

A way of mapping the informational needs of stakeholders across financial institutions to the process model artefacts

By Ivan Mikheev

Supervised By Dr. F. A. Bukhsh and Dr. R. Bemthuis

Table of contents

Abstract.....	2
1 Introduction and Problem Statement.....	2
1.2 Research Objectives and Questions.....	2
1.1 Research Approach.....	3
1.3 Thesis Structure.....	3
2 Background Research.....	4
2.1 Process Mining.....	4
2.2 Cognitive Load and Visualization in Stakeholder Communication.....	5
2.3 Stakeholder Roles in Financial Institutions and Information Needs.....	5
2.4 Governance and Structural Models for Stakeholder-Role Differentiation.....	5
2.5 Conclusion.....	7
3 Methodology.....	9
3.1 Process Modelling with CRISP-DM.....	9
3.2 Evaluation Rubric Development.....	11
4 Evaluation.....	19
4.1 Model-Level Evaluation Tables.....	19
4.2 Stakeholder–Artefact Suitability Matrix.....	20
4.3 Illustrative Figures.....	20
4.4 Evaluation Summary.....	21
5 Discussion and Limitations.....	21
5.1 Key Findings for Research Questions.....	21
5.2 Comparison with Prior Studies.....	21
5.3 Practical Implications.....	21
5.4 Methodological Reflection.....	21
5.5 Study Limitations.....	21
5.6 Future Research & Industrial Application.....	22
6 Conclusion.....	22
Appendix.....	24
References.....	26

Abstract

Process mining visualizations are typically produced, yet managers at different hierarchical corporate levels often require distinct views of the same process tied to their functions, which can potentially hinder effective decision-making. This study proposes a question-driven evaluation rubric that links data artefacts obtained with process mining to the information needs of stakeholder groups defined by Mintzberg's corporate model. Following the CRISP-DM life cycle, the 2017 BPI Challenge loan application log was modeled, resulting in five artefacts comprising three discovery algorithms (Alpha, Heuristic, Inductive) rendered in two model notations (Petri net, BPMN) plus one diagnostic view (performance-annotated DFG). Parameter thresholds were selected iteratively to avoid empty or trivial nets. Each model was scored against the rubric, resulting in a contingency table that records whether a visualization supports decision-making tasks. Results show that the performance DFG uniquely satisfies strategic-apex, middle-line, and operating-core requirements. In contrast, models generated by Heuristic and Inductive Miners fulfil technostructure needs. Evaluated artefacts left support staff's needs unmet. The study contributes an open repository of scripts to allow further use.

1 Introduction and Problem Statement

Process mining has emerged as a discipline that enables organizations to automatically derive models of processes from event logs and uncover the intricacies of their operational workflows (van der Aalst, 2016). This approach is particularly valuable in industries where operational processes are complex and multi-layered, such as financial institutions. Various stakeholders, from top-level executives to operational managers, interact with these models differently: while a CEO may need a high-level view to grasp overarching trends, a regional or departmental manager typically requires more detailed insights to address specific operational issues.

The challenges, however, lie in the universal presentation of process models and the unfamiliarity of specific viewers with different process notations. Most process mining tools, such as those in the PM4Py library, provide default settings that do not cater to the diverse informational needs of these stakeholder groups (van der Aalst, 2016; PM4Py, n.d.). As noted by Caire et al. (2013), an overly complex model can impose a significant cognitive load and

hinder effective communication, whereas an overly simplified model may omit critical details necessary for informed operational decision-making. This misalignment complicates stakeholder engagement and risks suboptimal decision quality.

Thus, the core focus of this research is the lack of a systematic approach that matches process artefacts to different corporate roles based on informational needs. This study aims to bridge this gap by investigating how various process mining algorithms and the resulting artefacts can be matched to the stakeholder groups within financial institutions. This study focuses on uncovering the perception of visual patterns. It therefore uses the term "artefact" to refer to both the visual outputs of the algorithms and the renditions (DFG, Petri Net, and BPMN) to facilitate visual comparison. In doing so, this research seeks to develop a practical mapping framework that ensures each stakeholder receives the appropriate level of detail for effective decision-making.

1.2 Research Objectives and Questions

The primary objective of this research is to determine which process mining artefacts best align with the needs of specific stakeholder groups in a financial institution. To achieve this, the study aims to:

1. Generate artefacts:
 - Discovery algorithms: Alpha, Heuristic, Inductive.
 - Models & views examined: Petri-net model, BPMN model, performance DFG.
2. Develop a question-driven evaluation rubric that links each stakeholder decision to an explicit visual criterion.
3. Pilot the rubric on a sample of graphs.
4. Produce a table that will indicate the best artefact for the information needs of each organisational tier.

Based on these objectives, the research is guided by the main question and its subquestions:

Main research question:

Which process model artefacts best align with the needs of specific stakeholder groups?

Sub-questions:

- What stakeholder groups exist within the corporate hierarchy of a financial institution?
- How can the complexity of process models be measured?
- How can the operational goals and management levels of stakeholders be translated into the corresponding model complexity?

1.1 Research Approach

This study employs a literature-based evaluation of stakeholder needs and process artefacts, informed by these needs. The CRISP-DM framework was selected to structure the findings into a narrative, as it can be applied to various domains and tools, and has been utilized in process mining research (Chapman et al., 2000; Bemthuis et al., 2025). After business understanding and log preparation, multiple artefacts are generated. The resulting visualisations are assessed with a question-driven rubric derived from the stakeholder decision

analysis.




Figure 1: The CRISP-DM Development Cycle (Jensen, 2012)

To quantify this, complexity metrics such as the number of nodes, arcs, and control-flow complexity will be considered (Cardoso, 2006). In parallel, the study will build upon academic literature to define the information needs of stakeholder groups within a corporate financial structure. Specifically, different models are compared, Mintzberg's five-part corporate structure model (Mintzberg, 1979), the IIA Three Lines Model (Institute of Internal Auditors [IIA], 2020), COSO ERM Framework (COSO, 2017), Basel Principles for Banks (Basel Committee on Banking Supervision [BCBS], 2015), and G20/OECD Principles of Corporate Governance (OECD, 2023) to select the one which separation principle matches the information needs separation. The resulting analysis will map the needs to the artefacts, forming a table that recommends artefacts for different management roles.

