Decisions, Advice and Explanation: an Overview and Research Agenda

Jan Vanthienen

ABSTRACT

Decisions are essential in business. The discovery, modeling and specification of decision knowledge through data-based or knowledge-based AI makes business operations more agile and more intelligent. Since the introduction of the Decision Model and Notation (DMN) standard (Object Management Group, 2019), decision management and modeling have become an important research subject and are increasingly being used in industry. There is, however, a wealth of research topics to be discovered. This chapter builds upon existing research in the field of decision modeling, knowledge representation and decision tables to outline a research agenda for the management, modeling and exploitation of decision knowledge.

Keywords: decision modeling, DMN, explainable decisions, decision knowledge, knowledge-intensive business processes

1 INTRODUCTION

Operational decisions are made on a daily basis, but they require a lot of knowledge (Vanthienen, 2015). Such decisions are taken frequently and are repetitive in nature, e.g., determining which insurance rate applies to a specific customer, deciding on eligibility, configuring a product to meet a customer's demands, etc. While the impact of a single decision is typically small for the organization, their volume means that there is a significant impact on the business. Operational decisions, therefore, have to be efficient, maintainable, consistent, reproducible, compliant, reliable and explainable.

The DMN (Decision Model and Notation) standard has emerged as a way to represent the knowledge of day-to-day operational decisions in business operations. Decision modeling is already heavily used in banking, insurance, social security and standard procedures, and numerous tools incorporate DMN modeling, making the standard available for industry. Also the research community is increasingly working on decision modeling (Aa et al., 2016; Calvanese et al., 2018; Dangarska et al., 2016; Figl et al., 2018).

2 DECISION MODELING AND MANAGEMENT

2.1 THE DECISION MANAGEMENT KNOWLEDGE CYCLE

Managing decision knowledge encompasses a number of stages, as illustrated in Figure 1. Decision knowledge is acquired from data, text or human expertise and then represented and modeled in a common notation. Knowledge quality is obtained through verification and validation. Based on the body of knowledge, decisions are then embodied in executables or services and knowledge-driven applications are built to support decision making. Finally, the execution trail will allow acquiring and refining new knowledge about the decisions in the application domain.

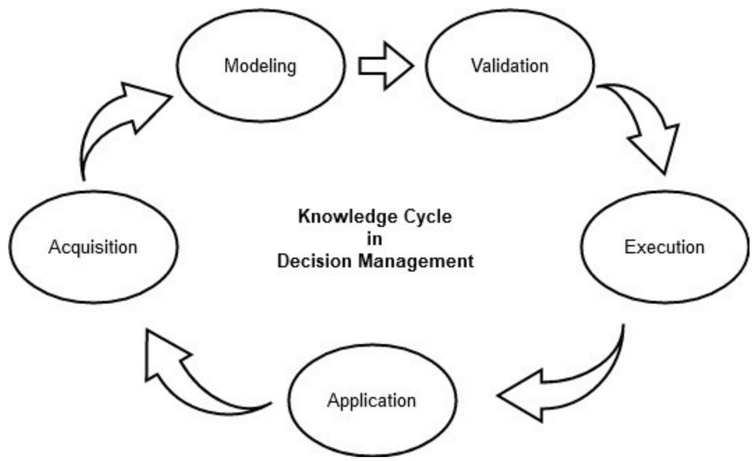

Figure 1. The knowledge cycle in decision management

2.2 DECISION MODEL AND NOTATION (DMN)

To address the need of a decision modeling standard, DMN (Decision Model and Notation) was introduced by OMG, the Object Management Group. The primary goal of DMN is to provide a common notation for all business users (from business analysts, to technical developers, and finally to business people), and to bridge the gap between business decision design and implementation. By explicitly identifying decisions and dependencies and by describing the decision logic, the decision can be managed separately from the process itself, thereby increasing the business agility of an organization.

DMN provides distinct, but related constructs for decision modeling: the decision requirements diagram, the decision logic, and the corresponding expression language, called FEEL – the Friendly Enough Expression Language. A brief overview of the most important elements is provided in Figure 2. DMN is designed to model decisions inside or outside the context of a business process model.

