



Designing a data mining process for the financial services domain

Veronika Plotnikova, Marlon Dumas, Alexander Nolte & Fredrik Milani

To cite this article: Veronika Plotnikova, Marlon Dumas, Alexander Nolte & Fredrik Milani (2022): Designing a data mining process for the financial services domain, Journal of Business Analytics, DOI: [10.1080/2573234X.2022.2088412](https://doi.org/10.1080/2573234X.2022.2088412)

To link to this article: <https://doi.org/10.1080/2573234X.2022.2088412>



© 2022 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 07 Jul 2022.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

Designing a data mining process for the financial services domain

Veronika Plotnikova^a, Marlon Dumas^a, Alexander Nolte^{a,b} and Fredrik Milani^a

^aInstitute of Computer Science, University of Tartu, Tartu, Estonia; ^bInstitute for Software Research, Carnegie Mellon University, Pittsburgh, PA, USA

ABSTRACT

The implementation of data mining projects in complex organisations requires well-defined processes. Standard data mining processes, such as CRISP-DM, have gained broad adoption over the past two decades. However, numerous studies demonstrated that organisations often do not apply CRISP-DM and related processes as-is, but rather adapt them to address industry-specific requirements. Accordingly, a number of sector-specific adaptations of standard data mining processes have been proposed. So far, however, no such adaptation has been suggested for the financial services sector. This paper addresses the gap by designing and evaluating a Financial Industry Process for Data Mining (FIN-DM). FIN-DM adapts and extends CRISP-DM to address regulatory compliance, governance, and risk management requirements inherent in the financial sector, and to embed quality assurance as an integral part of the data mining project life-cycle. The framework has been iteratively designed and validated with data mining and IT experts in a financial services organisation.

ARTICLE HISTORY

Received 13 June 2021
Accepted 1 June 2022

KEYWORDS

Data mining; financial services; design science; privacy; AI ethics; quality assurance

1. Introduction

Over the past decades, data mining practices have been widely adopted among organisations seeking to maintain and to enhance their competitiveness and business value (Davenport & Harris, 2017). This trend has led a number of large organisations to manage a rich portfolio of data mining projects (Davenport & Harris, 2017). The successful development, implementation, and management of data mining projects in such organisations requires a structured and repeatable approach. Accordingly, academia and industry practitioners have proposed several guidelines and standard processes for conducting data mining projects (Mariscal et al., 2010), most notably CRISP-DM¹ – a standard process that captures a wide range of recurrent data mining tasks and deliverables structured around a project life-cycle (Marban, Mariscal et al., 2009).

CRISP-DM is industry-agnostic. Organisations that wish to use CRISP-DM often need to adapt it to meet their domain-specific requirements (Plotnikova, Dumas & Milani (2019b)). Accordingly, adaptations of CRISP-DM have been developed in the fields of healthcare (Niaksu, 2015), education (Tavares et al., 2017), industrial engineering (Solarte, 2002), software engineering (Marban et al., 2007), logistics (Rahman et al., 2011) supply chain (Xiang, 2009), and e-commerce (Hang & Fong, 2009). However, to the best of our knowledge, there is no adaptation of CRISP-DM for guiding and structuring data mining projects in the financial services industry. Yet, just like the

abovementioned industry sectors, the financial services sector has its own set of domain-specific requirements, inherent to the financial services, which relate to governance, risk, and compliance, regulatory requirements, and ethics (cf., Cao (2021), Cao (2018)). Governance, risk, and compliance demand assurance for integrity, efficiency, fairness, and adequate risk management in any financial services organisation. In turn, ethics-associated challenges include addressing privacy and ethical concerns (cf., Cao (2021), Cao (2018)).

This paper addresses this gap by proposing a data mining process for the financial industry, namely FIN-DM (Financial Industry Process for Data Mining). FIN-DM adapts and extends CRISP-DM in order to address a set of requirements specific to the financial services sector, chiefly compliance, governance, and risk management requirements. In particular, FIN-DM extends CRISP-DM with tasks and guidelines designed to ensure privacy compliance, to address regulatory and ethics risks, to fulfill financial model risk management requirements, and to embed quality assurance as an integral component of the data mining life-cycle.

The FIN-DM process has been designed, starting from a set of financial industry requirements, by applying an iterative design science approach in collaboration with data mining and IT experts in a financial services organisation. This paper outlines the research design and documents the derivation of the FIN-DM process step by step.

The rest of this paper is structured as follows: in [Section 2, Background and Related Works](#), we introduce key data mining concepts, the CRISP-DM standard data mining process, and related work. Next, in [Section 3](#), we detail the *Research Design*, followed by [Section 4 FIN-DM Development and Evaluation](#). Finally, we discuss the findings in [Section 5](#), and draw conclusions, and outline directions for future work in [Section 6](#).

2. Background and related work

In this section, we introduce key concepts associated with data mining, provide an overview of existing data mining methodologies, and situate our research in the context of related works. We conclude this section by illuminating the research gap.

Data mining is defined as a set of rules, processes, and algorithms designed to generate actionable insights, extract patterns, and identify relationships from large data sets (Morabito (2016)). Data mining incorporates automated data extraction, processing, and modelling by using a range of methods and techniques. Data mining is typically conducted following a structured methodology that specifies inputs, outputs, tasks, and instructions for their execution to achieve a project's objectives (Mariscal et al. (2010), Marban, Mariscal et al. (2009)).

The foundations for structured data mining were first proposed in a series of works (Fayyad et al. (1996a), Fayyad et al. (1996b), and Fayyad et al. (1996c)) which introduced the industry-agnostic process model Knowledge Discovery in Databases (KDD). KDD is a conceptual process model developed in the computational field that supports information (knowledge) extraction from data (Fayyad et al. (1996a)). The KDD approach to knowledge discovery includes data mining as a specific step, and with its nine main steps (exhibited in [Figure 1](#)), KDD also has the advantage of considering data storage and access, algorithm scaling,

interpretation and visualisation of results, and human computer interaction (Fayyad et al. (1996a), Fayyad et al. (1996b)).

The KDD process gained dominance across both academia and industry (Kurgan and Musilek (2006), Marban, Mariscal et al. (2009)), and it served as the basis for other methodologies and process models, which addressed its various gaps and deficiencies (Plotnikova, 2021). These extensions ranged from process restructuring to complete change in focus. For example, a number of extensions concentrated on specific aspects, such as enhancements with practical steps and tasks: (1) to support iterativeness and interaction including stakeholders' decision-making (cf., Brachman and Anand (1996) and further Gertosio and Dussauchoy (2004)), (2) to detail data discovery and processing (Cabena et al. (1997)), or (3) to adapt standard process to specialised practical settings (e.g., entire life-cycle of online customer including web-based mining, e.g., Anand and Büchner (1998), Anand et al. (1998), and Buchner et al. (1999)). In contrast, the Two Crows data mining process model has defined the steps differently and re-modelled data mining process as a consultancy framework; however, it remained closed to the original KDD model. Lastly, SEMMA (Sample, Explore, Modify, Model, and Assess), based on KDD, is a vendor-specific process model related and adapted to the functional toolset of SAS Enterprise Miner (SAS Institute, (2017)), which in turn limits its application in other environments. In terms of adoption and use, KDD-based proposals have received limited attention across academia and industry (Kurgan and Musilek (2006), Marban, Mariscal et al. (2009)), and most of them converged into the CRISP-DM methodology.

In 2000, as a response to common issues and needs, and consolidating various KDD extensions, the Cross-Industry Standard Process for Data Mining (CRISP-DM) was developed (Marban, Mariscal et al. (2009)). Like KDD, CRISP-DM aims to provide practitioners




Figure 1. KDD process, as presented in Fayyad et al. (1996a) and Fayyad et al. (1996b).

with domain-agnostic guidelines to conduct data mining projects, while its iterative approach stood as its most distinguishing feature (Mariscal et al. (2010)).

CRISP-DM consists of six phases executed in iterations, as presented in Figure 2. The first phase is business understanding, which includes defining the problem, scoping, and planning. Phase 2 (data understanding) focuses on initial data collection, quality assurance and exploration, and hypotheses' detection and formulation. Next, in Phase 3 (data preparation), the final dataset is extracted from the raw data. Next, modelling techniques are selected and applied in Phase 4. In the model evaluation phase (Phase 5), decisions are based on the assessed data mining findings. Finally, in the deployment phase, the models are put into use. CRISP-DM emphasises data mining's cyclical nature within the concrete project, when discovered knowledge triggers new business questions, and beyond – when new data mining projects benefit from a previous projects' experience. Apart from the iterative approach and recognition of dependencies between the main phases, CRISP-DM allows flexibility in its application since redundant tasks and activities can be removed (Chapman et al. (2000)). CRISP-DM is currently considered a 'de facto' standard for the data mining process (Mariscal et al. (2010)).

There are two distinct patterns of how CRISP-DM is applied (Plotnikova (2020)). The first is "as-is", where CRISP-DM is used as stipulated. The second is with "adaptations", i.e., practitioners modify CRISP-DM by introducing various changes. The past decade has seen a steady decrease in the use of CRISP-DM as prescribed ('as-is'), in favour of adapting it to cater to domain-specific needs Plotnikova (2020). CRISP-DM

is industry-agnostic and, therefore, does not necessarily fulfill specific sectors' requirements. Typically, CRISP-DM adaptations address particular issues and problems arising from specific domain settings. For instance, Niaksu (2015) extended CRISP-DM to address issues specific to healthcare, such as mining non-standard datasets, addressing data interoperability challenges, and personal data privacy constraints. Likewise, Solarte proposed adapting CRISP-DM to address data mining aspects in the industrial engineering field (Solarte (2002)). In this adaptation, project roles and stakeholders' analysis of additional data requirements are emphasised to align data mining with organisational goals. In a similar vein, Marban et al. proposed adaptations to specifically target the industrial engineering industry by introducing new tasks, steps, and deliverables (Marban et al. (2007)). Thus, most adaptations focus on customising CRISP-DM for application in specific industries such as for healthcare (Niaksu (2015)), education (Tavares et al. (2017)), industrial engineering (Solarte (2002), Huber et al. (2019)), and software engineering (Marban et al. (2007); Marban, Segovia et al. (2009)).

A similar pattern has been observed in how CRISP-DM is applied in the financial sector (Plotnikova, Dumas, and Milani (2021)), and how practitioners adapt CRISP-DM when executing data mining projects in a bank (Sommerville (2005)). According to (Plotnikova, Dumas, and Milani (2019b)), three distinct types of adaptations have been observed, depending on the adaptation degree – "Modification", "Extension", and "Integration". Studies following "Modification" pattern typically concentrate on granular adjustments at the level of sub-phases, tasks, or

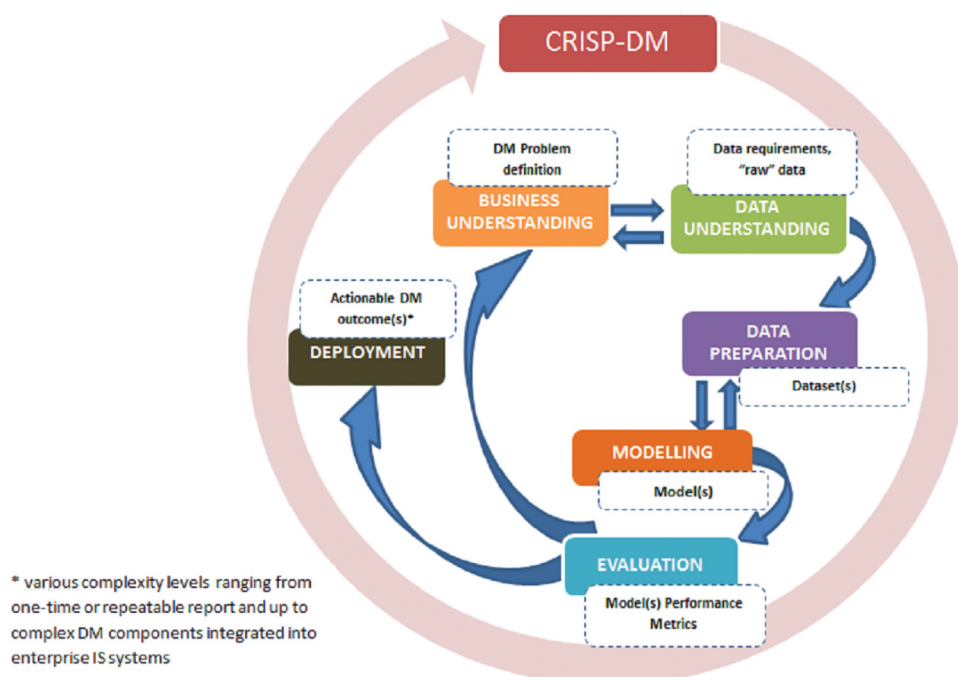


Figure 2. CRISP-DM phases and key outputs (adapted from Chapman et al. (2000))

deliverables within standard CRISP-DM phases, and they are primarily performed to address a specific use case or business problem, for example, to predict customers who likely to churn, customer credit score, or customer default (Plotnikova et al. (2019b)). This contrasts with 'Extension' pattern, where key adaptation proposals are CRISP-DM phases additions, substantial changes and extensions with the purpose to implement fully scaled, integrated data mining solutions into IS systems and business processes of organisations (Plotnikova 2021, Plotnikova et al. (2019b)), such as data mining-based solution for AML² implemented as a tool with respective IS architecture and investigative process, or knowledge-rich financial risk management process designed in addition to standard classification model and evaluation framework (Plotnikova et al. (2019b)). Hence, the key benefit achieved by 'Extension' in addition to business problem solution, is the usefulness and actionability of the data mining results (Plotnikova et al. (2019b))

Accordingly, 'Integrations' are performed at the highest abstraction level by combining CRISP-DM with methodologies and process models from other domains, such as Business Intelligence, Business Information, and Process Management (cf., Priebe and Markus (2015), Pivk et al. (2013)), with the purpose to improve data usage, business processes effectiveness, and deployment of data mining solutions in organisations (Plotnikova et al. (2019b)). 'Integrations' are also executed by adjusting to organisational (cf., Debuse (2007)) and other domain-specific aspects, for example, discrimination-awareness (cf., Berendt and Preibusch (2014)). In particular, the proposed solution includes tool support which increases correctness and usefulness of data mining results in the decision-making process, and supports monitoring, and avoidance of customers' discrimination in these decisions.

Thus, financial services organisations adapt CRISP-DM on a case-by-case basis. However, to the best of our knowledge, there is no general process model or adaptation of a standard model proposed to address the specifics of data mining in the financial services industry (Plotnikova et al. (2019b)). This paper aims to fill this gap by proposing an artefact called FIN-DM, which is a sector-specific data mining process model. For this purpose, we specify two research goals: (1) to design and develop the artefact and (2) to ascertain the artefact's utility, relevance in the application domain, and users' acceptance.

To design and develop a novel data mining process model, including sector-specific adaptations, the majority of related studies follow a common step-by-step approach (cf., Niaksu (2015), Marban et al. (2007)). Initially, based on the surrounding literature and in-depth domain reviews, improvement needs in the standard data mining process are identified, and

then solutions are proposed. Rarely, studies follow general design science principles (cf., Huber et al. (2019)), and only one domain adaptation followed a specialised design science research methodology (Tavares et al. (2017)). In this context, design science method could serve dual purpose, supporting artefact design and development in a structured manner, and providing for its evaluation with the potential users.

3. Research design

A significant part of Information Systems research is conducted within the settings of two complementary paradigms – behavioural and design science (cf., A. R. Hevner et al. (2004), A. Hevner and Chatterjee (2010)). While the behavioural paradigm is strongly orientated towards theory development, design science is an inherently problem-solving paradigm (A. Hevner and Chatterjee (2010)). For example, IS behavioural research frequently studies an IT artefact,³ predicting or explaining the artefact use, its usefulness, impacts on individuals and organisations, and similar factors based on theories (A. R. Hevner et al. (2004)). In contrast, design science research focuses on the development and evaluation of IT artefacts, usually intended to solve identified organisational problems (A. R. Hevner et al. (2004)). Furthermore, artefact utility and performance are put at the core of the evaluation. The knowledge gained in such research process highlights, for instance, how the artefact can more efficiently solve the problem (cf., J. Venable (2006), A. Hevner and Chatterjee (2010)).