1.3 Thesis Structure

This study delves into a literature review that synthesizes existing academic work on process mining, model complexity metrics, and stakeholder communication theories. Research proceeds using the BPI Challenge 2017 dataset as a basis for generating the

process artefacts. Insights from the analysis are synthesized into the stakeholder-artefact table. Finally, this thesis reflects upon the findings and limitations, discussing practical implications and future work.

2 Background Research

2.1 Process Mining

Process mining is designed to extract models from event logs, mapping the execution of processes within organizations or comparing them to the intended design (van der Aalst, 2016). An open-source library, such as PM4Py, implements a range of algorithms that differ in their treatment of noise, complexity, and the granularity of the output (van der Aalst, 2016).

2.1.1 Alpha Miner

Alpha Miner is one of the earliest process discovery algorithms. Based on the theory of Petri nets, it identifies causal relations between events in a log to construct a process model. While being simple, it is limited by being highly sensitive to noise and may struggle with complex or incomplete logs. Therefore, it is used for straightforward mapping of sequential dependencies (van der Aalst, 2016).

2.1.2 Heuristic Miner

Heuristic Miner addresses Alpha Miner's limitations by introducing frequency-based and dependency measures. It filters out infrequent behaviors and noise, focusing the model on dominant process flows. However, its reliance on threshold parameters may lead to oversimplification, potentially missing less frequent but essential process paths (Weijters, van Dongen, & van der Aalst, 2006).

2.1.3 Inductive Miner

Inductive Miner uses a recursive approach to process discovery. Decomposing event logs into smaller components, block-structured process models that are

easier to understand are constructed. The algorithm handles noisy data and complex control-flow patterns more effectively than earlier methods (Leemans, Fahland, & van der Aalst, 2013).

2.1.4 Inductive Miner with BPMN View

Because a stable Fuzzy-Miner implementation has become unavailable in the latest version of the selected library, this study adopts the Inductive Miner's BPMN export as a structurally diverse alternative. When rendered in BPMN notation, the resulting model preserves the block structure while providing a graphical language that is more familiar to business analysts and managers (Dijkman, Dumas, & Ouyang, 2008).

2.1.5 Direct-Follows Graph

A Direct-Follows Graph (DFG) is a frequency-annotated, activity-to-activity network that illustrates how often one event directly follows another in the log (van der Aalst, 2012). Unlike discovery algorithms that produce a Petri net, the DFG makes no structural assumptions about the underlying process.

2.1.6 Challenges of Process Mining

Two challenges emerge from the application of the discovery process mining. First, handling noise and incomplete data. As described by van der Aalst (2016), despite enhancements in algorithm design, noisy datasets can still obscure process patterns, leading to overly complex models or the loss of critical nuances. Second, a fundamental trade-off exists between the complexity of the process model and its comprehensibility (Dumas, M., *et. al.*, 2018). While detailed models capture more process variations, they impose a higher cognitive load on stakeholders, potentially hindering effective decision-making. Conversely, simplified models, while easier to understand, may omit essential information.

Financial institutions inherently operate in environments characterized by multifaceted and interdependent processes (Franzoi *et al.*, 2025). Managing concurrent operations, stringent regulatory requirements, and

disparate information systems is a complex task. These factors contribute to highly detailed and dynamic process structures that require careful analysis, making it challenging to design a one-size-fits-all process artefact. Instead, balancing oversight with operational details calls for an adaptive process mining framework that can be fine-tuned to produce models with varying levels of abstraction and complexity.

2.2 Cognitive Load and Visualization in Stakeholder Communication

Visualizations act as the interface between complex data and interpretation. Cognitive load theory posits that the human cognitive system is limited in its ability to process information. Excessively detailed or poorly structured visualizations can overwhelm users, leading to errors in comprehension and decision-making (Sweller, 1988). Empirical research in process model visualization highlights the adverse effects of high complexity in visuals. Caire et al. (2013) demonstrated that as the complexity of process models increases, the perceiver's cognitive burden grows significantly, impairing the ability to extract meaningful insights. Similarly, Reijers and Mendling (2011) found that the modularization and simplification of process models contribute to improved comprehension and usability.

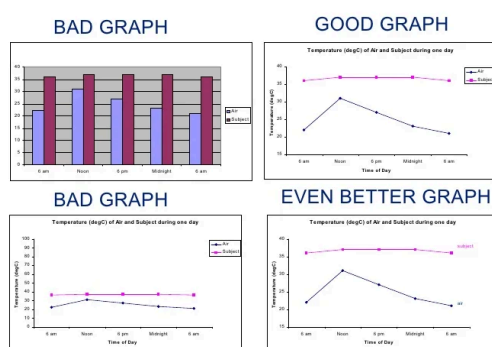


Figure 2: A representation of detail vs simplicity. Kukrety, V. (2010) .

Techniques that reduce visual clutter, such as grouping related activities, clear flow

indicators, and layered visualizations, can significantly enhance the comprehensibility of process models (Caire et al., 2013; Reijers & Mendling, 2011).

2.3 Stakeholder Roles in Financial Institutions and Information Needs

Financial institutions cater to a diverse range of stakeholders with varying roles and information needs (Kandel et al., 2012). At the apex of the organizational hierarchy, executives such as CEOs and senior management are primarily concerned with strategic insights, risk management, and long-term planning. They focus on high-level summaries that synthesize complex operational data into digestible visuals, emphasizing trends and overarching performance indicators. (Mintzberg, 1979). In contrast, mid-level managers and departmental leaders require granular representations of processes to identify inefficiencies, monitor compliance, and optimize resource allocation. The detailed process artefacts demanded by this group often incorporate more nodes and connections, which, while informative, can also introduce visual complexity and increase cognitive load (Caire et al., 2013). Operational staff require more detailed representations to perform day-to-day tasks and address process deviations.

The divergence in the information needs is a gap in process mining. While existing approaches often produce a single, uniform process model, the varying levels of detail required by different stakeholders might be omitted.