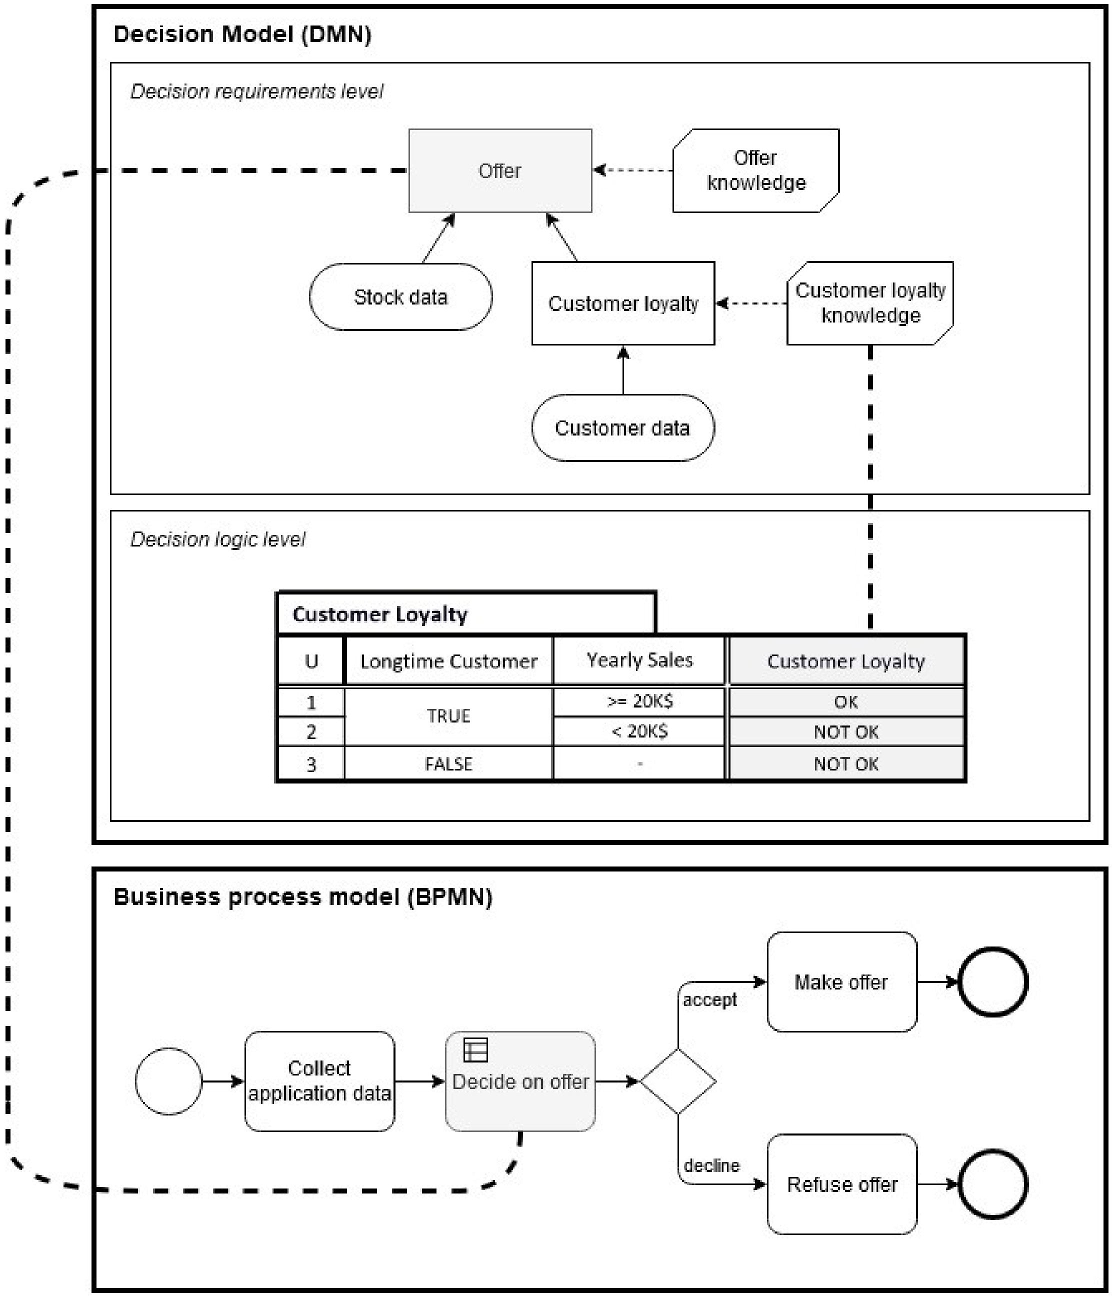

Figure 2. *Important modeling concepts in DMN*

2.2.1 The decision requirements level

A Decision Requirements Diagram (DRD) is used to portray the domain of decision making at a high level of abstraction, the decision requirements level, with only a few types of constructs: the decisions, input data, knowledge models and knowledge sources, together with the interdependencies, called requirements.

Rectangles are used to depict decisions, corner-cut rectangles for business knowledge models, and ovals to represent input data. In Figure 2, *Offer* and *Customer Loyalty* are decisions. They determine a value, based on input data and decision logic. *Offer* is the top decision. Its outcome is used in the business process model. Input data for the decisions are *Stock data* and *Customer data*. The *Offer* and *Customer Loyalty* decisions are made using knowledge, as indicated by the *Offer knowledge* and *Customer Loyalty knowledge* corner-cut rectangles.

The arrows represent requirements: solid arrows for information requirements and dashed arrows for knowledge requirements. An information requirement indicates that a decision needs the value of input data or the outcome of another decision. The *Offer* decision is dependent on *Stock data* and on the outcome of the *Customer Loyalty* decision, which in its turn is dependent on *Customer data.* A knowledge requirement indicates that a decision needs knowledge (e.g. in the form of rules) in order to determine an outcome. The *Customer Loyalty* decision requires *Customer Loyalty knowledge* to decide about the outcome. A third requirement (authority requirement) is not shown in Figure 2. An authority requirement indicates who or what is the source of the decision knowledge.

2.2.2 The decision logic level

The decision logic level specifies the underlying decision logic for each decision, very often in the form of decision tables (see Figure 2). Decision logic indicates what the decision outcome should be for specific combinations of the values of input information items. Decision tables traditionally visualize these rules with input-outcome combinations in a tabular format that is easy to use for business, guarantees completeness and consistency and offers straightforward automation (Huysmans et al., 2011).

The decision logic level also provides an expression language (called FEEL) for specifying detailed decision logic, by defining complex expressions, composed from simpler expressions. Moreover, this level offers a notation (boxed expressions) which allows these expressions to be associated with elements in the decision requirements level.

The two levels together specify a complete decision model, understandable by the business and detailed enough for automation.

3 A RESEARCH AGENDA FOR DECISION KNOWLEDGE ACQUISITION AND MODELING

Decision modeling with DMN finds its origin in decision tables, where rules for decision logic are represented in a structure of related tables. Each decision table maps combinations of input values to outcomes. Decision tables and the accompanying methodology have proven a powerful vehicle for acquiring the decision knowledge and for checking completeness, correctness, and consistency (Codasyl, 1982). DMN builds upon these concepts and standardizes decision table formats in use, standardizes the relations between decisions in a decision requirements diagram, and introduces a standard expression language (FEEL).

Based on earlier research and new developments, this chapter provides a set of guidelines and research topics applicable to the full trajectory of decision modeling and management (see Figure 3).