There are many approaches in design science research (cf., Ostrowski and Helfert (2011), Peffers et al. (2018)); however, the related methods, research standards, and guidelines have not been sufficiently systematised and consolidated (Ostrowski and Helfert (2011)). Further, there is a lack of the common validated and widely accepted uniform methodology to carry out design science research (cf., Alturki et al. (2012)); as well, there is little practical guidance on the selection of the most suitable design science method, except for J. R. Venable et al. (2017), and we relied on this guidance in our research design. J. R. Venable et al. (2017) differentiates and compares the most common design science research methods based on a broad IS research paradigm categorisation as Objectivist/Positivist⁴ or Subjectivist/Interpretative⁵ (cf., Adam (2014)), and other key research characteristics like objective, domain, scope, etc.

To this end, J. R. Venable et al. (2017) recommends that Objectivist/Positivist-based methods are the preferred choice to Subjectivist/Interpretivist approaches in case the research goal is to achieve the best and most suitable artefact. Further, three design science research methods are classified into the Objectivist/Positivist

class, and it is recommended that (1) Systems Development Research Methodology (SDRM) is used if the artefact outcome of the research should be an IT system, (2) Design Science Research Methodology (DSRM) is applied if extensive adaptation of artefact to daily use is needed, or (3) DSR Process Model (DSRPM) is used if the goal of the research to develop design theory. As this research aims to develop an artefact, as well as to ensure its utility, applicability, and relevance in the application domain, an extensive adaptation of the artefact to daily use in practical settings will be performed. Hence, DSRM (Peffers et al. (2008)) has been selected as the most suitable research method (presented in Figure 3).

DSRM is a six-step process model allowing for iterations. It starts with *Problem Identification and Motivation*, which aims to define the research problem and significance of the solution (Peffers et al. (2008)). As a next step, *Define Objectives of a Solution*, the objectives for the artefact, either quantitative (e.g., improvements over existing solutions) or qualitative (e.g., how the artefact is expected to provide solutions to problems), are derived. Then, in the *Design and Development* step, the artefact is constructed. These can be any design objects, such as constructs, models, methods, etc. (Peffers et al. (2008)). In this step, initially, the desired functionality and architecture of the artefact are determined, and its prototype is created (Peffers et al. (2008)). Then, in a *Demonstration* step, the artefact's use and how the problem is solved are presented. Typically, a demonstration is done by means of experiment, simulation, case study, or any other applicable method (Peffers et al. (2008)). Next, the formal *Evaluation* is executed – it aims to gauge how well the artefact assists in solving the problem. Similarly to the *Demonstration*, it can take many appropriate forms, e.g., comparison of artefact functionalities to solution objectives, performance measures, etc. (Peffers et al. (2008)). Based on the

evaluation results, the artefact is either reiterated and improved, or *Communication* step follows right away, while any other potential improvements are left for future research projects (Peffers et al. (2008)).

In applying DSRM, we defined and executed four iterative cycles, as presented graphically in Figure 4.

Cycle 1 – Problem and Objectives Formulation In this cycle, a complete nominal sequential order of DSRM was followed, starting with the very first activity (A1.1) of a *Problem-centred initiation* (Peffers et al. (2008)). This is the recommended procedure when a research problem is already observed or suggested being examined in prior research, but the artefact does not yet exist Peffers et al. (2008). *Problem-centred initiation* was addressed by three separate studies – two Systematic Literature Reviews (SLR I, SLR II) and a case study. The first cross-domain SLR investigated standard data mining approaches usage across a number of industries (SLR I – Plotnikova et al. (2020)), while the second sector-specific one concentrated on the financial services industry (SLR II – Plotnikova et al. (2019b)). An industrial case study identified and reported perceived 'gaps' in the CRISP-DM process when conducting data mining projects within the actual financial services organisation (Plotnikova et al. (2021)). Then, we proceeded with data collection and processing activity (A1.2) using data originating from these works as *Input 1* (Figure 4). The data and results of all the three studies were combined and triangulated, and a consolidated catalogue of the standardised data mining process gaps was constructed (*Consolidate Outcome 1* in Figure 4).

Cycle 2 – Artefact Design and Development (Figure 4) For this cycle, we have adopted a blended approach by extending DSRM with the artefact requirement concept, and such research approach is described and motivated in detail in section 4.2 *FIN-DM Conceptualisation*. An explicit

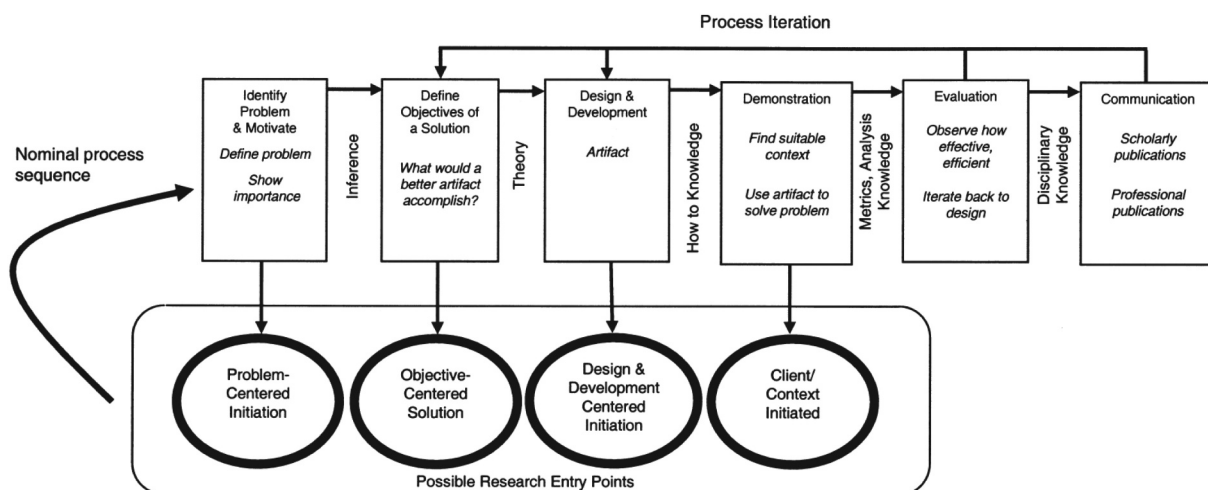


Figure 3. Design science research methodology (as in Peffers et al. (2008)).

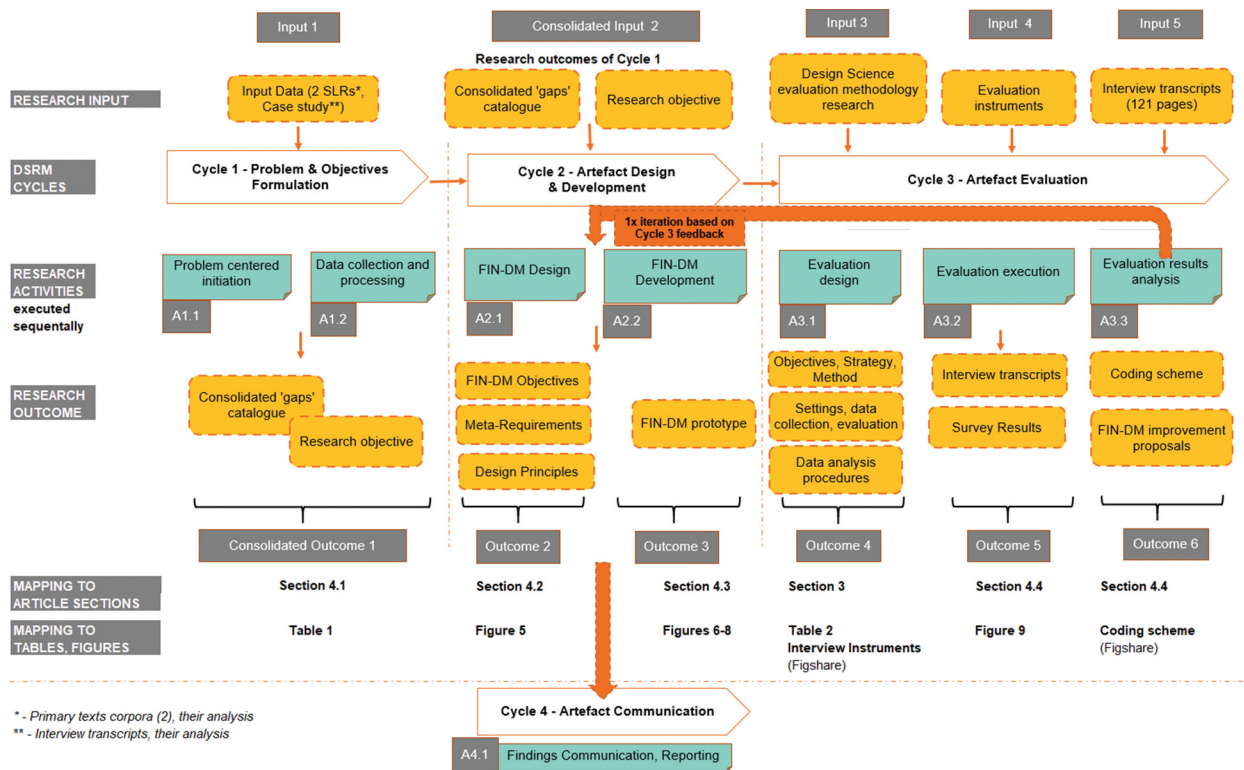


Figure 4. Research process.

elicitation of requirements and design principles was conducted as activity 2.1 (A2.1). To this end, based on the analysis of consolidated 'gaps' (*Consolidated Input 2*), we specified the *Research Outcome 2* containing three items: (1) artefact objectives, (2) artefact (meta)-requirements, and (3) design principles to address the given requirements. Then, based on *Outcome 2*, we determined the desired functionalities of the artefact and developed the first prototype as activity 2.2 (A2.2) with *Outcome 3* (see, Figure 4).

Cycle 3 – Evaluation Next, the artefact prototype was subjected to an *Evaluation cycle*, which comprised three distinct activities (see, Figure 4). First, the evaluation approach was designed and planned as activity 3.1 (A3.1). Subsequently, it was executed as activity 3.2 (A3.2) and collected data was analysed in activity 3.3 (A3.3).

In activity A3.1, the design of evaluation was detailed in *Research Outcome 4* (see, Figure 4). Artefact evaluation focused on examining it from users' perspective. Initially, a common set of evaluation criteria was derived based on the relevant artefact evaluation methods and models. The criteria set was subsequently broken down into lower-level constructs. Evaluation was conducted by the means of two methods, starting with the combined individual demonstration session followed by *semi-structured interviews*, and a *questionnaire*. An interview guide and questionnaire instruments were constructed as *Input 4* available at link⁶ towards Activity 3.2 (A3.2 *Conducting evaluation*).

As activity 3.3. (*Evaluation results analysis*), interviews were transcribed and evaluated jointly with the questionnaire results as *Input 5* to activity 3.3 (A3.3). Interviews' transcripts were coded iteratively based on the methods proposed by Saldana (2015).⁷ The coded interview responses were categorised into four distinct categories (*Minor, Medium, Major, and Critical*) based on the suggested improvements' complexity and impact. The suggested improvements were iterated back into the *Artefact design and development cycle*, which was repeated to produce its final, improved version demonstrated in **Cycle 4 – Artefact Communication** which comprises peer-reviewed publication, dissertation defence and broader sharing within research community.⁸

4. FIN-DM development

In this section, the FIN-DM design and development is presented following DSRM research cycles and activities (as in Figure 4). We start with the *Problem Formulation Cycle* – focusing on the derived catalogue of gaps. Next, we present the design cycle in *FIN-DM conceptualisation* – highlighting (meta)-requirements, design principles, and design features. This is followed by a presentation of the *FIN-DM process model*, itself, thus focusing on key components and design decisions. Subsequently, the *FIN-DM Evaluation* of potential users' is presented and discussed – highlighting key insights.

Table 1. Consolidated 'gaps' classes catalogue.

Gap name	Definition	Data source (input study)
G1 – Requirements management, & elicitation	A lack of tasks for validation and modification of existing requirements and elicitation of new ones	Case study, SLR I, ⁹ SLR II ¹⁰
G2 – Inter-dependencies	A lack of iterations between different phases of CRISP-DM	Case study
G3 – Universality	A lack of support for various analytical outcomes, unsupervised and specialised techniques, and deployment formats	Case study, SLR II
G4 – Regulatory & compliance	A lack of tasks to address regulatory compliance (in particular, GDPR)	Case study, SLR I
G5 – Validation	A lack of support for piloting models in real-life settings	Case study, SLR II
G6 – Actionability	A lack of support for piloting models in real-life settings	Case study, SLR I, SLR II
G7 – Process	Data mining process controls, quality assurance, and critical process enablers (data, code, tools, infrastructure, and organisational factors are not taken into consideration) required for the effective execution of data mining projects	Case study, SLR I

4.1. Problem formulation cycle

In this cycle, we initially consolidated the gaps inherent in the standard data mining process (chiefly, CRISP-DM) reported in input studies. Overall, there are seven gaps' classes in the gaps catalogue (as specified in Table 1). Most of the gaps were identified across at least two input studies. Approximately half of the gaps are not specific to financial services and were reported in research on other sectors, thus implying they are generic and universal by nature. We review each gap's class below.

The *Requirements management and elicitation Gap* encompasses all phases of the CRISP-DM, both in the context of financial services (Plotnikova et al. (2019b)) and in other sectors (Plotnikova et al. , (2020)). This gap was reported as the most critical. It is related to a lack of structured and explicit tasks for requirements elicitation, development, and management in the standard CRISP-DM model. To this end, a vast spectrum of requirements is concerned, including business requirements, technological aspects (tools, platforms, and implementation and deployment processes), and data mining itself – e.g., data, modelling, and evaluation requirements. This gap primarily causes (1) the risk of a mismatch between data mining outcomes and (2) additional efforts (especially, for model deployment) and project time overruns.

The *Universality Gap* has been discovered to be present in the *Modelling and Deployment phases* of the standard data mining process. This gap is related

to a lack of support for non-modelling analytical outcomes apart from traditional supervised modelling, such as interpretations, metrics, visualisations, and the like. Another related aspect is a lack of support for 'multi-modelling' (applying such techniques in combination), and for various deployment formats and their associated technical requirements. This gap results in the risk of a mismatch between the delivered data mining outcomes and business needs.

The *Validation Gap* has been identified in the *Evaluation phase*. This gap relates to a lack of validation in data mining outcomes in real business settings, causing a risk of (1) poor model(s) performance when used in actual business activities and processes and (2) a mismatch of the data mining solution to real business needs.

The *Actionability Gap* has been identified as missing tasks for discovering application scenarios and business settings in which data mining project outcome would be used. This gap, while narrow in scope, is most likely to produce a disproportionately high negative impact by inhibiting the data mining project's business value realisation.

The *Regulatory and Compliance Gap* has been broadly identified as a lack of support for privacy and regulatory compliance in the CRISP-DM. Especially, this gap impacts the *Data Understanding and Preparation phases* and project *Deployment* due to the risk of non-compliant processing of private customer data as stipulated by the GDPR.¹¹ Though GDPR is a general regulation, its highest impact is on industries collecting and processing customer data. Hence, financial services are among the most affected sectors.

Finally, the identified *Process Gaps* refer to a lack in CRISP-DM of three key components required for the effective execution of data mining projects. In particular, there is a lack of quality assurance (i.e., lack of internal controls and quality assurance in data mining development), governance mechanisms, and capabilities. To this end, capabilities refer to a lack of critical data mining process enablers including data, code, tools, infrastructure, and organisational factors. As well, the given aspects lack in the context of a stand-alone data mining lifecycle. They are also considered as prerequisites to enable the continuous, industrialised execution of data mining projects at scale.