2.4 Governance and Structural Models for Stakeholder-Role Differentiation

2.4.1 Rationale for Including Governance Frameworks

It is helpful to draw on well-established governance and organizational-structure models to classify stakeholder roles and justify differentiated information requirements. Five frameworks were selected: the IIA Three Lines Model (Institute

of Internal Auditors [IIA], 2020), G20/OECD Principles of Corporate Governance (OECD, 2023), COSO Enterprise Risk-Management Framework (COSO, 2017), Basel Committee Corporate-Governance Principles for Banks (Basel Committee on Banking Supervision [BCBS], 2015), and Mintzberg's five-part organizational structure (Mintzberg, 1979). These models emphasize the separation of concerns and clarity of information flow; each is viewed from a distinct disciplinary angle. This review provides theoretical grounding for process models to meet stakeholder needs (See table 1)

2.4.2 Thematic Synthesis: Actor-Separation as a Measure for Information Safety

Across the frameworks, stakeholder separation is:

- Horizontal segregation (IIA, Basel) protects the integrity of assurance activities.

- Vertical segregation (OECD) separates oversight from execution.
- Integrated segregation (COSO) embeds risk data in routine decision-making.
- Mintzberg's model offers structural differentiation: each organizational part is defined by its dominant function, naturally producing discrete layers of information abstraction.

While all five frameworks reinforce role clarity, Mintzberg's 5 part model is the most directly applicable foundation. Mintzberg's segmentation is seen in IIA, COSO, and Basel, which segment it into additional audiences (by risk, compliance, and audit). While the other frameworks supply additional governance details, Mintzberg models the structure of the corporate organization in optimal detail to meet the goal of this research.

Framework	Orientation & Core Focus	Segregation Logic	Immediate Relevance to Information Granularity
IIA Three Lines Model (Institute of Internal Auditors [IIA], 2020)	Assurance & risk governance	Horizontal separation of first-line management, second-line oversight, and an independent internal audit line	Ensures risk information escalates without distortion; highlights the need for progressively abstract reporting layers
G20/OECD Principles of Corporate Governance (OECD, 2023)	Capital-market integrity & board accountability	Vertical separation: Independent directors and specialised committees oversee management	Calls for concise, reliable board-level reports while safeguarding detailed operational disclosures
COSO ERM Framework (COSO, 2017)	Enterprise-wide risk integration	Five interlocking components embed risk thinking across all functions	Stresses consistent risk data and layered reporting from operations to strategy

Basel Committee Principles for Banks (BCBS, 2015)	Sector-specific governance & prudential control	Combines three-line logic with the statutory independence of CRO, compliance, and audit	Requires granular control information for assurance functions and high-level dashboards for boards
Mintzberg's Corporate Structure (Mintzberg, 1979)	Organizational design & information flow	Functional/structural differentiation into strategic apex, middle line, operating core, technostructure, support staff	Provides an intuitive one-to-one mapping between the managerial tier and the level of detail needed in process models

Table 1: Overview of the Five Frameworks

2.5 Conclusion

This literature review highlights the dimensions of process mining and its applications within financial institutions. First, an analysis of process mining algorithms illustrates that different process artefacts yield varying levels of model complexity. Each technique presents trade-offs between detail and comprehensibility, influencing the extent to which the generated models support stakeholder-specific decision-making.

Secondly, empirical studies on cognitive load and visualization emphasize the need for designing process models that are cognitively accessible to diverse user groups (Caire et al., 2013; Reijers & Mendling, 2011). Furthermore, investigations into stakeholder roles in financial institutions reveal that the requirements for process model detail vary significantly across the organizational hierarchy, from high-level strategic overviews for executives to detailed operational maps for frontline managers (Mintzberg, 1979).

Literature Matrix

Author(s) & Year	Title / Source	Main Themes	Methodology	Key Findings	Relevance to Research
van der Aalst (2016)	<i>Process Mining: Data Science in Action</i> (2nd ed)	Process mining fundamentals, process discovery algorithms	Theoretical analysis; empirical case studies	Outlines core process mining concepts and various discovery algorithms	Provides foundational knowledge and context for process mining in complex environments
Cardoso (2006)	Process control-flow complexity metrics	Process model complexity metrics	Quantitative analysis; metric derivation	Defines metrics such as node and arc counts that capture structural complexity	Offers quantitative measures for comparing and evaluating model outputs

Caire et al. (2013)	Visual notation design 2.0: Towards user-comprehensible requirements engineering notations	Cognitive load, visual complexity, process model design	Empirical research, case studies	Demonstrates that increased visual complexity elevates cognitive load, impeding comprehension	Informs design guidelines for creating visualizations that reduce cognitive burden for diverse stakeholders
Reijers & Mendling (2011)	Modularity in process models: Review and effects	Process model modularity, visualization, and cognitive effects	Literature review; theoretical analysis	Advocates for model modularization to enhance understandability and reduce complexity	Highlights the balance between retaining necessary details and clarity for various user groups
Mintzberg (1979)	<i>The Structuring of Organizations: A Synthesis of the Research</i>	Organizational structure, stakeholder roles, and information needs	Theoretical framework; empirical observations	Discusses differentiated information needs across hierarchical layers in organizations	Provides insights on how varied stakeholder roles influence process model requirements in financial institutions
Weijters et al. (2006)	Process mining with the heuristic miner algorithm	Heuristic Miner, process discovery, noise handling	Algorithm development; experimental evaluation	Introduces frequency-based dependency measures that filter noise in process discovery	Contributes to understanding parameter settings and trade-offs inherent to heuristic approaches
Leemans et al. (2013)	Discovering block-structured process models from event logs: A constructive approach	Inductive Miner; block-structured models; noise management	Constructive algorithmic approach; empirical validation	Presents a recursive method for decomposing complex processes into structured models that handle noise effectively	Provides an advanced methodological approach for generating understandable models from complex data
Sweller (1988)	Cognitive load during problem solving: Effects on learning	Cognitive load theory: information processing limitations	Experimental cognitive psychology	Demonstrates how limited working memory capacity affects the processing of complex information	Follows the argument for designing simplified, audience-specific visualizations to mitigate cognitive overload
Leemans et al. (2013); Dijkman et al. (2008)	Inductive Miner & BPMN readability studies	Block-structured discovery; business-oriented notation	Algorithmic development; user-comprehension experiments	Inductive Miner: how BPMN improves understandability for non-experts	Provide the BPMN view that supports stakeholder groups with low Petri-net literacy

van der Aalst (2012); Berti et al. (2020)	DFG concept & PM4Py performance view	Frequency-based visualization; bottleneck detection	Conceptual exposition; software implementation	DFG highlights dominant direct-follows relations; performance overlay adds cycle-time insight	Provides a strategic-level view
---	--------------------------------------	---	--	---	---------------------------------

3 Methodology

3.1 Process Modelling with CRISP-DM

3.1.1 Data Processing Framework

CRISP-DM (Cross-Industry Standard Process for Data Mining) remains the most widely adopted framework for transforming raw data into deployable artifacts (Chapman et al., 2000). The framework consists of six iterative phases: Business Understanding, Data Understanding, Data Preparation, Modeling, Evaluation, and Deployment. Each maps onto the tasks required to transform an event log into the deployable product, which in this study is a process visualisation. Table 2 outlines the implementation of each phase in this research.