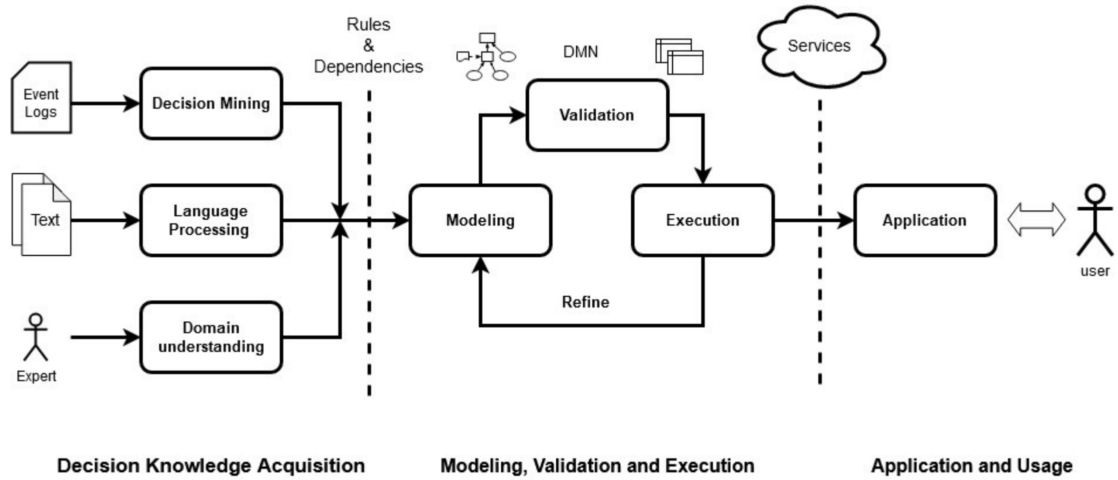

Figure 3. From decision knowledge to application

3.1 SINGLE DECISION TABLE MODELING APPROACHES

Operational decisions can be modeled according to different strategies, depending on what is the starting point of the modeling process: text (often augmented with expert knowledge), or historical case data. When historical case data are available, input information items and decision rules may be derived automatically from patterns in the log. Otherwise, relevant information items and rules will have to be extracted from the text, or from discussions with the domain expert.

3.1.1 Manual modeling

Usually, the modeling process starts from an available description in the form of a text, procedure, law, …. A domain expert is often within reach to deal with questions that turn up during the modeling process. If not all relevant information items or rules are available up front, the modeler and the domain expert gradually discover relevant criteria and outcomes in a dialogue mode and refine the table until a full description of the decision logic is obtained.

The following basic modeling steps in constructing decision tables are distinguished, as already described in (Vanthienen et al., 1998):

1. Define inputs (conditions) and outcomes of the decision situation.
2. Specify the problem in terms of decision rules.
3. Fill the decision table based on the rules.
4. If necessary, check the table for completeness, correctness, and contradictions.
5. Simplify the decision table and display it.

Note that some of these steps, including verification, can be automated, as illustrated by the Prologa tool (Vanthienen & Dries, 1994). For an overview of verification and validation research, relevant for step 4, see (Antoniou et al., 1998). Since the introduction of DMN, a lot of this research has been rediscovered in the form of verification tools for DMN models and tables, see e.g. (Calvanese et al., 2018).

Table simplification (step 5) can occur in multiple ways: reducing the number of rules by rule merging or by reordering, and splitting a table. Rule merging (table contraction) implies that rules with equal outcomes and complementary values for only one input item are joined together (Laurson & Maggi, 2017; Vanthienen & Dries, 1994). The number of rules can also be reduced by reordering the input information items (in combination with table contraction), which can be used to determine the order with the minimum number of rules. Finally, one decision table can (or should) be split into smaller tables if the table contains hidden dependencies. This is called factoring or normalization (Vanthienen & Snoeck, 1993), similar to database normalization.

3.1.2 Decision mining from case data

When historical data about case attributes and their outcome are available, the decision rules can be discovered from the case data and transformed into a decision table (Baesens et al., 2003; Wets et al., 1998). This is the area of data mining or business analytics. Predictive models, based on past data, are widely used in both research and business.(Baesens et al., 2009; Gopal et al., 2011; Liebowitz, 2013)

Most research, however, focuses on improving the accuracy or precision of these models and less research has been undertaken to increase their comprehensibility to the analyst or end user. Even if comprehensibility is of a subjective nature, some representation formats are generally considered to be more easily interpretable than others and decision tables score extremely well here in terms of comprehensibility, ease of use and confidence ((Huysmans et al., 2011; Martens et al., 2007).

One interesting form of analytics is business process mining, the discovery, monitoring and improvement of business process knowledge from event logs that are readily available in modern information systems, e.g. audit trails, electronic patient records, or the transaction logs of an enterprise resource planning system (Van Der Aalst, 2011). Process mining can be used to discover models describing processes, to monitor deviations, and to check and improve the performance of business processes.

As decisions are an important aspect of process models, it is clear that mining decisions is closely related to process mining. Mining decisions is not only about discovering the decision logic at a certain decision point in a process model. A decision is more than decision logic, it can be an entire decision model (see the later section on Full decision mining). Moreover, because DMN allows to separate processes and decisions, according to the separation of concerns principle, the integrated mining of decisions and processes offers very promising research topics (see the later section on Integrated mining of decisions and processes).

3.1.3 Decision mining from text

In many business cases, decision modeling starts from a law, text, manual, policy document or procedure. If a domain expert is not immediately available to support the modeling process, the text may be the only available source. Automatic extraction of models from text has been researched in many other modeling standards, such as process models, data models, rule models, etc. See e.g. (Friedrich et al., 2011) for the extraction of process models from text. Mining decision rules from text, using text mining, and transforming these rules into a decision table is a promising research direction.