On a higher abstraction level, CRISP-DM to some extent, misses *agility* via a lack of iterativeness and recognition of some dependencies. As well, it is perceived as not *complete*, in regard to content, to cover regulatory and practitioners' needs. Lastly, the *organisational perspective* is not completely addressed, as the standard data mining process model does not support repeatable and reproducible data mining. Hence, from a design science perspective, *Improvement research* is needed to create a *new solution for a known problem* within a known application context (Gregor and

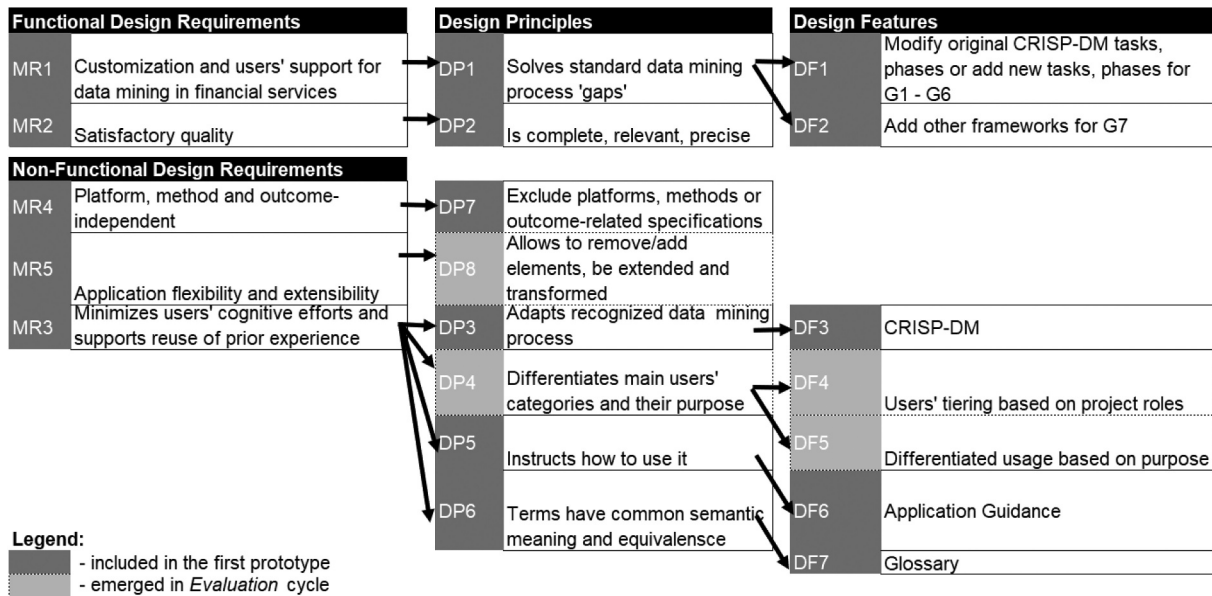


Figure 5. FIN-DM design requirements and translation into design principles and features.

Hevner (2013)). This would tackle the discovered gaps in the standard data mining process. Thus, the *Research objective* of FIN-DM is to provide practical solutions and mechanisms to mitigate or eliminate the identified gaps.

4.2. FIN-DM conceptualisation

Based on insights from the *Problem Formulation* and guided by related work practice, we have derived requirements for FIN-DM. Then, they have been translated into design principles and features, hence satisfying the requirements (as presented in Figure 5). To this end, requirements (as in Meth et al. (2015)) relate to generic requirements that any artefact instantiated from this design should meet. They define *what* needs to be satisfied. In turn, design principles are generic capabilities through which requirements are addressed (Meth et al. (2015)), and they define *how* (meta)-requirements are to be satisfied. Lastly, as defined by Meth et al. (2015), design features are specific ways to implement a design principle in an actual artefact, they are the key attributes of the artefact.

The common requirements' classification into functional and non-functional was used, cf., Bourque and Fairley (2014), Wieggers and Beatty (2013). In the context of the research, functional requirements define the content of FIN-DM – i.e., content attributes which enable users to execute data mining projects in the financial sector. These attributes make FIN-DM useful and relevant in this application domain. In turn, the non-functional requirements define FIN-DM from a user acceptance perspective, with importance placed on usability, flexibility, etc. (Wieggers and Beatty (2013)). Design principles are defined following the

pattern suggested by Chandra et al. (2015). Accordingly, it contains three types of information: (1) actions possible through the artefact, (2) the artefact's properties which facilitate the given actions, and (3) conditions under which such design will work (if applicable; Chandra et al. (2015)).

To satisfy the *Research objective* of FIN-DM, the first meta-requirement (MR1) and design principle (DP1) are as follows:

MR1. *FIN-DM provides customization and user support for data mining execution in financial services – DP1. FIN-DM solves standard data mining process' 'gaps' – enabling users to address specifics of data mining in the financial domain effectively, given that users have general data mining knowledge and experience*

Next, given that FIN-DM is a user-oriented solution, it should address the identified 'gaps' and satisfy key quality characteristics (MR2). In particular, FIN-DM should be precise and understandable by users. Also, it should be complete in each element so that it is useful and serves the purpose well.

MR2. *The quality of FIN-DM and its elements is satisfactory – DP2. FIN-DM and its elements are complete, relevant, and precise, enabling users to improve data mining projects structuring, execution experience, and delivery time, given users have basic knowledge on applying structured approaches.*

FIN-DM is designed for practitioners; therefore, it should be easy to use, promoting a wider adoption and usage of FIN-DM. In the context of data mining projects, "ease of use" could imply (1) minimising potential user's cognitive efforts to get acquainted with FIN-DM, (2) supporting the reuse of prior knowledge of the data mining execution experience when applying FIN-DM. To satisfy such requirements, we should not oblige the

user to learn a conceptually new process model, but rather to minimise familiarisation efforts. Also, users need to be supported to reproduce the existing experience. To this end, FIN-DM should be based on a recognised standard data mining process. Hence:

MR3. FIN-DM minimizes users' cognitive efforts to familiarize and apply FIN-DM by supporting the reuse of existing data mining process experience and knowledge. – **DP3.** FIN-DM adapts a recognized standard process, thus enabling users to apply FIN-DM effortlessly without specialized, prerequisite training, given that users have general knowledge about most common data mining approaches (e.g., CRISP-DM, KDD)

Further, to minimise users' efforts:

DP4. FIN-DM differentiates the roles of users, thus enabling users to apply FIN-DM effectively depending on their roles in data mining projects. This is under the assumption that the data mining project adheres to most common roles

DP5. FIN-DM includes instructions, thus enabling users to apply FIN-DM effortlessly. This is under the assumption that users have general knowledge about most common data mining approaches (e.g., CRISP-DM)

DP6. FIN-DM terms have common semantic meaning and equivalence for users, thus enabling users to interpret FIN-DM's elements correctly and consistently. This is under the assumption that users have basic knowledge of most common data mining and IT concepts

Next, to extensively support users and financial services organisations in data mining, FIN-DM possesses a number of flexibility and adaptability characteristics across broad ranges of contexts.

Firstly, while adapted towards financial services specifics, FIN-DM supports and guides potential users across a great variety of data mining projects and execution capabilities (flexibility in technical execution). Therefore, FIN-DM is independent of:

- environments – allows application in projects executed across a broad range of data mining environments, platforms, and tools,
- methods – is applicable to projects utilising a broad range of data mining and modelling methods,
- outcomes – allows for a broad range of data mining and modelling outcomes.

Accordingly, this results in the following meta-requirement:

MR4. FIN-DM is platform, method, and outcome-independent – **DP7.** FIN-DM and its elements exclude environment, methods, or outcomes-related specifications, thus enabling users to apply FIN-DM in any environment, with any methods, and for a broad range of data mining purposes

Secondly, to achieve complete flexibility in the project and organisational execution, FIN-DM needs to address two other flexibility dimensions, in particular: it needs (1) to be easily adaptable to the specifics of the particular project or organisation (context adaptability), (2) to be extensible, so that practitioners and researchers can easily integrate new aspects into the existing artefact (general adaptability).

Thirdly, to achieve flexibility on a higher abstraction level, FIN-DM needs to possess dynamic adaptation mechanisms to stay relevant and adapt to internal and external environmental changes (general agility).

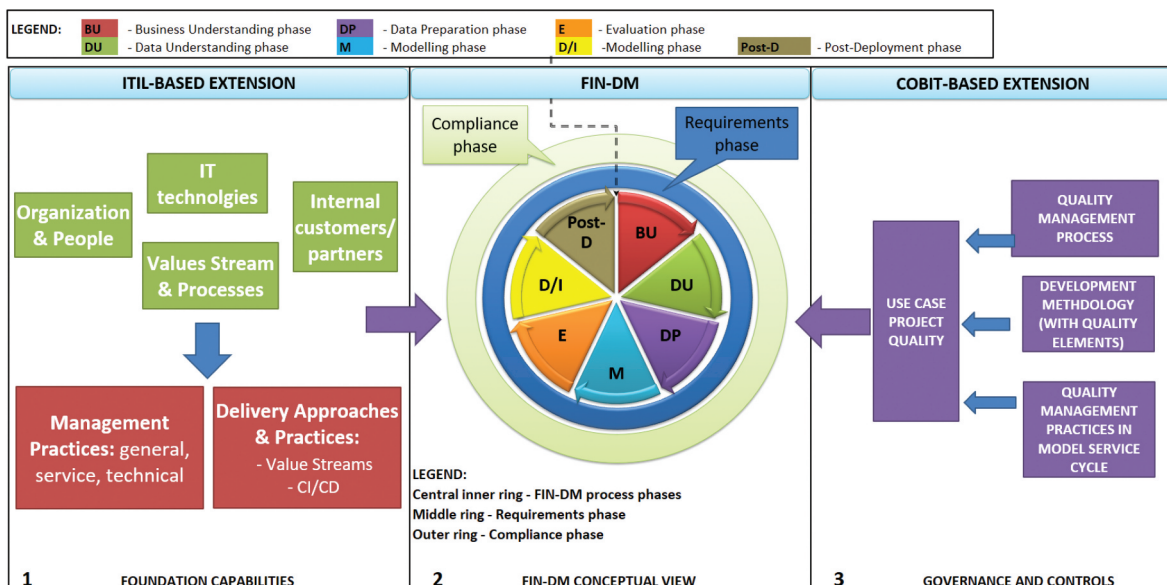


Figure 6. FIN-DM conceptual representation.

Consequently, the following meta-requirement is formulated:

MR5. *FIN-DM is flexible in application, extensible, and adaptable to internal and external environments changes – DP8.* *FIN-DM allows users to remove or add elements, accommodate extensions, make changes, and be permanently transformed. Thus, users are able to adapt FIN-DM without significant effort, given that users define the scope and purposes of such transformations*

We further specified the design features to satisfy design principles (as in Figure 5). They are reviewed in the following section in conjunction with the design decisions.

4.3. The FIN-DM process model

In this section, FIN-DM is presented and discussed, starting with a conceptual overview and a review of individual components. The solutions to 'gaps' and the respective design decisions are presented and discussed in detail.

FIN-DM consists of five items structured into two distinct parts: (1) two representations of the proposed process, in the form of a Conceptual View and a Hierarchical Process View, and accompanying checklists, as well as (2) supplementary materials – the Application Guidance and the Glossary.¹²

4.3.1. FIN-DM conceptual and hierarchical views

The FIN-DM Conceptual Representation consists of three components (visualised left to right in Figure 6): (1) the foundation capabilities, (2) the data mining process (FIN-DM), and (3) governance and controls.

The FIN-DM process model is the central component; it is a hierarchical process model. The FIN-DM process consists of three cyclic subprocesses visualised as rings – the inner, the middle, and the outer rings. The inner ring contains seven *phases*, adapted and extended from CRISP-DM. Acknowledging numerous iterations in between any phases, the process is represented as fully recurring. The middle ring is the *Requirements phase*, while the outer ring is the *Compliance phase* – they complement the data mining process (inner ring) end-to-end. Each phase consists of several generic *tasks* with respective *outputs* (see, Figures 7,8 below). By design, both *phases and generic tasks* have two characteristics, namely, independence (DP7) and flexibility (DP8). As noted, FIN-DM, itself, is based on a recognised standard data mining process, CRISP-DM (DP3, DF3); hence, FIN-DM retains key CRISP-DM terminology and structures (in the inner ring).

The *Foundation capabilities* element is positioned as a specialised *framework* extension of FIN-DM – it is relevant and applicable to the whole FIN-DM life-cycle. Therefore, it is placed on the same hierarchical level as the process model itself. The *Foundation capabilities* are based on frameworks or elements originating from ITSM,¹³ namely ITIL. The *Governance and Controls* element is a specialised *framework* extension of FIN-DM based on COBIT.¹⁴ These two higher-level supplementary domains provide guidance towards FIN-DM process elements and their execution (as visualised by two input arrows in Figure 6).

To satisfy FIN-DM's key research objective – to solve the standard data mining process' gaps – the proposed solutions (DP1, DP2) are integrated top-down at three levels with decreasing abstraction – *frameworks* and *phases* (both specified in *Conceptual View* in Figure 6), followed by *generic tasks* (specified in *Hierarchical View* in Figures 7, 8). FIN-DM's *Hierarchical Process Model* view (as visualised in Figures 7,8 below) remains generic, but at a lower abstraction level compared to the *Conceptual View* (central part in Figure 6). Here, the data mining process *phases and tasks* are detailed, including new elements in the form of *phases and tasks* introduced as solutions to the 'gaps' (DP1, DP2).

4.3.2. FIN-DM solutions to gaps

This section reviews FIN-DM (based on a detailed Hierarchical View in Figures 7, 8), focusing on the gaps solutions. We review them in a 'bottom-up' manner, starting with the new *tasks, phases*, and finally *domain extensions*. In developing solutions to gaps, we have followed a structured, five-step approach. Initially, we have reviewed the financial service industry's general regulatory and governance best practices, recommendations, and requirements from a domain-specific perspective. Then, from a data mining process perspective, we drew an analogy between IT delivery projects, software development practices, and data mining projects. We reviewed if any solutions were reported in the context of these domains. Next, we formulated a set of constructs or aspects for the respective gap and conceptualised the relationships between them and the characteristics of the data mining process. In this way, conceptualisations were derived for the *Compliance and AI Ethics phase*, which are presented in *Concepts universe* in Figure 8. As a final step, we addressed each of the required attributes or characteristics via the actions-tasks presented in *FIN-DM tasks and phases universe* in Figures 7, 8. Each action-task is tagged with the respective attribute and characteristics it tackles.