3.1.2 Business Context of the Dataset

The study relies on the BPI Challenge 2017 event log, an annual benchmark released by the Eindhoven Process Mining Group for reproducible experimentation BPI Challenge,

2017). The 2017 log captures the loan-application process of a Dutch financial institution: customers submit online or in-branch applications, the bank performs a background investigation, and accepted applications proceed to the offer, signing, and activation stages. Each trace records activity name, lifecycle transition, timestamp, and resource, making the log suitable for control-flow discovery and performance insight. Because the BPI Challenge is curated and versioned, the dataset provides a stable reference point for academic comparison while reflecting real-world regulatory constraints typical of European retail banking. For this thesis, the central goal is to generate process visualisations that can later be evaluated for their information potential to different managerial audiences. Detailed evaluation criteria are introduced in subsequent chapters; the present section focuses on how those visualisations are produced.

CRISP-DM Phase	Actions
Business Understanding	Identify the commercial process, the data artefacts to be compared, and the evaluation objectives (detailed in later chapters).
Data Understanding	Inspect the XML/XES log to verify the integrity of the case ID, timestamp, and activity fields.
Data Preparation	Unpack and clean the log, drop duplicates and null fields; compute supporting performance metrics.

Modelling	Generate process models.
Evaluation	Assess the resulting visualisations.
Deployment	Publish the work for industrial or academic reuse.

Table 2: Actions for each Phase in CRISP-DM

3.1.3 Data Preparation via Python Scripts

The public repository consists of a raw compressed dataset (BPI Challenge 2017.xes.gz), an `init.py` cleaning script, and one script per data artefact (Griboed007, 2025).

Running python `init.py` performs three operations:

- 1) Unpacks the zipped XES log to `/data/raw`.
- 2) Cleans duplicates and null attributes, exporting a trimmed XES file to a dedicated folder.

During modeling, the `.xes` dataset is converted into a log object. This object is not stored between algorithms; each modelling script re-imports the `.xes` file to isolate parameters and avoid cross-script artefacts.

3.1.4 Modelling: Generating the Five Process Artefacts

Direct-Follows Graph (DFG).

Using PM4Py's `dfg_discovery()` with the performance variant, a network is built where edges are weighted by average throughput time (See Figure 4 in the Appendix). Edge colouring (blue \rightarrow low, red \rightarrow high) exposes cycle-time dispersion, which is valuable for strategic and tactical review. A throughput-based colouring threshold of the 90th percentile distinguishes dominant bottlenecks without over-saturating the graph.

Alpha Miner (Petri-net).

The `alpha_miner.apply()` call produces a structurally clean net without tunable parameters. This baseline model offers crisp control-flow semantics but no performance overlays (See Figure 3 in the Appendix).

Heuristic Miner (Petri-net).

Due to the final script, iterative tests of dependency were completed. Noise thresholds varied from 0.5 to 0.9. Threshold 0.8 produced the richest actionable model: lower values (<0.7) yielded sparsely connected or empty nets, whereas higher values (>0.9) collapsed multiple alternative paths into single edges, adding no interpretive value (Weijters, van Dongen, & van der Aalst, 2006). The positive-observations threshold was fixed at 20 to filter noise without deleting infrequent yet valid paths (See Figure 5 in the Appendix).

Inductive Miner (Petri-net).

Following Leemans, Fahland, and van der Aalst (2013), a noise threshold of 0.20 was employed to strike a balance between completeness and structure. At < 0.10 , the miner produced excessive self-loops; at > 0.30 , variants disappeared. The resulting block-structured net allows recursive decomposition in later analyses (See Figure 7 in the Appendix).

Inductive Miner (BPMN export).

The Petri-net tree from the previous step is converted to BPMN via `pm4py.convert_to_bpmn()`. The concurrency threshold of 0.20 was retained to match the PN view. The BPMN notation provides a

business-friendly representation (See Figure 6 in the Appendix).

3.1.5 Evaluation and Deployment

Each representative image feeds into the evaluation rubric developed in later chapters. Threshold choices were set at the lowest level, resulting in a clear and actionable diagram. Lower thresholds generated near-empty outputs, while higher thresholds neither altered the structure nor improved interpretability. Finally, the cleaned dataset, scripts, and generated images are published under an MIT license in the GitHub repository for academic replication and industrial adaptation.

3.2 Evaluation Rubric Development

3.2.1 Mintzberg's Model

Mintzberg's framework decomposes a formal organisation into five mutually dependent functional parts, each characterised by a distinctive locus of authority, decision horizon, and information clearance. Applying this to a universal banking context yields the stakeholder catalogue, serving as the reference grid for all subsequent complexity-mapping activities.

Strategic Apex.

The strategic apex comprises actors who hold ultimate decision-making rights over institutional direction and risk management. This stratum typically includes the chief executive officer, the chief financial officer, the board of directors, and any board-level risk or audit committees in an economic setting. Their decisions are infrequent but far-reaching, including capital allocation, merger approval, and enterprise-wide risk tolerance, and therefore depend on highly aggregated, trend-oriented information that can be absorbed rapidly during periodic board meetings (Mintzberg, 1979).

Middle Line.

The middle line occupying the conduit between corporate strategy and day-to-day execution is populated by regional or divisional heads, line-of-business directors

(e.g., Head of Retail Banking, Head of Treasury), and large-branch managers. These actors translate strategic intent into departmental objectives and monitor performance against budgets and compliance targets. Their information needs are more granular than those of the apex: monthly or weekly dashboards that reveal exception volumes, service-level breaches, and control-break trends across their span of control.