3.2 FULL DECISION MODELING APPROACHES

A full decision model consists of the two levels: the decision requirements level and the decision logic level. Approaches towards building decision models, therefore, will have to construct elements at both levels, showing both the dependencies between decisions and the logic of each decision.

3.2.1 Manual decision modeling strategies

When a business analyst or a domain expert builds a decision model from a problem specification (usually a text), multiple starting points are possible: one can start from the general structure (and build the DRD first), or one can start from the detailed decision logic of each decision and work upwards towards a top decision. Mixed forms are, of course, also possible and very common. And while building the requirements diagram, it is always an option to immediately specify the corresponding logic for a decision, or postpone the detailed logic until the dependencies are completely specified. In reality, a mixture of all these approaches will be used. These strategies are similar to well established modeling approaches in the BPM community, i.e., bottom-up, top-down, and combined modeling approaches, adapted to suit DMN modeling.

3.2.2 Full decision mining from case data

Case data can be a wealthy source for discovering decision rules. This can be in the form of a DMN decision table, but a more complex challenge is the mining of an entire decision model from (event and) data logs, including dependencies between decisions, based on the data relations between them (Bazhenova & Weske, 2015; Smedt, Broucke, et al., 2017). Usually, however, this is in combination with process discovery from event logs (therefore see the later section on integrated mining of decisions and processes).

3.2.3 Full decision extraction from process flows

When a business process model is available, a decision model can also be extracted from the process model, based on split gateways.

In these approaches, the decision points in a process model are identified and the decision logic containing the data dependencies is derived from the process model, see (Aa et al., 2016; Batoulis et al., 2015; Bazhenova & Weske, 2015). The result is a decision model including the decision requirements diagram and decision logic. The process model is adapted accordingly, where the decision logic is now in the decision model, and not hidden in the process model.

3.2.4 Full decision model mining from text

Mining decision rules from text, using text mining, and transforming these rules into a decision table is one thing. It is even more challenging to mine dependencies between decisions, and other elements of the requirements level. This is a promising research direction that is just being explored.

3.3 ADDITIONAL GUIDELINES FOR DMN DECISION MODELS

Although DMN is mainly about notation, and is not meant to include a design methodology, there is a long history of decision modeling guidelines, offering guidance to structure decisions into separate tables, in order to build sound decision tables using a stepwise methodology and to avoid table anomalies. These guidelines deal with the form as well as with the contents of the decision tables.

Structure and content

1. *Basic structure:* Decision tables represent rules about related input information items and outcomes. All input information items in one rule are implicitly connected with AND.
2. *Completeness, consistency, consistency over multiple decisions* are important properties for maintainability, comprehensibility and correctness. The question of overlapping rules is a key issue in dealing with consistency and correctness.
3. *Multiple outcome items:* Decision tables can have multiple outcomes. If the purpose of the table is to assign an outcome to a (sub)decision, the main action will assign that outcome, e.g. true/false, classification results, values. There may also be additional outcomes, depending on the purpose of the table.

Form, conciseness and readability

4. *Table contraction:* Proper rule minimization enhances readability (still avoiding overlapping rules).
5. *Input order optimization:* a different overall input order may produce a smaller table because of contraction.

Normalization

6. *Normalization:* Decision tables can (or should) be split up if the outcomes are not dependent on all the input information items.

4 RESEARCH ISSUES IN DECISION MODEL VERIFICATION

4.1 VALIDATION AND VERIFICATION OF DECISION MODELS

Verification and validation of knowledge-based systems (including decision tables) has been a major area of research, e.g. in the EUROVAV series of conferences (European Conference on Verification & Validation of Knowledge-based systems) (Antoniou et al., 1998; Coenen et al., 2000). This research deals with typical rule anomalies, such as: redundancy (including duplicates and subsumption), ambivalence, circularity, and deficiency (missing rules) (Preece & Shinghal, 1994; Vanthienen et al., 1998). Numerous algorithms and tools are available for checking and eliminating rule anomalies for all possible values of the input variables.

4.1.1 Verification of single decision tables

Verification of decision tables mainly deals with completeness and consistency of the rules in a decision table.

- *Consistency of the rules:* The problem of consistency (including redundant, duplicate or subsumed rules), is closely related to the presence of overlapping rules. If the rules of a decision table are not mutually exclusive, there is at least one combination of input values that matches two rules. The outcome of the two rules can be compatible (called a multiple hit table), contradictory or equal. When the outcome is contradictory, a solution for the inconsistency has

to be provided (in DMN this is called the hit policy). Even when the rules produce the same outcome, the table is more difficult to maintain and validate manually. Because decision tables are relations, this is simply the requirement of normalization (Vanthienen & Snoeck, 1993).
Consistency can be obtained in three ways: (i) either by design, (ii) by signaling and repairing inconsistent rules, or (iii) by providing a policy that resolves overlapping and inconsistent combinations for the entire table (e.g. the first hit convention that gives priority to the first rule that matches the input data). Although all these approaches may finally produce a consistent table, especially the latter is known to be extremely complex and error prone.

- *Completeness:* Completeness implies that no combinations of input values are missing. It can be obtained in three ways: (i) either by design, (ii) by looking for missing combinations after the table is constructed, or (iii) by providing a remainder column which catches all missing rules. The latter solution, although complete by definition and even compact, is less elegant and much more difficult to understand, validate, and maintain.

Decision table methodology has shown that completeness and consistency is very important for comprehensibility and correctness of the decision model. Overlapping rules therefore, are considered harmful and reduce the power of the decision model. (Vanthienen et al., 1998).

4.1.2 Verification of single DMN decision tables

As DMN has included the decision table concept as one of the major decision logic components, all decision table research directly translates to DMN decision table research.
Actually, a lot of verification and validation research on DMN decision tables has inadvertently rediscovered decision table validation and verification research, see e.g. (Batoulis & Weske, 2018; Calvanese et al., 2018; Corea et al., 2019; Hinkelmann, 2016; Laurson & Maggi, 2017; Montali & Teinemaa, 2016; Ochoa & González-Rojas, 2017).