Addressing the Universality Gap (G3) We have mitigated this gap by (1) including an explicit requirements elicitation and development for analytical by-products and their associated deployments in separate *Requirements phases* (reviewed below) and (2)

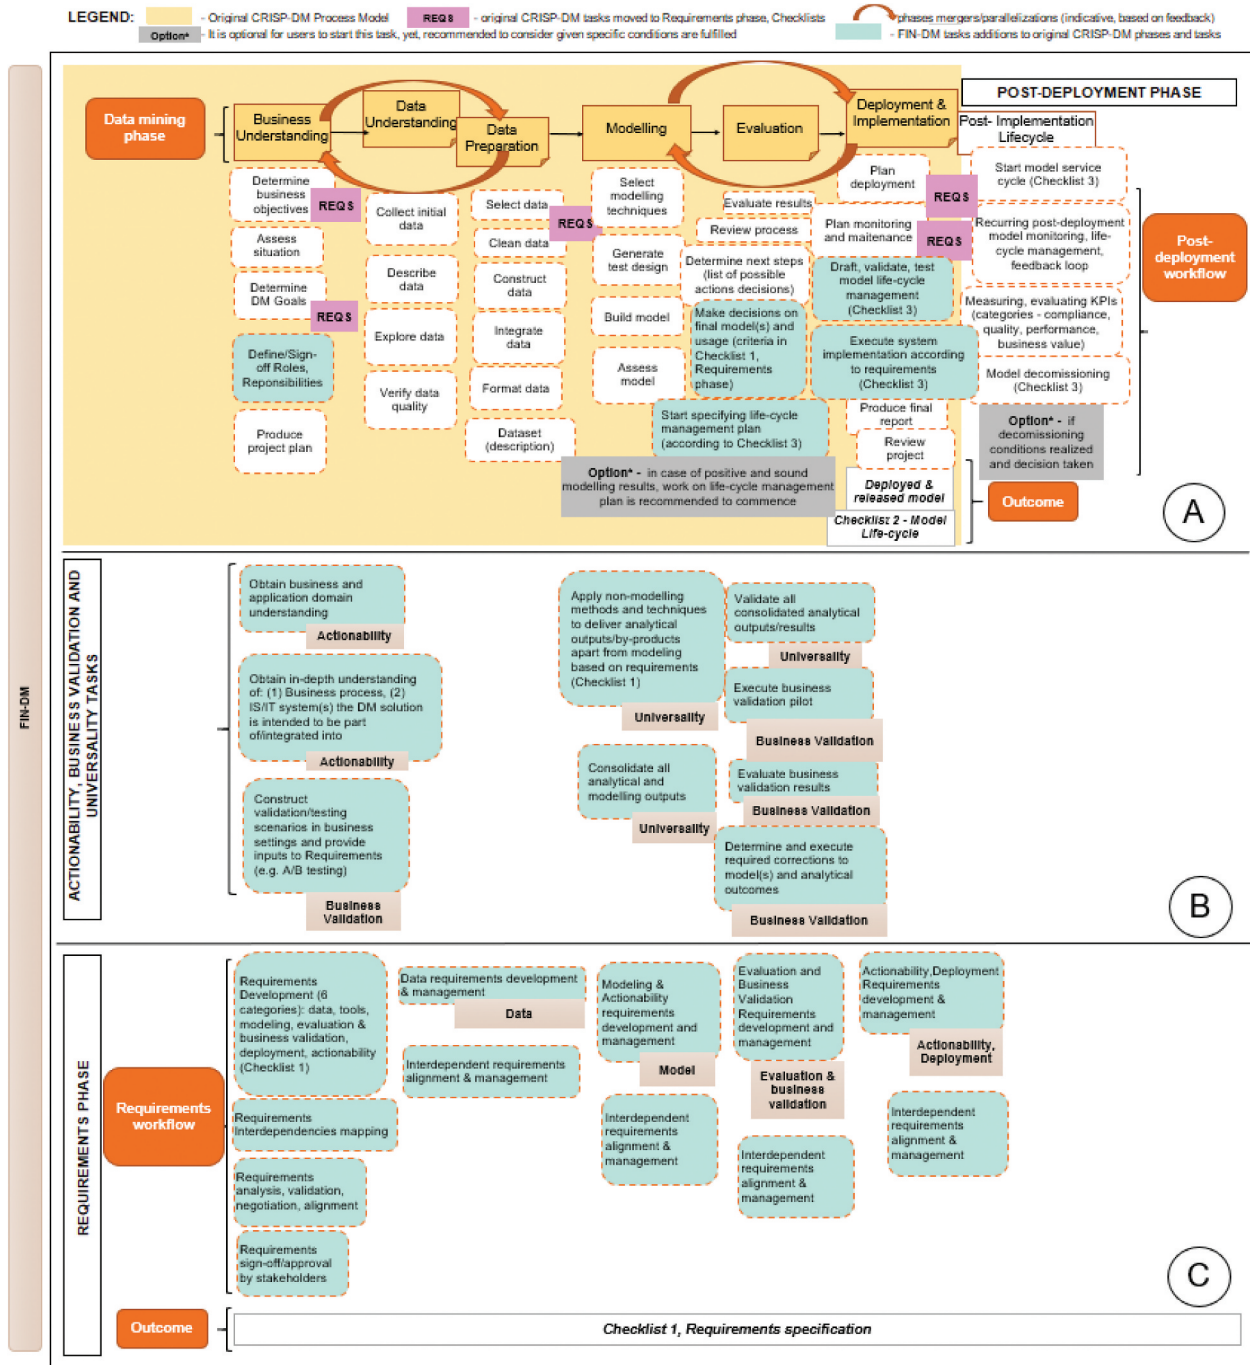


Figure 7. FIN-DM hierarchical view – additional tasks, requirements and post-deployment phases.

augmenting the original CRISP-DM phases with respective tasks of *Applying non-modelling methods and techniques*, *Consolidating all analytical and modelling outputs*, and *Validating all outputs with stakeholders*. These additional tasks have been integrated into FIN-DM’s *Modelling* and *Validation* phases, respectively (section B in Figure 7).

Addressing Validation Gap (G5) We have addressed this gap by embedding explicit *Business validation tasks* (section B in Figure 7). In particular, in the *Business Understanding* phase, we formulated *Construct validation/testing scenarios in*

business settings, for example, A/B testing as a common approach. Next, the respective scenarios are implemented in the *Evaluation* phase by *Executing and Evaluating business validation pilot*. Importantly, as a feedback loop, validation results are fed into the model’s final design through the *Determine and execute required corrections* task.

Addressing Requirements (management, elicitation (G1), Inter-dependencies (G2), Actionability Gap (G6)) In proposing a solution to fundamental gap of Requirements management, elicitation, we have considered various existing requirements engineering methods

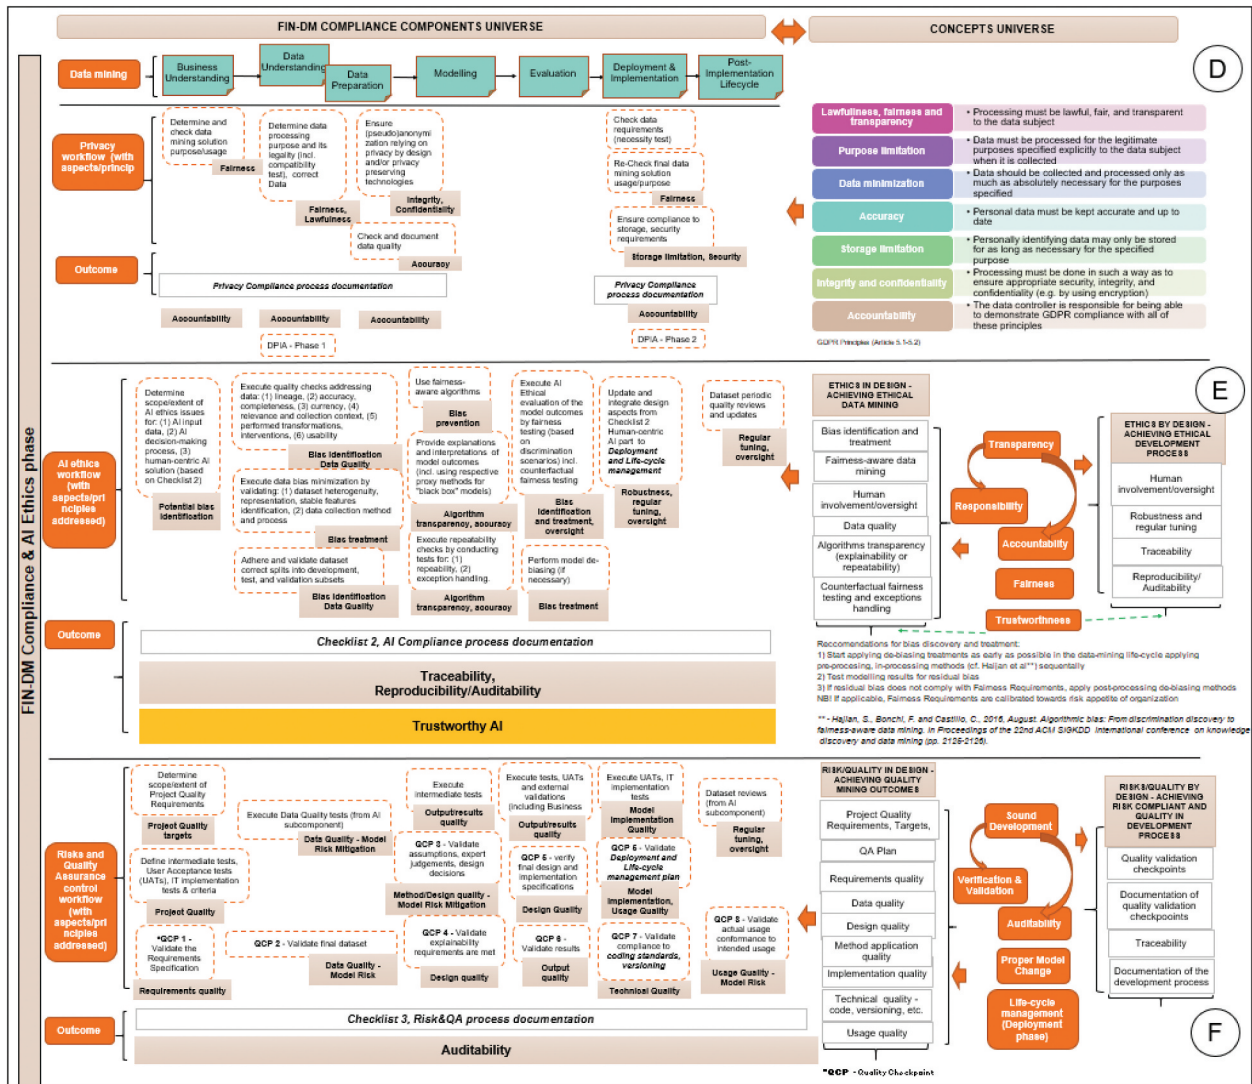


Figure 8. FIN-DM hierarchical view – compliance phase.

and approaches, including the well-known viewpoint-oriented approaches for elicitation and analysis, formal mathematical methods, and goal-oriented approaches. While adopting the method’s perspective could be beneficial, it might not support satisfying *Independence* (MR4, DP7) and *Application flexibility and extensibility*, (MR5, DP8) meta-requirements, and design principles. Furthermore, specific method implementation might require the involvement of potential users in mastering the particular requirements engineering method, which contradicts *Minimisation of users’ cognitive efforts* (MR3) meta-requirement of FIN-DM. Therefore, we opted to incorporate *Requirement activities* which are agnostic and well-recognised.

Hence, as the solution to the gap, we have proposed a new *Requirements phase* (section C in Figure 7) to augment the FIN-DM process and run in parallel to all process phases. Based on Sommerville (2005), in the *Business Understanding phase*, key requirements tasks include (1) *Requirements Development*, (2) *Requirements analysis and validation*, (3) *Requirements inter-dependencies mapping*, and (4) *Requirements*

negotiation and sign-off/approval with stakeholders. Throughout the rest of the data mining life-cycle, there are two *Requirement tasks* undertaken in each phase, namely – *Requirements development and management*, which is focused on the respective phase’s scope (data, modelling, evaluation, deployment, respectively), and *Requirements alignment*. Aligning to CRISP-DM’s standard process, the *Requirements phase and tasks* have a specific output, which is the *Requirements Specification*. As well, this phase has a supporting Checklist 1 to assist in documenting and managing requirements from the inception of data mining life-cycle up until its end. The checklist is structured based on a high-level area of a data mining project, to which respective questions/points are mapped. Overall, we have identified six key areas – data, tools, modelling, actionability, evaluation and business validation, and deployment.

Requirements phase running end-to-end of data mining life-cycle is an integral part of the overall solution to Interdependencies Gap. It is addressed by increased iterations and recommendations for phase mergers and

parallelisation (section A in Figure 7), which also provides for improved control and management of interdependencies in data mining life-cycle.

We also propose to address the *Actionability* problem as part of the *Requirements phase*, eliciting *Business usage requirements* which cover four interrelated items: (1) preliminary business usage scenario(s), (2) how the solution would be embedded/integrated into the business process, (3) what the prerequisite business process changes are, and (4) their potential implementation plan. This is reinforced by two additional tasks (in Section 1.A of Figure 7) aimed at obtaining an in-depth understanding of the business domain and business process, as well as how the respective data mining solutions will be used. Finally, as a means to address *Actionability*, we propose a separate *Implementation/Post-Deployment phase*, which explicitly formulates model deployment associated tasks and a data mining solution life-cycle management plan and activities (both not present in the original CRISP-DM model).

Addressing Regulatory and Compliance Gap (G4) As noted, this gap, in the context of financial services, is primarily associated with privacy regulation, GDPR in particular. At the same time, GDPR is not the only legislation applicable in the context of financial services, which traditionally have been heavily supervised and regulated. Therefore, we enhanced the regulatory and compliance scope of FIN-DM to tackle other core regulations applicable beyond privacy (fulfilment of MR1, MR2 and DP1, DP2), such as *AI ethics* and *Risk Management*.

In particular, many emerging policy initiatives and regulatory frameworks concern *AI ethics* and the associated risks inherent in the data mining process. Similar to privacy legislation, AI ethics guidelines are general and impact businesses and sectors developing and using AI-based solutions and products, including the financial services sector. FIN-DM also tackles financial sector-specific regulations. In particular, following the aftermaths of the 2007–2008 financial crisis, key regulatory requirements, implemented over the last decade, focused on the systemic stability and management of risks in the financial industry (Kashyap and Iveroth (2021)). These regulations concern enhancing risk governance in financial institutions, which in the context of data mining relates to governing and managing *model risk*¹⁵ emerging when data mining models are developed and used. Risk governance requirements have been addressed and implemented in the financial services sector via the generally accepted 3LoD (three Lines of Defence model) organisational model. Accordingly, data mining experts in financial services organisations belong to the first line of defence; they have to manage and control model risks as part of the model development process. Drawing on an analogy to the IT domain, a software developer is an example of a front-

line expert in the first line of defence, and managing risks in a software development life-cycle as part of daily operations.

To tackle the compliance gap and to address *privacy, AI ethics, and risk management* in data mining, we introduced the *Compliance phase* (visualised in Figure 8), consisting of three distinct sub-components for each area (areas D, E, F, respectively, in Figure 8). Similarly to the *Requirements phase*, this new phase covers the data mining life-cycle end-to-end.

In designing the *Privacy sub-component* (section D in Figure 8), we were guided by the conceptualised relationship between GDPR principles¹⁶ and data mining privacy risks (see *Concepts Universe* in section D of Figure 8). To this end, GDPR has been selected as the most notable and expansive privacy regulation currently in force. The regulation applies to the processing of personal data of individuals located in the European Economic Area (EEA) and is therefore applicable to any enterprise doing such processing inside the EEA. Also, GDPR core principles have been and are continued to be adopted in the new privacy regulations worldwide. Concerning risks, there are two key risks¹⁷ of transgressing privacy regulations in data mining. The first one concerns illegitimate personal data processing while developing the model or executing analytical tasks (impact on lawfulness and fairness). The second risk relates to the deployment and usage of data mining models in direct marketing, automated decision-making, and profiling without appropriate legal grounds (affecting lawfulness). In earlier works, these risks have been addressed with Privacy-Preserving Data Mining (PPDM; cf., Mendes and Vilela (2017)), which suggests specific methods and techniques for safeguarding privacy of data collection (lawfulness), data mining output (accuracy), data distribution, and data publishing (integrity, confidentiality, and security).

Combining both perspectives, in the *Privacy sub-component* (see, Figure 8), we initially propose to determine the purpose of the data mining solution and personal data processing. By assessing legitimacy, it is possible to prevent the risk of illegal data processing by updating and removing respective data requirements or, alternatively, ensuring legal grounds. As a next step, data requirements are checked for adherence to Data Minimisation principle.¹⁸ In the Data Preparation phase, explicit tasks for (pseudo) anonymisation are proposed, either by applying respective techniques or through privacy-preserving technologies (ensure privacy by design). That step is complemented by data quality checks to comply with accuracy principles. The legality of the final data mining results or model usage for the intended purpose is checked, along with security and storage requirements. The Accountability principle is

followed throughout the data mining life-cycle by documenting the privacy compliance process and outcomes.