Operating Core.

The operating core consists of employees who execute the institution's primary transaction-processing tasks: loan officers, tellers, trade-settlement analysts, and back-office clerks. Their perspective is strictly case-level and real-time: it requires thorough visibility into queue lengths, case status, and rule-based exception flags to resolve customer requests or transaction breaks. Process representations at this level must preserve full activity granularity and temporal ordering to be actionable (van der Aalst, 2016).

Technostructure.

Mintzberg designates the technostructure as the cohort of specialists who design, analyse, and standardise work processes but do not execute them directly. In universal banking, this includes quantitative risk models, business process engineers, IT architects, and regulatory reporting analysts. Their analytic scope spans the entire enterprise, yet their focus is technical; hence, they require both macro-level overviews for pattern detection and drill-down capabilities to inspect control-flow anomalies or data-quality issues (COSO, 2017).

Support Staff.

Finally, support staff deliver services that enable the primary value chain: human resources, legal, facilities, and corporate communications personnel. Although peripheral to core revenue processes, these units interact with operational workflows by handling escalations (e.g., employment checks, legal holds). They benefit from summary-level models that highlight interdepartmental touchpoints and workload

distribution, rather than transactional details (OECD, 2023).

Following Mintzberg's structural analysis, this five-part catalogue establishes clear boundaries for stakeholder-specific information requirements. In the following chapters, each part will be assigned quantitative complexity bands and algorithmic tuning rules to calibrate process mining artefacts to the audience's decision horizon and cognitive capacity.

3.2.2 Decision-Task and Data-Type Audit for Mintzberg Stakeholder Groups

To translate Mintzberg's functional parts into complexity requirements, the decision tasks performed by each group and the data attributes that best support those tasks were identified through a focused literature review drawing on management information systems (MIS) and banking operations (Laudon, K. C., & Laudon, J. P., 2021). It was found that the key dimensions are information granularity, time horizon, update cadence, content focus (KPI vs. transactional detail), and emphasis on exceptions or compliance. These attributes for each stakeholder group were synthesized into an Information-Attribute Matrix (See Table 3).

Strategic Apex. Research on executive information systems (Anthony, 1965; Simon, 1977) indicates that board-level actors prefer highly aggregated, externally oriented indicators that summarize performance over multi-quarter horizons. In banking, Haldane and Madouros (2012) stress the need for

dashboards that condense risk exposures and capital ratios into a handful of trend lines updated monthly or quarterly.

Middle Line. Divisional and regional heads translate strategy into tactical targets that require monthly drill-downs to compare budgeted versus actual volumes, service-level breaches, and control-break frequencies (Davenport & Harris, 2007). Their preferred granularity is intermediate: sub-process aggregation with the ability to isolate problematic branches or product lines.

Operating Core. Process execution staff work with real-time transaction-level data (van der Aalst, 2016). Studies of loan-origination desks and payment-operations centers (Beath & Orlikowski, 1994) highlight the need for case-level detail, queue length, and immediate exception flag attributes that demand the most detailed control-flow visualizations.

Technostructure. Risk modellers and process engineers switch between macro pattern detection and micro root-cause analysis. Literature indicates a need for enterprise-wide coverage, historical depth, and the ability to zoom in on anomalous variants for compliance testing. Accordingly, they require dual-resolution data and models.

Support Staff. HR and legal units consume cross-functional workload summaries, case handoffs, and policy exceptions on a weekly or monthly basis (OECD, 2023). They rarely require transaction-level traces except when handling escalations.

Mintzberg Part	Typical Decision Tasks	Granularity	Time Horizon & Cadence	Content Focus	Exception / Compliance Emphasis
Strategic Apex	Capital allocation, risk accommodation, long-range strategy	Very low (end-to-end process KPIs)	Multi-quarter horizon; monthly/quarterly updates	Aggregate KPIs, external risk indicators	Low: only aggregated breaches
Middle Line	Budget control, service-level tracking, and tactical compliance	Moderate (sub-process clusters)	Monthly/fortnightly	Departmental KPIs + exception counts	Medium: focus on recurring breaches
Operating Core	Case processing, real-time exception resolution	High (activity-level traces)	Intraday/real-time	Transaction details, queue metrics	High: immediate flagging
Technostructure	Risk modelling, process redesign, root-cause analysis	Dual (aggregate+drill-down)	Historical & on-demand	Control-flow patterns, variance stats	Very high compliance testing
Support Staff	HR workload, legal escalations, facilities planning	Low-to-moderate	Weekly/monthly	Cross-functional summaries	Low escalations only

Table 3 Information-Attribute Matrix

3.2.3 Implications for Complexity Mapping

Table 2 illustrates a monotonic relationship between the managerial tier and the

required details: complexity tolerance increases sharply from the strategic apex to the operating core, then changes for the technostructure, which demands both panoramic and microscopic views. These

attributes foster the metric thresholds adopted later when parameterising process mining artefacts, aiming to match each stakeholder with a model whose structural complexity aligns with its information-processing capacity.

3.2.4 Stakeholder Decision Inventory

This section aims to identify a single, high-impact decision or observation summarising each stakeholder group's information needs (See Table 4). These decisions will later be translated into five evaluation questions determining whether a given process mining graph fits.

Evaluation Framework for Decision Tasks

Three criteria guided the selection of one principal decision per stakeholder group. First, organisational impact: the decision must have a material influence on performance or risk. Second, frequency: it should recur often enough to warrant routine visual support. Third, data dependency must rely on process-generated information that a mining model can reveal (Simon, 1977). Where multiple candidate decisions met these criteria, the one most frequently cited across sources was chosen.

Strategic Apex Decision Analysis

In universal banking, the strategic apex board members, chief executive officer, and senior C-suite executives periodically set or revise the institution's risk policy and capital allocation (Haldane & Madouros, 2012). This decision is based on understanding end-to-end process efficiency and systemic bottlenecks that may threaten earnings or compliance. Consequently, apex stakeholders require high-level visualizations that surface cycle-time variance and macro bottleneck locations rather than granular task flows.