4.1.3 Verification of decision table networks or DMN requirements diagrams

DMN consists of the requirements level and the decision logic level. Inputs of the decision tables at the decision logic level are represented at the requirements level as information requirements. Outcomes of the decision table constitute information requirements to higher level decisions or form outcomes to the top decision. So the requirements level shows a visual representation of the relations between decision tables, and could be derived from them.
Whenever an information item in a decision table A is the outcome of another decision table B, obviously every possible value of the information item should be a possible outcome of table B. and every possible value of the information item in A should be the outcome of a rule in B. The opposite is not necessarily true: B can produce more outcomes than what is used in A if table B is reused somewhere else.

This type of verification refers to inter-tabular anomalies, anomalies that could arise due to the interaction between different tables. They are basically similar to the possible anomalies that could occur within one table: unfirable rules, missing rules, unusable outcomes, etc., but now between tables. See (Vanthienen et al., 1997) for an overview of inter-tabular verification and a toolset dealing with these anomalies.

4.1.3.1 Syntactic verification

Because the requirements level corresponds to the relations between decision tables, it basically contains no more information than what is present in the decision tables, if only tables are used. But it still useful to only model the information requirements if not all rules in the decision tables are fully specified yet. Obviously, when information requirements are modeled manually (not derived from the tables), there should be a full match between an information requirement in the DRD and an information item in the table. Most tools will ensure this automatically.

4.1.3.2 Verification over rule chains

While the previous verification of missing rules, missing information items or unusable outcomes is rather straightforward, because it is only based on the static description of the tables, things become more complicated when dependencies between information items are present.

When an input information item in one table is repeated in another table, e.g., some part of the decision logic in a certain decision may become unreachable or inconsistent for specific input values. Checking this consistency and completeness between interconnected decision tables, i.e. over rule chains, is a much more challenging problem than static verification or verification of single tables. See (Vanthienen et al., 1997) for a solution for inter-tabular verification.

5 RESEARCH ON DECISION AND PROCESS INTEGRATION

Business process management (BPM) and decision management (DM) improve the efficiency and effectiveness of organizations. While business processes are modeled in a structured and executable way, there is little attention, however, for the decision and knowledge aspect in business processes. Moreover, complex decisions are often modeled as processes, e.g. using cascaded gateways.

Decision management introduced an approach for modeling decisions independently (Batoulis et al., 2015, 2017; Biard et al., 2015; Mertens et al., 2017; Song et al., 2019a, 2019b; Taylor et al., 2013; Vanthienen & Caron, 2013) and aims at the separation of decision knowledge from business processes, thereby simplifying process modeling. This separation of concerns is crucial for the modeling and maintainability of both processes and decisions, but it raises the question how both approaches can be combined, both in modeling and mining (Janssens et al., 2016).

5.1 INTEGRATED MODELING OF DECISIONS AND PROCESSES

Decisions could be considered as local, not related to other elements of the process. A decision model is then a further refinement of a decision activity in a process model and multiple decisions in a process lead to isolated decision models.

But that is not the full intent of DMN. Decision models can contain multiple related decisions and top decisions in a single decision model. Related decisions have elements in common (e.g. decision logic, input data), and therefore belong in the same model, but are still different decisions at different places in the process model. These decisions, however, may extend over process modeling elements, produce intermediate events or data, or require a specific ordering in the process model, so the decision model is not completely isolated.

It is, therefore, important to apply an integrated approach for decision and process modeling and to ensure consistent integration between both models (Janssens et al., 2016). Potential inconsistencies are, e.g., unused decision outcomes, missing intermediate process actions, unnecessary decision activities, unsound ordering of decision activities, missing input data, etc. Consistent integration ensures the correct separation between decision and process models according to integration principles (Hasić et al., 2018).

5.2 INTEGRATED MINING OF DECISIONS AND PROCESSES

Decision mining in processes as introduced in (Rozinat & van der Aalst, 2006) is able to build predictive models that explain why certain paths are followed at fixed decision points in a process. This approach is control flow-driven and can be called decision point analysis. Additionally, and since the introduction of DMN, interesting new approaches have introduced the discovery of DMN models from process data (Batoulis et al., 2015; Bazhenova et al., 2016). The emphasis, however, is still on explaining the control flow, or the techniques at least incorporate control flow constructs in the models.

In accordance with the separation of concerns principle, control-flow agnostic techniques have been proposed for the integrated mining of both a process and a decision model based on extensive decision-process logs (Smedt, Broucke, et al., 2017; Smedt et al., 2019; Smedt, Hasic, et al., 2017). In this approach, mining decisions is independent from, but consistent with the control flow, which produces an integrated, but separated view of the decisions and process.

6 A RESEARCH AGENDA FOR DECISION MODEL EXECUTION AND USAGE

When properly specified, and now that appropriate tooling is available, decision models are executable. This means that, if the values for input information items can all be obtained, a straightforward execution will determine the outcome of the decision, inside or even outside a business process. This is the major application of decision modeling and DMN nowadays. Numerous business applications can be found in insurance, finance, healthcare, rules, laws and regulations, etc. See e.g. (Hasić & Vanthienen, 2020) for an income taxation case.

But there is more than straightforward input to output execution.

6.1 INCOMPLETE DATA

A decision model captures relevant decision knowledge, and current tools use this knowledge in one way: given all relevant input, what is the outcome of the decision.

In real-world applications, other functionalities are interesting and should be possible: Reasoning with missing data, e.g., could already provide useful consequences based on the data that is available. This would allow answering questions like: are certain decision outcomes still possible, given incomplete information? Or, which missing input information would be relevant in order to determine the outcome of a decision? The decision knowledge is already present in the model, but more powerful reasoning engines will be necessary. Current research in this area shows some very promising directions (Dasseville et al., 2016).