Addressing issues represented by common *AI ethics* constructs is the first step to establish a framework to build ethical AI solutions (Siau and Wang (2020)). The four constructs – *Transparency, Accountability, Responsibility, and Fairness* – are regarded as central in AI ethics.¹⁹ These constructs have been adopted when tackling *AI ethics* issues in IS development (cf., Vakkuri et al. (2019)). In a similar vein, Dignum (2017) proposed a so-called ART principles framework combining *Accountability, Responsibility, and Transparency (ART)* as a means to design and engineer ethical automated decision-making mechanisms in AI systems. ART principles serve as the basis for EU legislative discussions, augmented with Trustworthiness. Similarly, they are also covered in the IEEE EAD guidelines,²⁰ while *Fairness* is considered by professional organisations (e.g., ACM) and in other key geographies (e.g., US; Vakkuri et al. (2019)). In our solution, we addressed a combination of *ART principles, Trustworthiness, and Fairness*.

Furthermore, guided by Dignum (2018), from an AI ethics design perspective, we selected to tackle two categories of the AI ethics life-cycle *Ethics by Design*, which refers to achieving ethical data mining outcomes, and *Ethics in Design*, which refers to enhancing the FIN-DM process with tasks supporting the implementation of ethics. The third category of *Ethics for Design* is out of scope since it is not an integral part of the data mining process, but rather a part of regulating data mining experts' behaviour. We propose to govern that dimension in the form of a code of conduct, or similar principles within corporate or professional organisations (e.g., associations), to regulate ethical and compliance behaviours of employees or association members. Selected AI ethics categories and constructs are presented in *Concepts Universe* in section E in Figure 8.

In the *AI ethics sub-component* (section E in Figure 8) and the *Business Understanding phase*, we initially propose identifying the potential bias and risks of AI ethics issues for the data mining project by evaluating key areas where they might appear. In particular, AI ethics might be transgressed when using biased or incorrect input data, making discriminating decisions, and delivering unsafe solutions with a negative impact on well-being (ignoring human-centricity). This evaluation task is supported by *Checklist 3*, which specifies what key aspects to consider. For instance, for input data, these are quality and potential inherent bias (with the most common being selection and measurement biases). For the decision-making process, one of the key aspects to specify is potential discrimination scenarios, which can be tested when

evaluating data mining results. As a next step, in the *Data Understanding and Preparation phases*, a number of tasks to check and ensure data quality (incl. possible bias) are proposed. In the *Modelling phase*, we suggest tackling AI ethics either with Fairness Aware Data Mining²¹ applying adjusted techniques and algorithms (cf., Friedler et al. (2019)) or with traditional techniques. In the case of the latter, rigorous fairness testing (incl. counterfactual fairness testing²²) needs to be applied based on earlier defined scenarios. Then, concurrently identified bias and discrimination requires mitigation by applying model de-biasing techniques (cf., Hajian et al. (2016)). Further, in the *Modelling phase*, we proposed explanation and interpretation tasks to achieve algorithm transparency (via ensuring explainability) and providing accuracy. As a next step, in the *Deployment and Implementation phase*, human-centric design aspects are revisited, focusing on technical robustness, safety, and planning for the regular tuning and oversight of the data mining solution (as part of *Life-cycle management plan*). Finally, as part of the *Post-implementation life-cycle*, periodic quality reviews of datasets are proposed, while data mining solution monitoring is part of the service cycle. Traceability, reproducibility, and auditability aspects are ensured throughout the FIN-DM life-cycle by documenting the AI compliance process and outcomes that, in turn, underpin data mining's Trustworthiness.

While AI risks are heavily discussed and debated, *risk management* practices in the data mining process are under-investigated, practical guidance in the form of frameworks is scarce and fragmented, and at a higher abstraction level (tools and standards) are largely unavailable (Bradley (2020)). The risk management of data mining models is currently considered from a model governance perspective (cf., Kurshan et al. (2020)), a project management perspective (cf., Bradley (2020)), a specialised perspective, for instance, documentation (Richards et al. (2020)), or an automation perspective of model life-cycle conceptualised as ModelOps (cf., Hummer et al. (2019), Arnold et al. (2020)). Such conceptualisation is related to the adaptation of the software development and deployment life-cycle to data mining projects; however, this conceptualisation primarily focuses on enabling technologies and automation. In contrast, *quality assurance* in data mining, so far, has received more attention compared to risk management. In particular, guidelines for quality assurance in machine-learning-based applications have been proposed in (Hamada et al. (2020a)) and further extended in (Hamada et al. (2020b)). Ishikawa et al. defined key quality evaluation aspects and proposed a development process model for ML systems. In a similar vein, Studer et al. (2020)

expanded CRISP-DM for the development of machine-learning applications, which covered the entire project life-cycle addressing two key deficiencies of CRISP-DM: (1) lack of application (in our case deployment) tasks where a machine-learning model is maintained as the application, and (2) lack of guidance on quality assurance (Studer et al. (2020)).

Guided by software development practices in our design, we opted to combine *risk management and quality assurance* tasks as one component, thereby addressing both *Regulatory and Compliance* as well as portion of *Process gaps* related to quality assurance in CRISP-DM. Furthermore, by embedding and relating both elements, we expect to mitigate some model risk aspects with quality assurance tasks. In particular, we propose tasks to ensure *sound development methodology, intermediate testing, user acceptance tests, and proper model change policies and practices*. However, we replace the *testing* task with a much broader *verification and validation*, with a more extensive scope (e.g., assurance of requirements) and broader techniques. This is complemented by *auditability* requirements for data mining project artefacts and the data mining process itself, and *model monitoring* when a model is used. This conceptualisation is presented in *Concepts Universe* in section F in Figure 8, while solutions are specified in the *Risk Management and QA sub-component* in section F in Figure 8. There, we have made a distinction between two sub-processes – one workflow is inherent to the development life-cycle, itself, and constitutes tasks executed by project participants to ensure consistent quality in the development process. To improve development quality, the primary activity is testing based on respective plans on project quality targets and metrics. The other workflow is executed by peer reviewers who are external towards a given data mining project – i.e., they are not part of a development team. They provide independent quality assurance via *Quality Checkpoints* and aim to validate all key data mining project artefacts and outputs, including the requirements, data used, methods and design decisions, the model or solution prototype, and the model's implementation in systems. In the case of the latter, software verification and validation methods should be used. Lastly, model usage is independently validated to mitigate the model risk and ensure the model's conformance to model change procedures. Regular data oversight and tuning is ensured to keep control of model risk (as part of the *AI ethics* sub-component routines), complemented by regular model life-cycle management activities based on the plan defined in the *Deployment phase*. Deployed techniques include inspections, technical reviews, walkthroughs, conformance checks (adherence to processes), design and

product release reviews, etc., while the output is checked for the '3Cs' (correctness, consistency, and completeness; Bourque and Fairley (2014)).

Addressing Process Gaps (G7) We addressed *Process gaps* in two contexts. The lack of *controls and quality assurance mechanisms* for stand-alone data mining projects have been tackled in the *Risk Management and Quality Assurance sub-component*. Here, we present the mechanisms of the solution to address prerequisites for repeatable data mining at scale.

The most common approach to systematically address any business processes in the organisation is to consider it via the prism of the value chain or value delivery (cf., Dumas et al. (2018)). This paradigm is industry-agnostic and equally applicable to the creation of tangible products (e.g., manufactured goods), intangible assets (e.g., software), and services. Therefore, this methodology has been successfully adopted across various industries, including the IT domain, and widely incorporated into IT practices. Notably, ITSM (Information technology service management) is the most well-known and widely adopted approach to IT service delivery and management; it is process-centred (in contrast to technology-oriented) and focused on value delivery (Cusick (2020)). One of the ITSM frameworks, ITIL (Information Technology Infrastructure Library), is the most widely adopted approach for the management of IT services. ITIL defines IT delivery as a value delivery process with three core components – process inputs and outputs, process controls, and process enablers. Notably, it is successfully adopted and fits modern software development practices, like DevOps (Galup et al. (2020)), and has incorporated Agile and Lean delivery methods. Drawing an analogy to IT and software delivery, we can view data mining projects as similar IT delivery instances, which can be encompassed and supported by ITIL.

In the same vein, one more notable ITSM framework is COBIT,²³ which provides a comprehensive foundation for IT governance and management (cf., Haes et al. (2013)) in organisations. One of the key features of COBIT is its focus on governance and controls, including building comprehensive quality assurance governance and management.

In the context of achieving effective data mining at scale and tackling *Process Gaps* in CRISP-DM, we have selected the number of prerequisites and enablers *Foundation Capabilities* from the ITIL framework. Based on ITIL typology, the following attributes across 4 ITIL dimensions (part 1 in Figure 6) were chosen:

- Organisation and People – stakeholders management and data mining competencies (via knowledge management and codification systems)
- Information and Technology – data management, model development, deployment, and self-service technologies encompassing the whole data mining life-cycle including sharing the results with users
- Partners and Suppliers – planning of resources and internal coordination
- Value Stream and Processes – delivery models (horizontal vs. virtual cross-functional teams) and the design of service, delivery, and improvement processes

The relevant dimensions and attributes were complemented by sets of selected ITIL Management practices covering general management (portfolio, project, relationships, etc.), service (primarily related to implemented data mining models service management), technical (deployment and software development practices), and delivery approaches (including delivery models, and continuous integration, delivery, and deployment practices).

These prerequisites are well complemented by COBIT elements (part 3 in [Figure 6](#)), which refer to quality management at two levels: (1) data mining process' quality principles and policy (incl. quality management practices of data mining models), (2) the project's quality – quality management plans for respective data mining projects based on stakeholders' quality requirements.

As shown in [Figure 8](#), ITIL and COBIT elements at higher abstraction levels are cascaded down and natively integrated, where applicable, to the FIN-DM hierarchical process via respective phases, sub-components, and tasks. For COBIT elements, these are data mining project quality management tasks reflected in the *Risk Management and Quality Assurance sub-component of the Compliance phase*. For ITIL elements, these are data mining requirements reflected in the *Requirements phase*.

4.4. Fin-DM evaluation

In this section, we will report on the results of our evaluation of FIN-DM. We will particularly focus on the participants' perception of FIN-DM's quality (completeness, existence of overlaps, complexity, and presentation quality) as well as their perception of its usefulness, ease of use, future use intentions, and satisfaction. In addition, we will also discuss improvement suggestions by participants and elaborate on how we integrated them. The profiles of the participants are summarised in [Table 2](#).

4.4.1. FIN-DM quality evaluation

In the following, we will elaborate on our findings related to the quality perception of FIN-DM starting with its completeness, existence of overlaps, complexity, and presentation quality before discussing specifics about how it can be implemented in practice and connected to existing frameworks such as ITIL.

Completeness – “Nothing is missing for data mining . . . organisation and technology needs attention” The participants acknowledged the general completeness of FIN-DM in covering data mining aspects, especially the ones missing in the CRISP-DM model, as noted by one interviewee, *“It seems that most [. . .] of the things that were missing in CRISP-DM have been added here”* (DS2). In particular, the solutions addressing privacy, AI ethics “gaps”, and enhancements to deployment phase were very well received, as one interviewee noted *“I think it was very good that you captured the ethics and [. . .] the GDPR [. . .] or a AI and general perspective”* (DS5). At the same time, some organisational and technology aspects were acknowledged as missing, as one interviewee reflected *“As I mentioned, maybe these change management and real business implementations is missing, otherwise, [it] looks quite OK, at least all the bullet points and topics that need to be cared of [. . .]”* (Exp3).

Overlaps – “Iterations and Dependencies” Interviewees did not detect significant overlaps in the FIN-DM life-cycle; however, some interviewees classified such intersections as dependencies or necessary iterations intrinsic to data mining development as such. Some felt that *“no, I, couldn't say now if anything overlaps, I mean, to me this is a very iterative process, did you jump back and forth all the time during a project”* (DS5), while others considered, *“maybe there are some overlaps in the quality management and evaluation or validation part, and a little bit with a new deployment and monitoring parts. But it's like maybe not so much overlap, but more like dependencies”* (DS2).

FIN-DM complexity – “Generally adequate . . . needs to be simplified a bit” The majority of the interviewees perceived FIN-DM as balanced with respect to its complexity and the details it includes, as one interviewee put it *“I think it's quite balanced because there doesn't seem to be anything that is not needed [and] I immediately don't even see that many options, how it could be simplified that much further without losing important information.”* (DS2). However, the comment below suggests the need for some simplification to increase usability and foster easier adoption, *“I think that there are a lot of details that have been put into this new framework, and they are all very much needed [. . .] but, it is also if it should be really helpful and usable for the different organizations and easy to implement, then it needs to be*

Table 2. Experts' – participants in evaluation – profiles and key characteristics.

No.	Interviewees' profile	Primary area of expertise (*advanced academic degree, PhD, in the field)	Current role	Experience in the field (years)
1	Data Scientist 1 (DS1)	Data science, advanced analytics, and quant finance	Senior expert in data science	8+
2	Data Scientist 2 (DS2)	Data science and advanced analytics	Lead expert in data science	8+
3	Data Scientist 3 (DS3)	Data science and advanced analytics*	Senior expert in data science	10+
4	Data Scientist 4 (DS4)	Data science, advanced analytics, and quant finance*	Senior expert in data science	10+
5	Data Scientist 5 (DS5)	Data science, advanced analytics, and quant finance	Senior expert in data science	10+
6	Project Manager 1 (PM1)	Project management, data science, and advanced analytics	Lead project manager in data science	8+
7	Project Manager 2 (PM2)	Project management, data science, and advanced analytics	Lead project manager in data science	8+
8	Expert 1 (Exp1)	IT and software development	Agile coach and operational leader in IT	8+
9	Expert 2 (Exp2)	IT, software development, and risks in IT	Operational leader in IT	20+
10	Expert 3 (Exp3)	Project management, data science, advanced analytics, and tech	Analytics consultancy operational leader	10+
11	Expert 4 (Exp4)	Data science, advanced analytics, technology, and strategy	Strategic leader in data science and advanced analytics	20 +

a bit simplified maybe. But then, if you simplify it too much then I'm scared that we will lose out of a lot of very, very important aspects. So, yeah, somewhere in the middle" (PM2).

FIN-DM presentation quality – “Complicated but could be used with adequate support”

Interviewees provided valuable remarks on the presentation quality of the initial FIN-DM prototype, focusing on four key presentation aspects:

- terminology and definitions of elements – more clarity required (e.g., differentiation between functional and non-functional requirements, robustness of algorithms were vaguely defined) - *those non-functional requirements that we talked about again availability, robustness I don't know, also something redundancy of the data and so on, but these are for the IT system and the software system mainly.* (DS3)
- application guidance – too short and requires expansion
- communication design – difficult to grasp when looking - *"... from framework perspective, You should think on the way how to simplify the common content like presentation, means of presenting the content, showing the content, because otherwise it deserves too much time for explanation. And a lot of efforts of understanding."* (Exp 3)
- conceptual design – interrelations between elements (e.g., ITIL elements not natively integrated) - *Just the question is how to easily incorporate it in the overall framework so that it's clear that You need to pay attention to this part. I would say that maybe at this point it's not natively integrated to each other.* (Exp3)

As commented by one interviewee, *"I would say it's maybe from presentation perspective over-complicated, from content perspective not [it is] adequate, that's for sure, everything is needed and it should be checklisted for [the] project manager. But from framework perspective, you should think [about] how to simplify the common content like presentation, means of presenting the content, showing the content, because otherwise it deserves too much time for explanation. [...] And a lot of effort [for] understanding."* (Exp3).

Reference Process Instantiation²⁴ – “Requires adequate technology and documentation support”
Interviewees also mentioned how the instantiation of FIN-DM would benefit from the support of adequate tools, technologies, and process documentation. In reference to this issue, an interviewee said, *"Well, those I don't remember exactly about those architectural and business value verification documents, or checklists, or agreements, but I trust this can be done even in JIRA somehow or something like this by checkpoints. So, I think it can be realized by technology. So, I think this process has lots of underwater artifacts that are not mentioned here. Maybe you can even later visualize it somehow by BPM and diagram and adjust show what may happen and what may be the output"* (Exp1).