Middle-Line Decision Analysis

Regional presidents and line-of-business directors are responsible for allocating staff and budget to meet service-level agreements (SLAs) and revenue targets (Davenport & Harris, 2007). Decision support requires moderately granular views that reveal workload distribution, departmental handoffs, and emerging SLA breaches on a monthly or daily basis.

Operating-Core Decision Analysis

Loan officers, payment clerks, and other front-office or back-office staff comprise the operating core. Their critical decision is the immediate clearance of process exceptions—for example, resolving a failed know-your-customer (KYC) check or a payment mismatch (van der Aalst, 2016). They need transaction-level traces with real-time exception flags and queue-length indicators; any abstraction that suppresses these details undermines task execution.

Technostructure Decision Analysis

Risk modellers, process engineers, and IT architects' prime task is diagnosing low-frequency control-flow anomalies to inform root-cause analysis and regulatory submissions. They oscillate between birds-eye pattern detection and microscopic drill-down, necessitating dual-resolution models that retain seldom-seen variants while conveying overall structure.

Support-Staff Decision Analysis

Support units, including human resources, legal, and facilities, make workload-balancing and escalation decisions that depend on where and how often cross-functional handoffs occur (OECD, 2023 [25]). Weekly or monthly summaries of hand-off counts and escalation nodes suffice; overly detailed traces would add cognitive noise without improving decision quality.

Mintzberg Part	Core Decision ("X")	Primary Data Requirements
Strategic Apex	Set or adjust strategic risk policy	End-to-end cycle-time variance, macro bottlenecks, quarterly trend lines
Middle Line	Allocate resources to meet SLAs	Workload by department, SLA breach counts, monthly variance
Operating Core	Clear real-time process exceptions	Transaction-level traces, live exception flags, and queue length
Technostructure	Diagnose low-frequency anomalies for compliance	Variant frequencies, conformance gaps, drill-down traces
Support Staff	Monitor cross-functional hand-offs and escalations	Weekly hand-off counts, escalation nodes, workload summaries

Table 4: Synthesis of Stakeholder Decisions

3.2.5 Evaluation-Question Development

Implications for Evaluation-Question Formulation

Each decision distilled above will be reframed as an evaluation prompt: "Does this graph provide the information required to decide on X?" Five prompts will constitute the evaluation rubric applied to every process mining visualisation. A "yes" answer signifies that the model supplies sufficient insight for the stakeholder's core decision, enabling a direct, question-driven mapping.

Conceptual Principles for Question Design

The formulation of five evaluation questions is grounded in two complementary theories. Cognitive-fit theory (Vessey, 1991) posits that information representation is most effective when its format aligns with the mental operations required by a task. Task-technology fit theory extends this idea

by arguing that a tool's utility is maximized when its capabilities align with the information demands of the user's decision environment (Goodhue & Thompson, 1995). Applying these principles to process mining visualisations implies that a graph should be judged not by aesthetic criteria, but by its capacity to convey the specific data elements that a stakeholder needs to make a decision. Hence, a yes/no question, coupled with that decision, provides a test of cognitive and technological fit.

Transformation of Stakeholder Decisions into Question Prompts

The previous section identified one high-impact decision for each functional part of Mintzberg's structure. To transform decisions into evaluation prompts, each statement was recast in the interrogative form "Does the graph provide the information required to decide on X?". The predicate "provide the information" signals a sufficiency test rather than a presence test,

while “decide on X” preserves explicit linkage to the stakeholder’s task. Linguistic ambiguity was minimised by specifying the key informational element: cycle-time variance, SLA breaches, live exceptions, low-frequency variants, or hand-off counts, wherever feasible, inside the question.

Draft Question Statements

Q1 (Strategic Apex) Does this graph display end-to-end cycle-time variance and the dominant process bottlenecks needed to set or adjust strategic risk posture?

Q2 (Middle Line) Does this graph show departmental workload distribution and service-level breach counts required to allocate resources effectively?

Q3 (Operating Core) Does this graph surface live queue length and explicit exception-flag paths necessary to clear real-time process exceptions?

Q4 (Technostructure) Does this graph reveal low-frequency control-flow variants and conformance gaps needed for root-cause diagnosis and compliance analysis?

Q5 (Support Staff) Does this graph highlight cross-functional hand-off counts and escalation nodes that allow support units to monitor workload and escalations?

Before applying the rubric, a pilot assessment is planned to verify the clarity and sufficiency of the five questions. Five data artefacts will be scored (See Table 5). The pilot is designed to expose threshold rules that may be overly strict or ambiguous.

3.2.6 Operationalisation of Yes/No Criteria

Mapping Questions to Observable Model Features

To render the yes/no judgements, each question was linked to artefacts detectable in the visual output or the underlying model

statistics. For Q1, the requisite artifacts are a cycle-time distribution (e.g., a box plot or histogram) and a bottleneck heat map that highlights activities with congestion values exceeding the 90th percentile. Q2 requires a frequency view that partitions cases by department, along with an SLA-breach counter exported from delay annotations. Q3 demands explicit exception transition paths marked in red and a queue-length indicator derived from waiting-time metrics. Q4 is satisfied only when a variant list enumerates all paths with a frequency below 1 % and a conformance overlay identifies deviations from the normative model. Finally, Q5 calls for a hand-off matrix or chord diagram showing activity-to-activity transfers between organisational units, with escalation points labelled.

Metric- and Heuristic-Based Thresholds

A “yes” judgment is awarded when the graph meets presence and sufficiency thresholds. Presence means the visual element exists; sufficiency means it meets minimum informational depth. The apex heat map, for example, must display no more than three dominant bottlenecks to satisfy cognitive-fit guidance, as executives tend to cope poorly with numerous focal points (Haldane & Madouros, 2012). For middle-line SLA charts, the breach counter must display numerical values, not colored status, and breaches must be grouped by department. Low-frequency variants for Q4 follow 1 % rarity rule. Hand-off counts for Q5 are considered sufficient if at least 90 % of inter-unit transfers are captured.

Tool Implementation Notes

Observable features are extracted via manual procedures. Cycle-time distributions and queue-length metrics are computed with PM4Py 2.7.0 functions `case_duration` and `waiting_time`. All scripts and parameter files are version-controlled in a public repository by Griboed007. (2025).