6.2 OPTIMAL EXECUTION

In a number of cases, attention could be paid to execution efficiency or more flexible forms of code generation. By generating least-cost execution trees dealing with condition test times and case frequencies, the average execution time of a decision can be minimized, by transforming decision tables into optimal test sequences. See e.g. (Codasyl, 1982) for an overview of optimization algorithms.

6.3 EXPLANATION

Explainability is becoming a hot topic in AI. When decisions are made by intelligent systems and algorithms, trust is of utmost importance. One of the major reasons to trust a model or system is the ability to understand and explain in detail the underlying knowledge. The ability to explain is not only desirable, it is often required by regulators for accountability reasons. Black box decisions will not offer this advantage. Explainability is also important because it allows evaluating and improving the decisions, correcting unwanted effects and including missing decision logic.

When it comes to explainability of the decisions taken, DMN offers a number of advantages: separation of concerns, modular structure and a comprehensible representation of the decision logic. Decision tables have a proven record in ease of use, completeness and consistency (Huysmans et al., 2011).

6.4 DECISION ANALYSIS, SIMULATION, ADVICE AND OPTIMIZATION

DMN offers a business-friendly, but still limited representation of business decision knowledge. The advantage is that it can be directly built and maintained by business experts, but for more sophisticated applications, the expressive power and reasoning mechanisms will have to be extended. The challenge here is to preserve the ease-of-use for domain experts, and extend the functionality by linking it to knowledge representation and reasoning achievements, optimization techniques, constraint satisfaction methods, etc. Interesting developments in this area can be found in (Dasseville et al., 2016; Deryck et al., 2018; Feldman, 2016; Paschke & Könnecke, 2016).

Consider for instance the application domain of eligibility for loans in a bank. That is the decision knowledge. Ideally, this knowledge should be comprehensible to the business experts, well-organized, explainable (e.g. for legal reasons) and multi-purpose for different types of applications or questions. Typical questions might be:

- *Decision*: Is this person, given all relevant data, eligible for a loan?
- *Explanation:* Why can this person not get a loan?
- *Incomplete inputs*: Given what we know already, what is the maximum loan amount this person might get?
- *Simulation:* What would be the result if the values of a few information items change?
- *Advice:* What are important information items to get a loan?
- *Goal seeking*: What would have to change for this person to be eligible for a loan?
- *Optimization*: What are the parameter values for this person that maximize the loanable amount.

The knowledge remains the same, but the questions are different. Answering these questions requires powerful knowledge representation and reasoning techniques, see e.g. (De Cat et al., 2018). On the other hand, it is important that business domain experts are still able to formulate, understand and validate the relevant knowledge.

When decision knowledge is represented in a standard and comprehensible way, other advantages appear. It now becomes possible to analyze the knowledge using advanced, but generic business intelligence techniques, answering questions such as:

- *Fairness*: Does the knowledge correspond to what one would expect in terms of changes in information item values?
- *Compliance*: Is the decision knowledge compliant with existing rules and regulations?
- *Decision monitoring*: How many cases actually obtained a certain decision outcome?
- *Policy evaluation*: Given the number of historical cases, do we have to change the decision rules?
- *Simulation and prediction*: What would be the aggregated outcome of a policy change?
- *Policy optimization*: What can we do to increase certain decision outcomes?

7 CONCLUSION

The introduction of the Decision Model and Notation (DMN) standard triggered decision management and modeling as important research subjects. There is a wealth of research topics to be discovered. for the management, modeling and exploitation of decision knowledge.

DMN offers a business-friendly representation of business decision knowledge. The advantage is that it can be directly built and maintained by business experts, but for more sophisticated applications, the expressive power and reasoning mechanisms can still be extended. The challenge here is to preserve the ease-of-use for domain experts, and extend the functionality by linking it to knowledge representation and reasoning achievements.

REFERENCES

Aa, H. van der, Leopold, H., Batoulis, K., Weske, M., & Reijers, H. A. (2016). Integrated Process and Decision Modeling for Data-Driven Processes. In M. Reichert & H. A. Reijers (Eds.), *Business Process Management Workshops 2015* (Vol. 256, pp. 405–417). Springer.

Antoniou, G., Harmelen, F. van, Plant, R., & Vanthienen, J. (1998). Verification and Validation of Knowledge-Based Systems: Report on Two 1997 Events. *AI Magazine*, *19*(3), 123–126.

Baesens, B., Mues, C., Martens, D., & Vanthienen, J. (2009). 50 years of data mining and OR: upcoming trends and challenges. *JORS*, *60*(S1). https://doi.org/10.1057/jors.2008.171

Baesens, B., Setiono, R., Mues, C., & Vanthienen, J. (2003). Using Neural Network Rule Extraction and Decision Tables for Credit—Risk Evaluation. *Management Science*, *49*(3), 312–329. https://doi.org/10.1287/mnsc.49.3.312.12739

Batoulis, K., Haarmann, S., & Weske, M. (2017). Various Notions of Soundness for Decision-Aware Business Processes. *International Conference on Conceptual Modeling*, 403–418.

Batoulis, K., Meyer, A., Bazhenova, E., Decker, G., & Weske, M. (2015). Extracting decision logic from process models. *International Conference on Advanced Information Systems Engineering*, 349–366.

Batoulis, K., & Weske, M. (2018). Disambiguation of DMN Decision Tables. In W. Abramowicz & A. Paschke (Eds.), *Business Information Systems* (pp. 236–249). Springer International Publishing.

Bazhenova, E., Buelow, S., & Weske, M. (2016). Discovering decision models from event logs. *International Conference on Business Information Systems*, 237–251.

Bazhenova, E., & Weske, M. (2015). Deriving Decision Models from Process Models by Enhanced Decision Mining. *Business Process Management Workshops*, *256*, 444–457.