FIN-DM integration into existing frameworks
Interviewees also proposed improving and distinctly positioning FIN-DM by specifying the purpose, and the issues it intends to solve, more clearly. In the same vein, interviewees also suggested differentiating FIN-DM from other software development frameworks, as noted *"The difference between this framework and software frameworks could be a bit more emphasized explicit, and also*




Figure 9. Results of the questionnaire. All responses were given on a 5-point scale, which were anchored between *strongly disagree* (1) and *strongly agree* (5). The bars indicate the mean (m) and standard deviation (SD) for each scale.

I think explainability, slash interpretability could be also a more explicit and more bold, more visible in the framework” (DS3).

4.4.2. FIN-DM acceptance evaluation

Related to the acceptance of FIN-DM we will subsequently report on the participants’ perception of its ease of use and usefulness as well as their intentions to continue using it and their overall satisfaction.

FIN-DM ease of use The questionnaire results indicate that the participants generally perceived FIN-DM to be easy to use, as evident by a mean value above 4 (as presented in Figure 9). This score, however, ranks the lowest of all the questionnaire scores. Taking this together with the previously discussed findings related to the participants’ perception of the presentation quality of FIN-DM, however, suggests that the understandability can still be improved, e.g., by reducing complexity or providing adequate guidance and support.

FIN-DM usefulness Our participants generally perceived FIN-DM as useful, as evident by it receiving the highest score and lowest standard deviation among the aspects in the questionnaire (Figure 9). This finding is underpinned by interviewees stating that, “[...] of course there is no doubt that [the new elements] will be useful. They will help to achieve more transparency, clearance and maybe more mistake-proof environment. And of course, understanding between counterparties, and no missing processes or no gaps and disagreements as well” (Exp1). This perception is also supported by other interviewee, “Absolutely, undoubtedly, the artifacts that you have brought they add value to the whole process, and make it more manageable. For the further industrialization and scaling, it’s absolutely must” (Exp2).

In the following, we will go into more details regarding the interviewee perception of the usefulness of FIN-DM to address existing gaps before discussing

its different components both from a micro (individual FIN-DM component or data mining project) and macro (life-cycle or repeatable data mining settings) perspective.

FIN-DM addresses existing gaps Participants also extensively discussed how FIN-DM solves CRISP-DM ‘gaps’. The central issue for data mining projects in data scientists’ view was business value realisation. They, therefore, perceive the positive impact and potential of FIN-DM to lie in increasing the business value realisation of data mining projects. As noted by one interviewee, “as I understand this CRISP-DM, this [FIN-DM] is more integral. Basically, it’s like trying to concentrate all the required things from the beginning. And, I think that the main problem now is that we built an impressive analytic solution and that sometimes does not add business value. So, defining like this definitely helps to bring this business value on the table” (DS4). Furthermore, FIN-DM was characterised as an up-to-date, modern, and reusable data mining process that solves typical data mining process problems by design.

FIN-DM Components View – “contribution and positive impact per component” In terms of individual FIN-DM components, interviewees also emphasised the importance and necessity of the proposed Requirements elements in conjunction with Business validation later in the data mining life-cycle. “I would like to emphasize that the relationship to business, like a business understanding, business validation on the outputs, et cetera, that was really very valuable to see” (DS5). The participants perceived the Requirements as an integral part of the data mining project to timely discover, capture, and specify key data mining project prerequisites, such as data, tooling, deployment patterns, and infrastructure, as commented: “I think it’s something that should be investigated [...] at the very beginning any data mining project [...] and there have been many examples in the past for us where we start with a project [...] and we discover half-way that we

lack something” (DS5). Another interviewee concurred: *”I liked that idea that immediately in the beginning requirements not only towards business but deployment, data [...] identified since the start [...] that’s valuable, that is frequently missing in business, that business is starting to [deploy] requirements quite lately, and they witness themselves a lot of problems because in many cases their infrastructure is not ready to deployment so that is definitely valuable”* (Exp3).

As an aspect closely related to the *Requirements*, participants reflected on the *Deployment and Post-deployment* additions introduced in FIN-DM. In particular, the model governance and maintenance aspects and follow-up activities of FIN-DM were overlooked in other standard processes as the following comment shows *”So, I would say that this is extremely valuable part, especially for companies that are starting their analytical journey. Because when they start, they mostly focus on modelling, on results, and just delivering the model that works and predict something that they expect. But they don’t think about the long-term maintenance of [...] their solution. So, I would say that if in the beginning they would start thinking about the ownership of the model, like governance, about the principles, how their maintenance assured, model is giving good results, it would result in better ROI for their models. Because the frequent problem that analytics is done, but it’s not fully implemented in the business process. And even if it’s implemented, in many cases, nobody follows-up on the model performance. So, that’s very important part which needs to be discussed for each project somewhere in the beginning even, like owner of the model”* (Exp3). In the same way, the importance of the *Post-deployment* phase, especially to ensure the data mining solution has been properly used, has been emphasised by other participants: *”And yeah, I think that this is a very nice way of thinking about data mining projects.[...] I especially like the last phase, [...] post-deployment, because without that, what’s the point of all the products. So, it’s like I’m creating something that will never be properly used [...]”* (DS3).

Another critical element discussed extensively by interviewees has been the *AI Ethics and Compliance* phase. Interviewees especially emphasised the AI ethics context, which is very new in data mining. Also, there is a lack of guidelines and practices in traditional process models. As noted by one interviewee, *”I think this is quite necessary to get the CRISP-[DM] or the methodology up to date. Since both the GDPR and ethics are very important at the times today. I absolutely think that we should incorporate things like this ... Absolutely”* (DS5).

The critical importance of AI ethics for the financial services industry, in conjunction with GDPR compliance and personal data handling, has been highlighted as well, *”I believe it’s both adequate and that it serves its*

purpose. I think this is quite necessary in order to get the Crisp-[DM] up to date. Since both the GDPR and ethics are, are very important at the times today. I absolutely think that we should incorporate things like this and have this in mind throughout like all of our cases.” (DS5). Furthermore, the imminent need for larger organisations and industries with a lot of data to control and tackle specific AI ethics risks, e.g., biases were emphasised. As one interviewee put it, *”I think this is super well-appreciated, welcome, welcome component here, because I, as I mentioned earlier, I don’t think this part has ever been covered, or at least not extensively in these traditional frameworks. [...] it makes the approach more solid both with regards to ethics as such, but also with regards to, as you highlighted here compliance. Because those are becoming imminently more important in our work, especially if you consider [this] from a banking perspective – banking is about trust and these two components directly address underpinning a bank’s trust with data and customers”* (Exp4). Lastly, participants recognised and appreciated *Quality management* as a required activity. Concurrently, the critical role of quality management in two different contexts was emphasised, (1) as a stand-alone data mining project and (2) as an integral part of overall organisational practices. For example, *”Yeah, and that is absolutely needed. Yes, I mean, just both on the high level, as you mentioned, the more quality management plan, more if you have these more regulatory models that need to be in place, we have to ensure certain qualities and there are different departments involved. But also once the model is built, that we have these peer review processes, making sure that we are always having four eyes on each step”* (PM2).

Macro Settings – Industrialised Data Mining at Scale The usefulness of FIN-DM in broader settings for scaling and industrialising data mining as an internal service has been proposed. In particular, FIN-DM would help establish repeatable data mining and thus resemble IT established processes of IT services’ portfolio management. On this point, the interviewee commented, *”Because I see that [the] data mining role is becoming more and more essential in each organization, and it’s not anymore hidden somewhere in the corner like it was before, just some kind of R&D.[...] And of course the services, because your data mining, it’s closely related with using the technology, and maybe later it will become some kind of a service, which will also require some SLA. And SLAs mean that, of course it has, has to be managed by ITIL, [...] And of course, technical things, I think they also are based on good practices and they are relevant. So, basically it states everything that is needed, and it will be a guide for people who are unsure what they are doing or want to establish. So, this would serve as a framework even for those who would like to establish something like this. It’s my opinion, because you have all the components, it’s*

like the establishment guide” (Exp1). The other participant resonated, “absolutely, undoubtedly, the artifacts that you have brought adds value to the whole process, and makes it more manageable. For the further industrialization and scaling, it’s absolutely must” (Exp2).

Micro Settings – Stand-alone Project The project participants extensively discussed FIN-DM’s usefulness in terms of goals and benefits achieved in stand-alone data mining projects. They particularly mentioned the integrated, holistic approach for data mining projects from an organisational perspective. In particular, the crucial alignment between different functions, stakeholders, and required capabilities to support end-to-end project execution was emphasised and perceived well. For example, one interviewee said “[...] for me, this is just my point of view, CRISP-DM lacks the bridge of three things – business, IT, and data science, because we are doing the models, but if the model gets too complex, then there is nobody else that will ensure that it would be operationalized. And this is where IT should jump in, but [...] there’s no link [...] or somehow it gets lost [...]. And on the other hand, sometimes you develop some models, whatever, but the link with business is lost as well. [...] there is a lack of clarity in the business goal, or that either we do not understand the business objective, or that we want to minimize overhead and leap into the interesting bit of the project, analyzing the data and so on, and to [open] these results in very interesting models that don’t meet a real business need. So, a methodology that connects these three parts, [...] the real end-to-end process that has to touch all these three parts” (DS4).

FIN-DM intention to use and satisfaction Compared to the *Perceived Usefulness* results, the interviews and questionnaire exhibited a higher variability in answers on other *Acceptance Perception* aspects, such as the *Intention to use* and *Satisfaction* of FIN-DM. Evaluation participants expressed firm *Usage intentions* both in interviews and in the questionnaire (Figure 9). In the latter case, the *Usage intention* is preceded by the overall *Satisfaction* of potential users with FIN-DM – it is on average in the *Very satisfied* category. Lower *Satisfaction* relative to the *Intention to use* of FIN-DM can be explained by, and was captured in, the *Ease of use* measure previously discussed in conjunction with *Quality perception*. Related to future usage intentions, the participants included the overall logic and concept of the FIN-DM, as noted “I think, anyways, later it will be visualized more and more. I take this as a draft material, so I have no objection, objections towards this format and looking absolutely logically . . .” (Exp1). “[...] it is so detailed but it is good. [...] but I understand the main picture, the main big picture, which is this – to connect these three main parts, and to make it like full end to end process. This

part I understand well, maybe just the particularities, the specific artifacts, and so on and I want to emphasize that, that is very good” (DS4).

Also, interviewees further contextualised potential usage scenarios and proposed the usage of FIN-DM by various data mining projects participants and stakeholders, as noted: “I think it’s also very important that this framework is not only used by data scientists, but by the whole cross-functional team, including product owners and, and so on. And, and even maybe, perhaps, stakeholders and so on. So, I think it’s like very important that everyone understands it” (DS2). Lastly, participants also expressed conditional usage – provided concrete improvements suggestions are fulfilled – for example, “If it were simplified, I will definitely take it to my next project” (Exp3).

4.4.3. Suggested improvements

Here, we present and discuss key improvement proposals received in FIN-DM’s evaluation (*Outcome 6* in Figure 4). We group them by design aspects, missing elements, components, and prerequisites. A complete list of all suggested improvements as formulated by interviewees and how they were addressed in the final FIN-DM version is available on figshare (<https://figshare.com/s/37eaab47cb024e3a2023>). Overall, we have received 55 improvement proposals, with 51 unique suggestions. We classified them into four distinct categories – *Minor, Medium, Major, and Critical* – based on the suggested improvements’ complexity and impact. Close to 80% (40 suggestions out of 51) of all the proposals have been categorised as *Minor and Medium*. The remaining 11 items were *Major and Critical*. There have been no conflicting suggestions except one requiring substantial simplification, which originated from the strategic leader. We have resolved this conflicting suggestion by introducing a new design principle (DP4) and integrating clear-cut process model adaptations to differentiate users’ categories based on data mining project roles (DF4). Below, we proceed with the analysis of the *Major and Critical* suggestions, their potential adverse impacts, and how we addressed or mitigated them for the final version of FIN-DM.

Design Aspects Some interviewees reflected on FIN-DM’s complexity, which in turn could hamper its adoption. Talking about these issues, an interviewee noted, “CRISP-DM [is] extremely simple, so here you have a combination, so you really need to think how to put it together to make it simple, because otherwise all the content is correct, but it’s complex. And if it’s complex, it means people will not understand, will not adopt, and will not stick to it, so that’s the risk” (Exp3).

Furthermore, interviewees underscored two key trade-offs and balances to achieve in FIN-DM. One was associated with equalising data mining-related

elements and other elements from the IT domain, which some interviewees perceived to be over-emphasised: *"It was a lot on the requirements in the beginning, and it was a lot on the deployment phase, but those steps in between, I think those are the main tasks of the data scientist, because of course the data science project needs all kinds of skills"* (DS3). The other trade-off we found was between simplicity and details. As pinpointed by one interviewee, *"[...] I think that there are a lot of details that have been put into this new framework, and they are all very much needed. The old framework, the CRISP-DM kind of lacked a lot of phases and tasks and, and so, but I mean, it is also to be really, if it should be really helpful and usable for the different organizations and easy to implement, then it needs to be a bit simplified maybe. But with, then, if you simplify it too much, then I'm scared that we will lose out of a lot of very, very important aspects. So, yeah, somewhere in the middle [...]"* (PM2).

Another aspect proposed for improvement relates to FIN-DM's flexibility and its ability to accommodate external changes, i.e., to stay dynamic. To this end, changes in IT domain-related elements have been emphasised as the ones requiring an update when the underlying ITIL and COBIT frameworks change. As one interviewee put it, *"couple of years back, we didn't even consider to include general IT processes into this, and, and nor did we. You know, we were not talking about CI/CD and software development practices. So, so there is a bit of a so to say in, there is a bit of so to say trends that come and influence the formations of these frameworks. The question is now, that process to also re-engineer or revitalize or update the framework itself, for a use case, for models portfolio, how does that process look like? And how do you make sure that, because now you refer to, you know, different frameworks and practices and so, how do you make sure that you always stay on top and make sure you bring in what is most relevant into your processes and frameworks and, and omit the ones that are not. It's more of a philosophical question, it's more of a meta question, than rather than the more of a specific use case question"* (Exp4).

Interviewees also recognised the need for change management approaches towards FIN-DM e.g., based on external requirements. The interviewees perceived the emergence of external disruptive technologies and business models as a core source for such changes. As one interviewee commented, *"[...] one is that how do we ensure [...] in the processes and stages stated here capture for example, all the new technology enablement, that brings times, reduce times to market, and it includes efficiencies. So, how do you include that, those components into this? How do you include the new frameworks around Open Banking, APIs that, that so to say, and cloud as such, the technology enablement that opens up new ways of consuming information. And with that comes also certain*

level of controls and qualities and process checks that need to be introduced, which traditional approaches maybe perhaps or not really taking into account" (Exp4).

Also, the required adaptation of FIN-DM towards new business models can be related to the emergence of new business ecosystems,²⁵ and associated collaboration between participants. This is especially relevant for the financial services sector with the active emergence of open banking ecosystems fostered by respective regulations, most notably the PSD2.²⁶ Talking about the issues, as one interviewee noted, *"The other thing is that in, I guess this is more of a business understanding stages [...] and data readiness. Maybe this is how to say, this is to what extent your frameworks take into account both the internal and external perspective. And with external, probably more external perspective, I mean, if you are operating in an open banking environment, for example, in our case, and you will need to adhere to bunch of other companies connecting to your data and exchanging data with other partners and so on. How, what sort of requirements that external environment will have on your use cases? And, and as part of your business understanding, you need to have steps, activities that capture that and make sure that further requirements internalized. This is, I feel, probably all of these frameworks can improve upon, or at least clarify the standpoints on"* (Exp4).

Missing Elements and Pre-requisites Interviewees also identified missing elements in FIN-DM. In particular, a necessity to transform business processes (via adequate change management) to embed and integrate data mining solutions has been emphasised, *"And actually one important things what is maybe a bit missing, [...], it's regarding deployment. What we observe quite frequently, this is change management because in many cases, business might have put understanding why they need like product owner has a good vision. The model is designed like according to the best practice. (The) model is good enough, it's, it deployed, it works even maybe on a real-time, but all organization is not adapted. Then, they're just not using in their daily processes, they follow old, old methods, old reports, old practices, and et cetera. So, I would say (a) crucial part after deployment, or in parallel with deployment should be business change management. [...] to make it happen you need really to make a change inside the proposed system, not just make mathematics"* (Exp3).