Question	Observable Features	Threshold for Yes	Automated / Manual
Q1	Cycle-time histogram; bottleneck heat-map	≤ 3 dominant bottlenecks visualised; histogram present	Presence check: manual
Q2	Departmental workload view; SLA-breach counter	Workload chart partitioned; breach counter numeric	PM4Py workload script (auto); visual check manual
Q3	Exception-flag paths; queue-length bar	Exceptions marked in red; queue metric available; timestamp gap ≥ 2× median	Exception path: manual; queue metric: PM4Py (auto)
Q4	Variant list: conformance overlay	Variants with frequency < 1 % listed; deviation overlay visible	Variant list: PM4Py (auto); conformance: manual
Q5	Hand-off matrix/chord; escalation labels	Matrix captures ≥ 90 % transfers; escalation node labels present	Label check manual

Table 5: Pilot Evaluation Rubric

3.2.7 Pilot Results

The initial scoring produced one ✓ mark out of a possible twenty-five (See Table 6). Most false negatives occurred in Q2 (Middle-Line) and Q4 (Technostructure), where arc-frequency cues and dotted edges were present but not considered sufficient evidence to support the presence of a structure. The initial rubric (Version 1) required explicit KPI overlays, including cycle-time charts for Q1 and SLA counters for Q2, among others.

During the pilot, it was observed that several control-flow-only models nonetheless conveyed stakeholder-relevant cues, such as arc thickness indicating frequency, even though they lacked full KPI instrumentation. Consistent with the cognitive-fit literature (Vessey, 1991), the rubric was revised to accept such cues, provided they met minimal sufficiency thresholds. Table 7 presents the refined rubric (Version 2) used in the full evaluation.

Artefact/ Stakeholder	Strategic Apex	Middle Line	Operating Core	Techno-struct ure	Support Staff
Performance DFG	✓	x	x	x	x
Alpha PN	x	x	x	x	x
Heuristic PN	x	x	x	x	x
Inductive PN	x	x	x	x	x
IM-BPMN	x	x	x	x	x

Table 6: Pilot Results

Threshold Adjustment

Guided by these observations and cognitive-fit arguments (Goodhue & Thompson, 1995; Vessey, 1991), the sufficiency rules for Q2-Q4 were relaxed to accept frequency-weighted arcs. Exception and bottleneck criteria (Q1 & Q3) remained unchanged:

Q1 (Strategic Apex) Does this graph display end-to-end cycle-time dispersion and highlight the dominant process bottlenecks needed to set or adjust strategic risk posture?

Q2 (Middle Line) Does this graph depict major activity-frequency patterns or delay hotspots sufficient to guide divisional

resource allocation and service-level management?

Q3 (Operating Core) Does this graph identify explicit exception or denial paths, enabling operational staff to track and identify real-time process exceptions?

Q4 (Technostructure) Does this graph expose alternative or low-frequency paths, via branching structures or dotted edges, adequate for technostructure analysts to perform root-cause or compliance analysis?

Q5 (Support Staff) Does this graph reveal cross-activity hand-offs or escalation points that support support-staff units in monitoring workload and escalations?

Question	Observable Feature(s) Acceptable for ✓	Minimum Sufficiency Threshold	Detection Mode
Q1	Any performance overlay that visualizes cycle-time dispersion or arc-level throughput color coding/width, plus at least one bottleneck indicator (e.g., the thickest arc, a red edge).	The dispersion element present and ≥ 1 bottleneck is visibly emphasized.	Colour/width inspection (manual).

Q2	Frequency-weighted arcs or node-delay shading that indicates workload concentration or delay hotspot.	At least three activities or paths are visibly contrasted by their thickness or shading.	PM4Py frequency stats (auto) + manual visual check.
Q3	Explicit “exception”, “denied”, “cancelled”, “refused” transitions or red-coloured exception arcs.	≥ 1 dedicated exception path/transition visible.	Label scan (auto) + manual confirmation.
Q4	Branching structures, dotted edges, or variant paths rendered at reduced opacity/line-style or a textual variant list.	At least one alternative path whose frequency < 5 % is displayed (visual or textual).	Heuristic-Miner dotted edges, Inductive branch, or PM4Py variant list (auto).
Q5	Multi-incoming/outgoing gateways or cross-band arcs that imply hand-offs or labelled escalation nodes.	There is at least one hand-off or escalation point that is visually identifiable.	Manual inspection; optional custom script to count cross-band transitions.

Table 7: Revised Evaluation Rubric

3.2.8 Final Evaluation Procedure

The complete evaluation utilizes the Version 2 rubric. For every model, answers to Q1–Q5 were recorded in a structured table that logs ✓/✗.

Validation

From the final score sheet, a contingency table was compiled, crossing artefacts (DFG, Alpha, Heuristic, Inductive-PN, IM-BPMN) with stakeholder suitability (as assessed by five questions). A second pass using the revised rubric yielded 12 ✓ marks, indicating increased coverage (See Table 9).

4 Evaluation

4.1 Model-Level Evaluation Tables

Table 6-1 through Table 6-5 present the question-level scores (✓ = criterion satisfied; ✗ = criterion not satisfied) for each representative model generated in Chapter 5. A brief discussion follows each table to

clarify which stakeholder questions were met.

Question	Q1	Q2	Q3	Q4	Q5
Score	✓	✓	✓	✗	✗

Table 8-1: Q1–Q5 Scores for Performance DFG

The performance-annotated Direct-Follows Graph (DFG) meets the strategic-apex, middle-line, and operating-core criteria due to its colored throughput arcs and explicit refusal path, which lacks variant visibility.

Question	Q1	Q2	Q3	Q4	Q5
Score	✗	✗	✓	✓	✗

Table 8-2: Q1–Q5 Scores for Alpha-Miner Petri Net

Although the Alpha Miner offers no performance layer, it exposes explicit

exception transitions and alternative branches, supporting operating-core and technostructure needs.

Question	Q1	Q2	Q3	Q4	Q5
Score	x	✓	✓	✓	x

Table 8-3: Q1–Q5 Scores for Heuristic-Miner Petri Net

Arc-thickness cues reveal workload delays (Q2), dotted edges expose rare variants (Q4), and labelled exception paths satisfy operating-core requirements (Q3).