Biard, T., Le Mauff, A., Bigand, M., & Bourey, J.-P. (2015). Separation of decision modeling from business process modeling using new “Decision Model and Notation”(DMN) for automating operational decision-making. *Working Conference on Virtual Enterprises*, 489–496.

Calvanese, D., Dumas, M., Laurson, Ü., Maggi, F. M., Montali, M., & Teinemaa, I. (2018). Semantics, analysis and simplification of DMN decision tables. *Information Systems*, *78*, 112–125.

Codasyl. (1982). A modern appraisal of decision tables. *Report of the Decision Table Task Group*.

Coenen, F., Bench-Capon, T. J. M., Boswell, R., Dibie-Barthélemy, J., Eaglestone, B., Gerrits, R., Grégoire, E., Ligeza, A., Laita, L. M., Owoc, M. L., Sellini, F., Spreeuwenberg, S., Vanthienen, J., Vermesan, A. I., & Wiratunga, N. (2000). Validation and verification of knowledge-based systems: Report on EUROVAV99. *Knowledge Eng. Review*, *15*(2), 187–196.

Corea, C., Blatt, J., & Delfmann, P. (2019). A Tool for Decision Logic Verification in DMN Decision Tables. *Proceedings of the Dissertation Award, Doctoral Consortium, and Demonstration Track at the 17th International Conference on Business Process Management, BPM 2019*, *2420*, 169–173. http://ceur-ws.org/Vol-2420/papeDT11.pdf

Dangarska, Z., Figl, K., & Mendling, J. (2016). An Explorative Analysis of the Notational Characteristics of the Decision Model and Notation (DMN). *Enterprise Distributed Object Computing Workshop (EDOCW), 2016 IEEE 20th International*, 1–9.

Dasseville, I., Janssens, L., Janssens, G., Vanthienen, J., & Denecker, M. (2016). Combining DMN and the knowledge base paradigm for flexible decision enactment. *CEUR Workshop Proceedings*, *1620*.

De Cat, B., Bogaerts, B., Bruynooghe, M., Janssens, G., & Denecker, M. (2018). Predicate logic as a modeling language: The IDP system. In *Declarative Logic Programming: Theory, Systems, and Applications* (pp. 279–323).

Deryck, M., Hasić, F., Vanthienen, J., & Vennekens, J. (2018). A case-based inquiry into the decision model and notation (DMN) and the knowledge base (KB) paradigm. *International Joint Conference on Rules and Reasoning*, 248–263.

Feldman, J. (2016). What-If Analyzer for DMN-based Decision Models. *RuleML (Supplement)*.

Figl, K., Mendling, J., Tokdemir, G., & Vanthienen, J. (2018). What we know and what we do not know about DMN. *Enterprise Modelling and Information Systems Architectures (EMISAJ)*, *13*, 2–1.

Friedrich, F., Mendling, J., & Puhlmann, F. (2011). Process model generation from natural language text. *International Conference on Advanced Information Systems Engineering*, 482–496.

Gopal, R. D., Marsden, J. R., & Vanthienen, J. (2011). Information mining—Reflections on recent advancements and the road ahead in data, text, and media mining. *Decision Support Systems*, *51*(4), 727–731. https://doi.org/10.1016/j.dss.2011.01.008

Hasić, F., De Smedt, J., & Vanthienen, J. (2018). Augmenting processes with decision intelligence: Principles for integrated modelling. *Decision Support Systems*, *107*, 1–12.

Hasić, F., & Vanthienen, J. (2020). From decision knowledge to e-government expert systems: The case of income taxation for foreign artists in Belgium. *Knowl. Inf. Syst.*, *62*(5), 2011–2028. https://doi.org/10.1007/s10115-019-01416-4

Hinkelmann, K. (2016). Business Process Flexibility and Decision-Aware Modeling—The Knowledge Work Designer. In D. Karagiannis, H. C. Mayr, & J. Mylopoulos (Eds.), *Domain-Specific Conceptual Modeling, Concepts, Methods and Tools* (pp. 397–414). Springer. https://doi.org/10.1007/978-3-319-39417-6

Huysmans, J., Dejaeger, K., Mues, C., Vanthienen, J., & Baesens, B. (2011). An empirical evaluation of the comprehensibility of decision table, tree and rule based predictive models. *Decision Support Systems*, *51*(1), 141–154.

Janssens, L., Bazhenova, E., Smedt, J. D., Vanthienen, J., & Denecker, M. (2016). Consistent Integration of Decision (DMN) and Process (BPMN) Models. In S. España, M. Ivanovic, & M. Savic (Eds.), *Proceedings of the CAiSE'16 Forum, at the 28th International Conference on Advanced Information Systems Engineering (CAiSE 2016), Ljubljana, Slovenia, June 13-17, 2016* (Vol. 1612, pp. 121–128). CEUR-WS.org. http://ceur-ws.org/Vol-1612/paper16.pdf

Laurson, Ü., & Maggi, F. M. (2017). A Tool for the Analysis of DMN Decision Tables. In L. Azevedo & C. Cabanillas (Eds.), *Proceedings of the BPM Demo Track 2016 Co-located with the 14th International Conference on Business Process Management (BPM 2016), Rio de Janeiro, Brazil, September 21, 2016.* (Vol. 1789, pp. 56–60). CEUR-WS.org. http://ceur-ws.org/Vol-1789

Liebowitz, J. (2013). *Big data and business analytics*. CRC press.

Martens, D., Baesens, B., Gestel, T. V., & Vanthienen, J. (2007). Comprehensible credit scoring models using rule extraction from support vector machines. *European Journal of Operational Research*, *183*(3), 1466–1476. https://doi.org/10.1016/j.ejor.2006.04.051

Mertens, S., Gailly, F., & Poels, G. (2017). Towards a decision-aware declarative process modeling language for knowledge-intensive processes. *Expert Systems with Applications*, *87*, 316–334.