The other aspect is associated with ITIL and COBIT's more native integration with the data mining life-cycle, and embedding concrete ITIL and COBIT practices into the FIN-DM process. This aspect is complementary to balancing and specifying data mining and IT perspectives discussed above. As noted by one interviewee, *"it would be nice to direct like these both COBIT and ITIL frameworks towards data mining, because this, for me [...] are [...] too*

general, which applies to every project of course, but I can't right away see what the specific for data mining project in when we are making use of this COBIT and ITIL. So, it would be interesting to see what are the best practices in ITIL or COBIT that we can specifically use for data mining projects" (DS3).

Some interviewees expressed the necessity for detailing roles and responsibilities and governance aspects in the proposed reference process. As one interviewee commented, "And maybe responsibility [...] who is responsible for doing what." (Exp1), while other noted "with regards to more governance of the different stages that you might have some more questions. In regard to the roles and responsibilities needed in various stages to execute these tasks" (Exp1).

Further, interviewees emphasised the risks and aspects which would hamper the straightforward application and adoption of FIN-DM. In particular, the risk of moving far away from the lean process and being too restrictive was pinpointed, "Well, I think this is a very important step, but at the same time [...] the threat is that we move far away from a lean process, like it's if we just add lots of checklists and documents, and lots of meetings, and lots of discussions about around this. If it's too much, and it can easily be, get too much, because you want to make sure you cover all the requirements for [...] functional requirements, what they expect, the nonfunctional requirements, the deployment plan requirements, the maintenance requirements, and this can get a lot" (DS3).

Also, it was emphasised that it might be easier to apply FIN-DM in organisations with established IT foundations and practices, as one interviewee considered, "Yeah, I would not say that it's too much details, on the contrary, under each bullet point that you have provided there is like much deeper area of activities and knowledge. And then basically if you imagine starting from the scratch, I would say that it would be massive work even to you know, get ready the framework as such and the practices, et cetera. If we already apply this thing on the existing for example, organization, where (the) IT foundation is quite established, then it would fly much easier" (Exp2).

5. Discussion

In this research, we addressed the non-existence of a specialised data mining process for the financial service sector. As a result, for financial industry practitioners, the research outcome is a domain-specific process model – FIN-DM – which supports the execution of data mining. FIN-DM is contextualised and is up-to-date with the latest technological and regulatory developments impacting financial services. To this end, FIN-DM captures and provides data mining experts with concrete phases, tasks, and activities to tackle a number of critically important regulatory

requirements and recommendations that emerged over the last decade. Some given regulations are not financial-services specific, but financial services have been one of the key sectors impacted. In particular, FIN-DM supports a systematic approach to effectively address *privacy regulations* (e.g., GDPR) and execute fully *privacy compliant data mining life-cycles*. It also provides a practical approach to execute best practices of *AI ethics and AI ethical risks mitigation* in all key phases of the data mining project. The proposed tasks support effective tackling of the key, and most widely recognised, AI ethics principles and concerns – *Transparency, Responsibility, Accountability, Fairness, and Trustworthiness*. Also, FIN-DM will be helpful to practitioners in the financial services industry to design the operations of data mining function in line with modern *internal control operating models* (3 Lines of Defence Model) and practically fulfill *model risk management requirements* in data mining *model development, validation, and use*. Lastly, FIN-DM supports practitioners with practical solutions to embed *quality assurance* as an integral component of the whole data mining life-cycle – introducing both internal project *verification* tasks and external project *quality controls (validations)*.

Additionally, FIN-DM supports improving *business validation and actionability* of data mining results and provides an improved and structured data mining life-cycle *requirements elicitation and management* end-to-end. Furthermore, two other key CRISP-DM gaps – the lack of actual data mining model deployment/the transition of data mining models into software products and model life-cycle management – have been specified explicitly. To ensure adequate technological support and governance to data mining projects, we have incorporated a number of ITIL and COBIT elements in life-cycle practices and as a broader enablers' set. Finally, FIN-DM has multiple advantages and limitations when applied. It can serve as an establishment guide or blueprint for scaling and industrialising data mining functions; nevertheless, it is likely to be more effective in organisations that have already established IT foundations and operate within standard IT industry frameworks (e.g., ITIL). Moreover, knowledge of the CRISP-DM life-cycle and prior experience would be beneficial to foster more effective organisational adoption. Some of the proposed solutions to gaps might not be entirely financial-sector specific but rather applicable in the context of other sectors tackling similar challenges (e.g., telecom). A good example of such generic solutions is *privacy* and *AI ethics* components.

FIN-DM has undergone an initial evaluation, but it still needs to mature by being tested in practice. At the same time, as underscored in Czarnecki and Dietze (2017), the evaluation of reference models constitutes a problem and research challenge. In particular, the




Figure 10. Design science contributions and knowledge categorisation (Gregor and Hevner (2013)).

reference model, by nature, is generalisable; thus, its development is decoupled from its application. Thus, the evaluation in practical settings of data mining projects will be focused on FIN-DM’s applicability. To this end, as pinpointed in Czarnecki and Dietze (2017), each application of reference process increases their chance of adoption.

From the design science perspective, based on Gregor and Hevner (2013) classification, FIN-DM classifies as *improvement research*, as we have attempted to create a solution for the pre-discovered gaps in CRISP-DM within the application context (financial services sector; see, Figure 10). Additionally, our research has also contributed to the theory in the form of *nascent design theory* (Gregor and Hevner (2013)). Concurrently, by creating FIN-DM we have enhanced Lambda knowledge space²⁷ (as per Figure 10) with a more in-depth understanding of data mining models, methods, and their instantiation specifics in the context of financial services. FIN-DM design and development is also one of the first works where the formal design science methodology was applied to derive a process model.

This research also has a number of inherent limitations, primarily associated with conducted evaluation and techniques. In particular, the evaluation has been performed in one financial institution with a small set of selected participants, which limits the ability for the results to be generalised beyond the context of the study (external validity risk as per Runeson et al. (2012)). We mitigated this concern by selecting interviewees with different backgrounds and roles in data mining projects. We also invited external experts from outside the organisation. However, more studies are recommended in other financial companies. Next, FIN-DM has undergone an initial evaluation, but it requires a complete evaluation by applying FIN-DM in actual data mining projects. Apart from external validity, our study has construct validity risks and reliability risks. Construct validity refers to the extent

to which what is studied corresponds to what is intended and defined to be studied. In our research, the interview method can be a source of construct validity risk. We mitigated this concern by verifying the interviewees’ understandings of the questions and reconfirming the responses we received. Finally, reliability risks concern the researcher’s biases in generating and interpreting study results. We have tackled this risk by logging coding procedures and interviews and by adopting an iterative research process with regular validations within our research group. We also ensured the replicability of the research steps and results by keeping evidence-tracking the research materials and process.

6. Conclusion and outlook

This research has presented the design and evaluation of a data mining reference process for the financial services domain – FIN-DM. Guided by DSRM methodology, and combining behavioural science and design science paradigms, we developed and proposed design artefact. FIN-DM was designed to tackle the gaps in the CRISP-DM standard process, which have been identified in previous studies on the use of CRISP-DM in the financial industry. We, first, consolidated evidence from earlier studies on the gaps of CRISP-DM. We proceeded with formulating the requirements and design aspects of the new data mining process, FIN-DM, and developing its prototype. Next, we evaluated the prototype by conducting demo sessions and semi-structured interviews with the experienced data mining and IT practitioners actively engaged in data mining projects in the financial services industry. We also constructed and distributed a qualitative questionnaire among the pre-selected group of data scientists. Finally, we integrated feedback received from the evaluation to improve the final version of FIN-DM.

We implemented a number of new design principles in FIN-DM. We also addressed the somewhat restrictive nature of CRISP-DM by incorporating full iterativeness among all life-cycle elements. We also embedded higher flexibility and adaptability than in the standard process (CRISP-DM). We proposed dynamic adaptation mechanisms that would allow FIN-DM to stay relevant and adapt to its composing frameworks (ITIL, COBIT) and external environment changes, such as disruptive technologies and the emergence of new business models.

As mentioned, the immediate direction for future research could be testing and evaluating FIN-DM in the settings of an actual data mining project(s). In this context, FIN-DM applicability is the most crucial aspect to investigate. Another direction for future research could be associated with the broader scope of data mining problems and solutions embedded in FIN-DM, which go beyond the specifics of the financial sector (e.g., *privacy* and *AI ethics* components). Therefore, the potential application of FIN-DM in a broader solution space, i.e., other industries (telecom) and organisations, could also be investigated.

Notes

1. Cross-industry standard process for data mining.
2. Anti-Money Laundering.
3. E.g., computer-based IS, or IS/IT design artefacts, such as techniques, tools, development, planning and management methods, IS/IT security and risk management practices, etc. (J. Venable (2006)).
4. The key principle behind Objectivist/Positivist research paradigm is that the researcher is independent of the investigated phenomenon, and the purpose of the research is to design an artefact achieving objectively defined results (J. R. Venable et al. (2017)).
5. In the Interpretivist paradigm, researchers engage in the social setting and attempt to understand phenomena through the meanings and interpretations that people assign to them (Adam (2014)).
6. <https://figshare.com/s/1bda7ccadaa254fcabe1>.
7. The initial and final coding schemes are available on <https://figshare.com/s/d2bf5084f3e6cfc1af80>.
8. Permanent FIN-DM materials repository is placed at <https://figshare.com/s/e9fed5237a8577c8d0c>; website for user-friendly navigation of the FIN-DM is available at <https://fin-dm.info>.
9. SLR I – cross-domain systematic literature review.
10. SLR II – systematic literature review in the financial services domain.
11. General Data Protection Directive (2016/679) – EU Regulation on data protection and privacy in EU and EEA countries.
12. All FIN-DM materials are available at repository <https://figshare.com/s/e9fed5237a8577c8d0c> and on the website <https://fin-dm.info>.
13. Information Technology Service Management.
14. Control Objectives for Information and Related Technology.

15. Model risk is defined as a subtype of operational risk and refers to the potential for adverse financial, reputational, or regulatory consequences for decisions based on incorrect or misused model outputs and reports.
16. GDPR stipulates 9 core principles – Fairness, Lawfulness, Minimisation, Integrity, Confidentiality, Accuracy, Data Storage Limitation and Security, Accountability.
17. The other interrelated aspect is personal data collection without legitimate grounds, but since data is used as input in the data mining life-cycle it is out of scope.
18. GDPR requirement to collect and process only as much data as absolutely necessary for the purposes specified.
19. There is not, yet, universal agreement achieved on these values being the core of AI ethics (Vakkuri et al. (2019)), though there are alternative principles proposed (cf., Morley et al. (2020)) with Beneficence, Non-maleficence, Autonomy, Justice, and Explicability constructs.
20. Guidelines on Ethically Aligned Design.
21. Developing data mining models and systems which are discrimination-conscious by design (Hajian et al. (2016)).
22. AI decision is counter-factually fair towards an individual if it is the same in (a) the actual world and (b) a counterfactual world where the individual belonged to a different demographic group (Kusner et al. (2017)).
23. Control Objectives for Information Technologies.
24. General approach for mapping the generic process model involves specialising (or instantiating) generic contents according to concrete characteristics of the data mining project context (Chapman et al. (2000)).
25. Defined as the network of organisations – including suppliers, distributors, customers, competitors, government agencies, etc. – involved in the delivery of a specific product or service through both competition and cooperation (cf., Palmi'e et al. (2020), Adner (2017)).
26. Payment Services Directive 2015/2366, promotes a harmonised payments services market open to operate for payment providers beyond traditional incumbent banks, requiring close cooperation between payment services providers in the execution of payments and data exchange (cf., Cortet et al. (2016)).
27. Consists of constructs, models, methods, and instantiations (Gregor and Hevner (2013)).

Disclosure statement

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

Funding

This work was supported by the European Regional Development Fund [EXCITE].

References

- Adam, I. O. (2014). The ontological, epistemological and methodological debates in information systems research: A partial review. *Epistemological and Methodological*

- Debates in Information Systems Research: A Partial Review (March 2014)*. <https://dx.doi.org/10.2139/ssrn.2411620>.
- Adner, R. (2017). Ecosystem as structure: An actionable construct for strategy. *Journal of Management*, 43(1), 39–58. <https://doi.org/10.1177/0149206316678451>
- Alturki, A., Gable, G. G., Bandara, W., & Gregor, S. (2012). Validating the design science research roadmap: Through the lens of “the idealised model for theory development”. In S. L. Pan & T. H. Cao (Eds.), *16th Pacific Asia Conference on information systems, PACIS 2012, ho chi minh city, Vietnam, 11-15 July 2012* (pp. 2). <http://aisel.aisnet.org/pacis2012/2>
- Anand, S. S., Patrick, A., Hughes, J. G., & Bell, D. A. (1998). A data mining methodology for cross-sales. *Knowledge-based Systems*, 10(7), 449–461. [https://doi.org/10.1016/S0950-7051\(98\)00035-5](https://doi.org/10.1016/S0950-7051(98)00035-5)
- Anand, S. S., & Büchner, A. G. (1998). Decision support using data mining. *Financial Times Management*. Financial Times Pitman.
- Arnold, M., Boston, J., Desmond, M., Duesterwald, E., Elder, B., Murthi, A., et al. (2020). Towards automating the AI operations lifecycle. *CoRR, abs/2003.12808*. <https://arxiv.org/abs/2003.12808>
- Berendt, B., & Preibusch, S. Better decision support through exploratory discrimination-aware data mining: Foundations and empirical evidence. (2014). *Artificial Intelligence and Law*, 22(2), 175–209. <https://doi.org/10.1007/s10506-013-9152-0>
- Bhattacharjee, A. (2001). Understanding information systems continuance: An expectation-confirmation model. *MIS Quarterly*, 25 (3), 351–370. <http://misq.org/understanding-information-systems-continuance-an-expectation-confirmation-model.html>
- Bourque, P., & Fairley, R. E. (2014). *Guide to the software engineering body of knowledge (swebok (r)): Version 3.0*. IEEE Computer Society Press.
- Brachman, R. J., & Anand, T. (1996). The process of knowledge discovery in databases. In *Advances in knowledge discovery and data mining* (pp. 37–57). American Association for Artificial Intelligence.
- Bradley, P. (2020). Risk management standards and the active management of malicious intent in artificial superintelligence. *AI & SOCIETY*, 35(2), 319–328. <https://doi.org/10.1007/s00146-019-00890-2>
- Buchner, A. G., Mulvenna, M. D., Anand, S. S., & Hughes, J. G. (1999). An internet-enabled knowledge discovery process. *Proceedings of the 9th international database conference*, 1836. San Diego, California, USA: User-Driven Navigation Pattern Discovery from Internet Data, Web Usage Analysis and User Profiling, International WEBKDD'99 Workshop, 74--91. https://doi.org/10.1007/3-540-44934-5_5 .
- Cabena, P., Hadjinian, P., Stadler, R., Verhees, J., & Zanasi, A. (1997). International Business Machines Corporation (San Jose, C., & International Technical Support Organization). *Discovering data mining: from concept to implementation*. Prentice Hall PTR.
- Cao, L. (2018). . AI in Finance: A Review. *ACM Comput Article* 39. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>.
- Cao, L. (2021). AI in finance: Challenges, techniques and opportunities. *CoRR, abs/2107.09051*. <https://arxiv.org/abs/2107.09051>
- Chandra, L., Seidel, S., & Gregor, S. (2015). Prescriptive knowledge in IS research: Conceptualizing design principles in terms of materiality, action, and boundary conditions. In S. X. Bui & R. H. S Jr. (Eds.), *48th Hawaii international conference on system sciences, HICSS 2015, kauai, Hawaii, usa, January 5-8, 2015* (pp. 4039–4048). IEEE Computer Society. <https://doi.org/10.1109/HICSS.2015.485>
- Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., and Shearer, C., et al. (2000). *Crisp-dm 1.0 step-by-step data mining guide*, 9(13). SPSS Inc, pp. 1–73.
- Cortet, M., Rijks, T., & Nijland, S. (2016). Psd2: The digital transformation accelerator for banks. *Journal of Payments Strategy & Systems*, 10(1), 13–27.
- Cusick, J. J. (2020). Business value of ITSM. Requirement or mirage? *CoRR. abs/2001.00219*. <http://arxiv.org/abs/2001.00219>.
- Czarnecki, C., & Dietze, C. (2017). Domain-specific reference modeling in the telecommu- nications industry. In A. Maedche, J. Vom Brocke, & A. R. Hevner (Eds.), *Designing the digital transformation - 12th international conference, DESRIST 2017, karlsruhe, Germany, may 30 - June 1, 2017, proceedings*, 10243, 313–329. Springer. <https://doi.org/10.1007/978-3-319-59144-5/19>
- Davenport, T., & Harris, J. (2017). *Competing on analytics: The new science of winning*. Harvard Business Press.
- Debus, J. (2007). Extending data mining methodologies to encompass organizational factors. *Systems Research and Behavioral Science: The Official Journal of the International Federation for Systems Research*, 24(2), 183–190. <https://doi.org/10.1002/sres.823>
- Diaz, D., Theodoulidis, B., & Abioye, E. (2012). Cross-border challenges in financial markets monitoring and surveillance: A case study of customer-driven service value networks. *2012 annual srii global conference* (pp. 146–157).
- Dignum, V. (2017). Responsible autonomy. In C. Sierra (Ed.), *Proceedings of the twenty- sixth international joint conference on artificial intelligence, IJCAI 2017, melbourne, Australia, August 19-25, 2017* (pp. 4698–4704). ijcai.org. <https://doi.org/10.24963/ijcai.2017/655>
- Dignum, V. Ethics in artificial intelligence: Introduction to the special issue. (2018). *Ethics and Information Technology*, 20 (1), 1–3. <https://doi.org/10.1007/s10676-018-9450-z>
- Dumas, M., Rosa, M. L., Mendling, J., & Reijers, H. A. (2018). *Fundamentals of business process management, second edition*. Springer. <https://doi.org/10.1007/978-3-662-56509-4>
- Fayyad, U. M., Piatetsky-Shapiro, G., & Smyth, P. (1996a). From data mining to knowledge discovery in databases. *AI Magazine*, 17(3), 37–54. <https://doi.org/10.1609/aimag.v17i3.1230>
- Fayyad, U. M., Piatetsky-Shapiro, G., & Smyth, P. The KDD process for extracting useful knowledge from volumes of data. (1996b). *Communications of the ACM*, 39(11), 27–34. <https://doi.org/10.1145/240455.240464>
- Fayyad, U. M., Piatetsky-Shapiro, G., & Smyth, P. (1996c). Knowledge discovery and data mining: Towards a unifying framework. *Proceedings of the second international conference on knowledge discovery and data mining (kdd-96)* (pp. 82–88). portland, oregon, USA. Association for the Advancement of Artificial Intelligence (AAAI). <http://www.aaai.org/Library/KDD/1996/kdd96-014.php>
- Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., Choudhary, S., Hamilton, E. P., & Roth, D. (2019). A comparative study of fairness-enhancing interventions in machine learning. *Proceedings of the conference on fairness, accountability, and transparency, fat* 2019, Atlanta, ga, usa, January 29-31, 2019* (pp. 329–338). ACM. <https://doi.org/10.1145/3287560.3287589>
- Galup, S. D., Dattero, R., & Quan, J. What do agile, lean, and ITIL mean to DevOps? (2020). *Communications of the ACM*, 63(10), 48–53. <https://doi.org/10.1145/3372114>

