop Proceedings, 2196(September), 76–80.

de Jong, R., Leewis, S., & Berkhout, M. (2021). Decision Mining versus Process Mining: a Comparison of Mining Methods. 2021 5th International Conference on Software and E-Business (ICSEB), 28–32. https://doi.org/10.1145/3507485.3507490

De Smedt, J., vanden Broucke, S. K. L. M., Obregon, J., Kim, A., Jung, J.-Y., Vanthienen, J., B, J. D. S., Broucke, S. K. L. M., & Obregon, J. (2017). Decision Mining in a Broader Context: An Overview of the Current Landscape and Future Directions. In Lecture Notes in Business Information Processing (Vol. 281, pp. 197–207). Springer International Publishing. https://doi.org/10.1007/978-3-319-58457-7_15

Dijkstra, E. W. (1974). On the Role of Scientific Thought. In Selected Writings on Computing: A personal Perspective (pp. 60–66). Springer New York. https://doi.org/10.1007/978-1-4612-5695-3_12

Etikala, V., Veldhoven, Z. Van, & Vanthienen, J. (2020). Text2Dec: extracting decision dependencies from natural language text for automated DMN decision modelling. … Management Workshops: BPM …. https://doi.org/10.1007/978-3-030-66498-5_27

Gibbs, G. (2007). Thematic Coding and Categorizing. In Analyzing Qualitative Data (pp. 38–55). SAGE Publications, Ltd. https://doi.org/10.4135/9781849208574.n4

Goossens, A., Claessens, M., Parthoens, C., & Vanthienen, J. (2021). Deep Learning for the Identification of Decision Modelling Components from Text (pp. 158–171). https://doi.org/10.1007/978-3-030-91167-6_11

Goossens, A., Claessens, M., Parthoens, C., & Vanthienen, J. (2022). Extracting Decision Dependencies and Decision Logic from Text Using Deep Learning Techniques (pp. 349–361). https://doi.org/10.1007/978-3-030-94343-1_27

Goossens, A., De Smedt, J., & Vanthienen, J. (2023). Extracting Decision Model and Notation models from text using deep learning techniques. Expert Systems with Applications, 211(August 2022), 118667. https://doi.org/10.1016/j.eswa.2022.118667

Graham, I. (2006). Business rules management and service oriented architecture a pattern language (1st ed.). John Wiley & Sons.

ISO. (2014). ISO/IEC 25000:2014. https://www.iso.org/standard/64764.html

Kotonya, G., & Sommerville, I. (1998). Requirements engineering: processes and techniques. Wiley Publishing.

Leewis, S., Berkhout, M., & Smit, K. (2020). Future Challenges in Decision Mining at Governmental Institutions. AMCIS 2020 Proceedings., 6. https://doi.org/https://aisel.aisnet.org/amcis2020/adv_info_systems_research/adv_info_systems_research/6

Leewis, S., Smit, K., & Berkhout, M. (2022). Business Rules Management and Decision Mining - Filling in the Gaps. Proceedings of the 55th Hawaii International Conference on System Sciences, 6229–6238. https://doi.org/10.24251/HICSS.2022.755

Leewis, S., Smit, K., & Zoet, M. (2020). Putting Decision Mining into Context: A Literature Study. In Lecture Notes in Information Systems and Organisation (Vol. 38, Issue September, pp. 31–46). https://doi.org/10.1007/978-3-030-47355-6_3

Mircea, M., Ghilic-Micu, B., & Stoic, M. (2012). An Agile Architecture Framework that Leverages the Strengths of Business Intelligence, Decision Management and Service Orientation. In Business Intelligence - Solution for Business Development.

Ossher, H., & Tarr, P. (2001). Using multidimensional separation of concerns to (re)shape evolving software. Communications of the ACM, 44(10), 43–50. https://doi.org/10.1145/383845.383856

Rozinat, A., & van der Aalst, W. M. P. (2006). Decision Mining in ProM. In S. Dustdar, J. L. Fiadeiro, & A. P. Sheth (Eds.), Business Process Management: 4th International Conference, BPM 2006, Vienna, Austria, September 5-7, 2006. Proceedings (pp. 420–425). Springer Berlin Heidelberg. https://doi.org/10.1007/11841760_33

Rula, A., Maurino, A., & Batini, C. (2016). Data and Information Quality: Dimensions, Principles and Techniques (1st ed.). Springer. https://doi.org/10.1007/978-3-319-24106-7

Sarno, R., Sari, P. L. I., Ginardi, H., Sunaryono, D., & Mukhlash, I. (2013). Decision mining for multi choice workflow patterns. Proceeding - 2013 International Conference on Computer, Control, Informatics and Its Applications: "Recent Challenges in Computer, Control and Informatics", IC3INA 2013, 2007, 337–342. https://doi.org/10.1109/IC3INA.2013.6819197

Scheibel, B., & Rinderle-Ma, S. (2022). Decision Mining with Time Series Data Based on Automatic Feature Generation. In Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics): Vol. 13295 LNCS (Issue c, pp. 3–18). https://doi.org/10.1007/978-3-031-07472-1_1

Schlosser, S., Baghi, E., Otto, B., & Oesterle, H. (2014). Toward a Functional Reference Model for Business Rules Management. 2014 47th Hawaii International Conference on System Sciences, 3837–3846. https://doi.org/10.1109/HICSS.2014.476

Smit, K. (2018). Organization and Governance of Business Rules Management Capabilities. Open University. https://research.ou.nl/en/publications/organization-and-governance-of-business-rules-management-capabili

Smit, K., Zoet, M., & Berkhout, M. (2017). Verification capabilities for business rules management in the Dutch governmental context. 2017 International Conference on Research and Innovation in Information Systems (ICRIIS), 1–6. https://doi.org/10.1109/ICRIIS.2017.8002499

Sommerville, I., & Sawyer, P. (1997). Requirements engineering: a good practice guide. John Willey & Son Ltd.

The Open Group. (2011). TOGAF v9.1 standard. http://pubs.opengroup.org/architecture/togaf9-doc/arch/

van der Aalst, W. M. P. (2011). Process Mining. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-19345-3

Vanthienen, J. (2021). Decisions, advice and explanation: An overview and research agenda. In A Research Agenda for Knowledge Management and Analytics (pp. 149–170). https://doi.org/10.4337/9781800370623.00016