Montali, M., & Teinemaa, I. (2016). Semantics and Analysis of DMN Decision Tables. *Business Process Management: 14th International Conference, BPM 2016, Rio de Janeiro, Brazil, September 18-22, 2016. Proceedings*, *9850*, 217.

Object Management Group. (2019). Decision Model and Notation, Version 1.3. https://www.omg.org/spec/DMN/1.3/PDF

Ochoa, L., & González-Rojas, O. (2017). Analysis and Re-configuration of Decision Logic in Adaptive and Data-Intensive Processes (Short Paper). *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems,"* 306–313.

Paschke, A., & Könnecke, S. (2016). RuleML-DMN Translator. *RuleML (Supplement)*.

Preece, A. D., & Shinghal, R. (1994). Foundation and application of knowledge base verification. *International Journal of Intelligent Systems*, *9*(8), 683–701.

Rozinat, A., & van der Aalst, W. M. (2006). Decision mining in ProM. *International Conference on Business Process Management*, 420–425.

Smedt, J. D., Broucke, S. K. L. M. vanden, Obregon, J., Kim, A., Jung, J.-Y., & Vanthienen, J. (2017). Decision Mining in a Broader Context: An Overview of the Current Landscape and Future Directions. In M. Dumas & M. Fantinato (Eds.), *Business Process Management Workshops—BPM 2016 International Workshops, Rio de Janeiro, Brazil, September 19, 2016, Revised Papers* (Vol. 281, pp. 197–207). https://doi.org/10.1007/978-3-319-58457-7

Smedt, J. D., Hasic, F., Broucke, S. K. L. M. vanden, & Vanthienen, J. (2019). Holistic discovery of decision models from process execution data. *Knowl. Based Syst.*, *183*. https://doi.org/10.1016/j.knosys.2019.104866

Smedt, J. D., Hasic, F., Broucke, S. K. L. M. vanden, & Vanthienen, J. (2017). Towards a Holistic Discovery of Decisions in Process-Aware Information Systems. In J. Carmona, G. Engels, & A. Kumar (Eds.), *Business Process Management—15th International Conference, BPM 2017, Barcelona, Spain, September 10-15, 2017, Proceedings* (Vol. 10445, pp. 183–199). Springer. https://doi.org/10.1007/978-3-319-65000-5

Song, R., Vanthienen, J., Cui, W., Wang, Y., & Huang, L. (2019a). A DMN-Based Method for Context-Aware Business Process Modeling Towards Process Variability. *Business Information Systems - 22nd International Conference, BIS 2019, Seville, Spain, June 26-28, 2019, Proceedings, Part I*, *353*, 176–188. https://doi.org/10.1007/978-3-030-20485-3_14

Song, R., Vanthienen, J., Cui, W., Wang, Y., & Huang, L. (2019b). Context-Aware BPM Using IoT-Integrated Context Ontologies and IoT-Enhanced Decision Models. In J. Becker & D. Novikov (Eds.), *21st IEEE Conference on Business Informatics, CBI 2019, Moscow, Russia, July 15-17, 2019, Volume 1—Research Papers* (pp. 541–550). IEEE. https://doi.org/10.1109/CBI.2019.00069

Taylor, J., Fish, A., Vanthienen, J., & Vincent, P. (2013). *Emerging standards in decision modeling*.

Van Der Aalst, W. (2011). *Process mining: Discovery, conformance and enhancement of business processes* (Vol. 2). Springer.

Vanthienen, J. (2015). On Smart Data, Decisions and Processes. In A. L. N. Fred, J. L. G. Dietz, D. Aveiro, K. Liu, & J. Filipe (Eds.), *KDIR 2015—Proceedings of the International Conference on Knowledge Discovery and Information Retrieval, part of the 7th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2015), Volume 1, Lisbon, Portugal, November 12-14, 2015* (p. 5). SciTePress. http://ieeexplore.ieee.org/document/7526887/

Vanthienen, J., & Caron, F. (2013). Modeling Business Decisions and Processes—Which Comes First? *KDIR/KMIS 2013 - Proceedings of the International Conference on Knowledge Discovery and Information Retrieval and the International Conference on Knowledge Management and Information Sharing, Vilamoura, Algarve, Portugal, 19 - 22 September, 2013*, 451–456. https://doi.org/10.5220/0004623904510456

Vanthienen, J., & Dries, E. (1994). Illustration of a Decision Table Tool for Specifying and Implementing Knowledge Based Systems. *International Journal on Artificial Intelligence Tools*, *3*(2), 267–288. https://doi.org/10.1142/S0218213094000133

Vanthienen, J., Mues, C., & Aerts, A. (1998). An Illustration of Verification and Validation in the Modelling Phase of KBS Development. *Data Knowl. Eng.*, *27*(3), 337–352. https://doi.org/10.1016/S0169-023X(98)80003-7

Vanthienen, J., Mues, C., & Wets, G. (1997). Inter-Tabular Verification in an Interactive Environment. In J. Vanthienen & F. van Harmelen (Eds.), *Proceedings of the Fourth European Symposium on the Validation and Verification of Knowledge-Based Systems, EUROVAV'97, June 26-28, 1997, Katholieke Universiteit Leuven, Leuven, Belgium* (pp. 155–165). Katholieke Universiteit Leuven, Belgium.

Vanthienen, J., & Snoeck, M. (1993). Knowledge factoring using normalisation theory. *International Symposium on the Management of Industrial and Corporate Knowledge (ISMICK'93), October*, 27–28.

Wets, G., Vanthienen, J., & Timmermans, H. J. P. (1998). Modelling Decision Tables from Data. *Research and Development in Knowledge Discovery and Data Mining, Second Pacific-Asia Conference, PAKDD-98, Melbourne, Australia, April 15-17, 1998, Proceedings*, 412–413. https://doi.org/10.1007/3-540-64383-4_48