- Gertosio, C., & Dussauchoy, A. Knowledge discovery from industrial databases. (2004). *Journal of Intelligent Manufacturing*, 15(1), 29–37. <https://doi.org/10.1023/B/%3AJIMS.0000010073.54241.e7>
- Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37 (2), 337–355. <http://misq.org/positioning-and-presenting-design-science-research-for-maximum-impact.html>
- Haes, S. D., Grembergen, W. V., & Debreceeny, R. S. (2013). COBIT 5 and enterprise governance of information technology: Building blocks and research opportunities. *Journal Information Systems*, 27(1), 307–324. <https://doi.org/10.2308/isis-50422>
- Hajian, S., Bonchi, F., & Castillo, C. (2016). Algorithmic bias: From discrimination discovery to fairness-aware data mining. In B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, & R. Rastogi (Eds.), *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and datamining, San Francisco, ca, usa, August 13-17, 2016* (pp. 2125–2126). ACM. <https://doi.org/10.1145/2939672.2945386>
- Hamada, K., Ishikawa, F., Masuda, S., Myojin, T., Nishi, Y., Ogawa, H., et al. (2020a). Guidelines for quality assurance of machine learning-based artificial intelligence. In R. Garcia-Castro (Ed.), *The 32nd international conference on software engineering and knowledge engineering, SEKE 2020, KSIR virtual conference center, usa, July 9-19, 2020* (pp. 335–341). KSI Research Inc. <https://doi.org/10.18293/SEKE2020-094>
- Hamada, K., Ishikawa, F., Masuda, S., Myojin, T., Nishi, Y., and Ogawa, H., et al. (2020b). Guidelines for quality assurance of machine learning-based artificial intelligence. *International Journal of Software Engineering and Knowledge Engineering, Special Issue: Best Paper from SEKE 2020*, 11 (12), 1589–1606. doi:10.18293/SEKE2020-094.
- Hang, Y., & Fong, S. (2009). A framework of business intelligence-driven data mining for e-business. *2009 fifth international joint conference on inc, ims and idc* (pp. 1964–1970).
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
- Hevner, A., & Chatterjee, S. (2010). Design science research in information systems. *Design research in information systems* (pp. 9–22). Boston, MA: Springer.
- Huber, S., Wiemer, H., Schneider, D., & Ihlenfeldt, S. (2019). Dmme: Data mining methodology for engineering applications – A holistic extension to the crisp-dm model. *Procedia CIRP*, 79, 403–408. Elsevier. <https://doi.org/10.1016/j.procir.2019.02.106>
- Hummer, W., Muthusamy, V., Rausch, T., Dube, P., Maghraoui, K. E., Murthi, A., et al. (2019). Modelops: Cloud-based lifecycle management for reliable and trusted AI. *IEEE international conference on cloud engineering, IC2E 2019, Prague, Czech Republic, June 24-7, 2019* (pp. 113–120). IEEE. <https://doi.org/10.1109/IC2E.2019.00025>
- Kashyap, S., & Iveroth, E. (2021). Transparency and accountability influences of regulation on risk control: The case of a Swedish bank. *Journal of Management and Governance* 25(2), 475–508. <https://doi.org/10.1007/s10997-020-09550-w>
- Kurgan, L. A., & Musilek, P. (2006). A survey of knowledge discovery and data mining process models. *The Knowledge Engineering Review*, 21(1), 1–24. <https://doi.org/10.1017/S0269888906000737>
- Kurshan, E., Shen, H., & Chen, J. (2020). Towards self-regulating AI: Challenges and opportunities of AI model governance in financial services. *Proceedings of the First ACM International Conference on AI in Finance* (p. 8).
- Kusner, M. J., Loftus, J. R., Russell, C., & Silva, R. (2017). Counterfactual fairness. In I. Guyon et al. (Eds.), *Advances in neural information processing systems 30: Annual conference on neural information processing systems 2017, December 4-9, 2017* (pp. 4066–4076). CA, USA: Long Beach. <https://proceedings.neurips.cc/paper/2017/hash/a486cd07e4ac3d270571622f4f316ec5-Abstract.html>
- Marban, O., Mariscal, G., Menasalvas, E., & Segovia, J. (2007). An engineering approach to data mining projects. In *International conference on intelligent data engineering and automated learning 4881*, 578–588. doi:10.1007/978-3-540-77226-2/59.
- Marban, O., Mariscal, G., & Segovia, J. (2009). A data mining and knowledge discovery process model. In P. Julio & K. Adem (Eds.), *Data mining and knowledge discovery in real life applications* (pp. 438–453). Paris, I-Tech.
- Marban, O., Segovia, J., Menasalvas, E., & Fernandez-Baizan, C. (2009). Toward data mining engineering: A software engineering approach. *Information Systems*, 34(1), 87–107. <https://doi.org/10.1016/j.is.2008.04.003>
- Mariscal, G., Marban, O., & Fernandez, C. (2010). A survey of data mining and knowledge discovery process models and methodologies. *The Knowledge Engineering Review*, 25(2), 137–166. <https://doi.org/10.1017/S0269888910000032>
- Mendes, R., & Vilela, J. P. (2017). Privacy-preserving data mining: Methods, metrics, and applications. *IEEE Access*, 5, 10562–10582. <https://doi.org/10.1109/ACCESS.2017.2706947>.
- Meth, H., Müller, B., & Maedche, A. (2015). Designing a requirement mining system. *Journal of the Association for Information Systems*, 16 (9), 2. <http://aisel.aisnet.org/jais/vol16/iss9/2>
- Montiel, J., Bifet, A., & Abdessalem, T. (2017). Predicting over-indebtedness on batch and streaming data. In J. Nie et al. (Eds.), *2017 IEEE international conference on big data, bigdata 2017, Boston, ma, usa, December 11-14, 2017* (pp. 1504–1513). IEEE Computer Society. <https://doi.org/10.1109/BigData.2017.8258084>
- Morabito, V. (2016). *The future of digital business innovation*. Springer Books.
- Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2020). From what to how: An initial review of publicly available AI ethics tools, methods and research to translate principles into practices. *Science and Engineering Ethics*, 26 (4), 2141–2168. <https://doi.org/10.1007/s11948-019-00165-5>
- Niaksu, O. (2015). Crisp data mining methodology extension for medical domain. *Baltic Journal of Modern Computing*, 3(2), 92.
- Ostrowski, L., & Helfert, M. (2011). Commonality in various design science methodologies. In M. Ganzha, L. A. Maciaszek, & M. Paprzycki (Eds.), *Federated conference on computer science and information systems - fedcsis 2011, szczecin, Poland, 18-21 September 2011, proceedings* (pp. 317–320). <http://ieeexplore.ieee.org/document/6078172/>

- Palmi'e, M., Wincent, J., Parida, V., & Caglar, U. (2020). The evolution of the financial technology ecosystem: An introduction and agenda for future research on disruptive innovations in ecosystems. *Technological Forecasting and Social Change*, 151(151), 119779. <https://doi.org/10.1016/j.techfore.2019.119779>
- Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2008). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24 (3), 45–77. <http://www.jmis-web.org/articles/765>
- Peffers, K., Tuunanen, T., & Niehaves, B. Design science research genres: Introduction to the special issue on exemplars and criteria for applicable design science research. (2018). *European Journal of Information Systems*, 27(2), 129–139. <https://doi.org/10.1080/0960085X.2018.1458066>
- Pivk, A., Vasilecas, O., Kalibatiene, D., & Rupnik, R. (2013). On approach for the implementation of data mining to business process optimisation in commercial companies. *Technological and Economic Development of Economy*, 19 (2), 237–256. <https://doi.org/10.3846/20294913.2013.796501>
- Plotnikova, V., Dumas, M., & Milani, F. (2019b). Data mining methodologies in the banking domain: A systematic literature review. In M. Pankowska & K. Sandkuhl (Eds.), *Perspectives in business informatics research - 18th international conference, BIR 2019, katowice, poland*, (Vol. 365, pp. 104–118). Springer. Retrieved from <https://doi.org/10.1007/978-3-030-31143-8/8>
- Plotnikova, V., Dumas, M., & Milani, F. (2020). Adaptations of data mining methodologies: A systematic literature review. *PeerJ Comput. Sci.*, 6, e267. Retrieved from. <https://doi.org/10.7717/peerj-cs.267>
- Plotnikova, V., Dumas, M., & Milani, F. (2021). Adapting the CRISP-DM data mining process: A case study in the financial services domain. In S. S. Cherfi, A. Perini, & S. Nurcan (Eds.), *Research challenges in information science - 15th international conference, RCIS 2021, limassol, cyprus*, (Vol. 415, pp. 55–71). Springer. <https://doi.org/10.1007/978-3-030-75018-3/4>
- Priebe, T., & Markus, S. (2015). Business information modeling: A methodology for data-intensive projects, data science and big data governance. *2015 IEEE international conference on big data, big data 2015, Santa Clara, ca, usa, October 29 - November 1, 2015* (pp. 2056–2065). IEEE Computer Society. <https://doi.org/10.1109/BigData.2015.7363987>
- Rahman, F. A., Desa, M. I., & Wibowo, A. (2011). A review of kdd-data mining framework and its application in logistics and transportation. In *The 7th international conference on networked computing and advanced information management* (pp. 175–180).
- Richards, J. T., Piorkowski, D., Hind, M., Houde, S., & Mojsilovic, A. (2020). A methodology for creating AI factsheets. *CoRR, abs/2006.13796*. <https://arxiv.org/abs/2006.13796>
- Runeson, P., Höst, M., Rainer, A., & Regnell, B. (2012). *Case study research in software engineering – Guidelines and examples*. Wiley. <http://eu.wiley.com/WileyCDA/WileyTitle/productCd-1118104358.html>
- Saldana, J. (2015). *The coding manual for qualitative researchers*. Sage publications.
- SAS Institute, I. (2017). *Sas® enterprise miner™ 14.3: Reference help*. SAS Institute Inc.
- Siau, K., & Wang, W. (2020). Artificial intelligence (AI) ethics: Ethics of AI and ethical AI. *Journal of Database Management*, 31(2), 74–87. <https://doi.org/10.4018/JDM.2020040105>
- Solarte, J. (2002). *A proposed data mining methodology and its application to industrial engineering* [Master's thesis]. University of Tennessee.
- Sommerville, I. Integrated requirements engineering: A tutorial. (2005). *IEEE Software*, 22(1), 16–23. <https://doi.org/10.1109/MS.2005.13>
- Studer, S., Bui, T. B., Drescher, C., Hanuschkin, A., Winkler, L., and Peters, S., et al. (2020). Towards CRISP-ML(Q): A machine learning process model with quality assurance methodology, 3(2), 392–413. *CoRR, abs/2003.05155*. <https://arxiv.org/abs/2003.05155>
- Tavares, R., Vieira, R., & Pedro, L. (2017). A preliminary proposal of a conceptual educational data mining framework for science education: Scientific competences development and self-regulated learning. *2017 international symposium on computers in education (siie)* (pp. 1–6).
- Vakkuri, V., Kemell, K., & Abrahamsson, P. (2019). AI ethics in industry: A research framework. In M. M. Rantanen & J. Koskinen (Eds.), *Proceedings of the third seminar on technology ethics, tethics 2019, turku, finland, October 23-24, 2019*, 2505, 49–60. CEUR-WS.org. <http://ceur-ws.org/Vol-2505/paper06.pdf>
- Venable, J. (2006). The role of theory and theorising in design science research. *Proceedings of the 1st international conference on design science in information systems and technology (desrist 2006)* (pp. 1–18).
- Venable, J. R., Pries-Heje, J., & Baskerville, R. L. (2016). FEDS: A framework for evaluation in design science research open. *European Journal of Information Systems*, 25(1), 77–89. <https://doi.org/10.1057/ejis.2014.36>
- Venable, J. R., Pries-Heje, J., & Baskerville, R. L. (2017). Choosing a design science research methodology. *Australasian conference on information systems 2017*, hobart, Australia. ACIS2017 Conference Proceeding.
- Wieggers, K., & Beatty, J. (2013). *Software requirements*. Pearson Education.
- Xiang, L. (2009). Context-aware data mining methodology for supply chain finance cooperative systems. *2009 Fifth International Conference on Autonomic and Autonomous Systems* (pp. 301–306). <https://doi.org/10.1109/ICAS.2009.48>.