Question	Q1	Q2	Q3	Q4	Q5

4.2 Stakeholder–Artefact Suitability Matrix

Artefact/ Stakeholder	Strategic Apex	Middle Line	Operating Core	Techno-struct ure	Support Staff
Performance DFG	✓	✓	✓	x	x
Alpha PN	x	x	✓	✓	x
Heuristic PN	x	✓	✓	✓	x
Inductive PN	x	x	✓	✓	x
IM-BPMN	x	x	✓	✓	x

Table 9: Stakeholder–Artefact matrix.

Only the DFG and Heuristic Miner cover three stakeholder tiers; no visualization satisfies the Support-Staff criterion (hand-offs), revealing a gap for future modeling work.

Score	x	x	✓	✓	x
-------	---	---	---	---	---

Table 8-4: Q1–Q5 Scores for Inductive-Miner Petri Net

The block-structured Inductive net lacks frequency cues but marks exception and low-frequency paths, aligning with Q3 and Q4.

Question	Q1	Q2	Q3	Q4	Q5
Score	x	x	✓	✓	x

Table 8-5: Q1–Q5 Scores for Inductive-Miner BPMN

IM-BPMN replicates Inductive Petri net branching in a business-friendly notation, yet still lacks performance overlays and hand-off indicators.

4.3 Illustrative Figures

Figure 6-1 shows the performance DFG, with red/blue arcs highlighting bottlenecks that satisfy Q1. Figure 6-2 displays the Heuristic-Miner Petri net, where dotted edges and thick arcs illustrate the rare-variant and workload cues used to

answer Q2 and Q4. Both figures demonstrate how visual elements map directly to rubric criteria.

4.4 Evaluation Summary

The evaluation demonstrates that the choice of process artefact influences how well stakeholder information needs are met. The performance-enhanced DFG uniquely supports strategic, tactical, and operational decisions (Q1–Q3). Petri nets generated by Heuristic and Inductive miners are essential for technostructure analysis (Q4). No current model satisfies the support-staff requirement (Q5), indicating the need for future work.

5 Discussion and Limitations

5.1 Key Findings for Research Questions

The primary objective of this study was to identify which process artefacts satisfy the information needs of stakeholder groups. First, the performance-annotated DFG met the information requirements of the strategic apex, middle line, and operating core (Q1–Q3). Second, Petri-net models generated by the Heuristic and Inductive Miners, with the BPMN rendition of the Inductive Miner, consistently exposed low-frequency paths and transitions, thereby addressing technostructural needs in drill-down traces (Q4).

Third, no visualization captured data required to satisfy support staff requirements (Q5).

5.2 Comparison with Prior Studies

The superiority of the performance DFG for senior stakeholders aligns with van der Aalst's (2012) argument that DFGs provide "instant insight" into throughput bottlenecks by avoiding model-theoretic detail. This study verifies the findings of Weijters, van Dongen, and van der Aalst (2006), who demonstrated that frequency-weighted Heuristic nets assist line managers in workload diagnostics. For technostructure roles, results support the recommendation that risk analysts require explicit variant disclosure, dotted edges, and branching

structures in Heuristic and Inductive nets fulfilled without overlays.

5.3 Practical Implications

From a managerial standpoint, the stakeholder-artefact table offers a decision support tool that may be implemented as a dynamic dashboard. From a research perspective, the question-driven rubric offers a modular evaluation template that can be applied to future datasets or extended with new stakeholder questions (e.g., environmental, social, and governance-oriented metrics).

5.4 Methodological Reflection

By broadening sufficiency rules to recognise arc thickness, dotted edges, and labeled exception transitions, the refined rubric (Version 2) achieved higher coverage. This outcome aligns with the Task-Technology Fit theory (Goodhue & Thompson, 1995) by taking into account more generalizable details.

5.5 Study Limitations

The potential dataset bias arises from the fact that the findings rely solely on the BPI Challenge 2017 log. Other financial event logs, especially those containing richer resource attributes, could alter the observed suitability patterns. Regarding the algorithm scope, the BPMN export of the Inductive Miner serves as a substitute for a dedicated Fuzzy-Miner implementation; as a result, models optimized for abstraction may be under-represented in the analysis. The coverage of the parameter grid is limited in scope, with the exploration focusing on a subset of parameter combinations to select the best artefact. However, the entire hyperparameter space is considerably larger, and extreme or edge-case settings might significantly impact the metric values and, consequently, the assessment of suitability. Additionally, a support-staff gap is evident, leaving Question 5 unsatisfied. This shortcoming highlights a construct validity issue, as the current rubric cannot be fully operationalized without specialized analytics

for handoffs. Finally, there is a potential bias introduced by having only one rater conduct the model evaluation, which could lead to inter-rater bias.

5.6 Future Research & Industrial Application

This work demonstrates that the choice of process artefact influences how well different managerial tiers can interpret process visualizations. Nevertheless, it relies on a single event log. For example, mortgage servicing logs, credit-card dispute workflows, or insurance-claim life cycles could reveal whether the suitability patterns persist when case volumes and activities vary. A multi-log comparison could also illuminate whether certain stakeholder groups remain sensitive to the same visual cues (e.g., throughput colouring) or require alternative abstractions under different operational contexts.

While the current evaluation rubric is based on a literature analysis of stakeholder roles, empirical work with human subjects would enhance its diagnostic power. Controlled laboratory studies that challenge representatives of different stakeholder groups with alternating visualizations and record comprehension time, error rates, or eye-tracking data (cf. Caire et al., 2013) could identify which glyphs, colors, or layouts best match the required cognitive load. Survey or interview research could further refine the evaluation questions by asking stakeholders to articulate what information they lack when interpreting a chart. Incorporating such evidence would enable rubric thresholds to move from “literature-defined” to “peer-calibrated,” thereby increasing both validity and acceptance of a study.

Moreover, the gap for support-staff information needs (Q5) persists. Variant-clustering techniques might produce a simplified “handoff heat-map” highlighting escalation loops, offering a visual grammar more suited to HR and legal audiences than traditional BPMN gateways.

Beyond academia, the stakeholder-artefact perception table offers a tool for financial institutions. Consider a bank planning to modernise its internal audit dashboards. The bank could:

1. Profile stakeholder roles (board, regional management, operations centre, model-risk unit, HR/legal) against the information categories in the perception table.
2. Select default views (e.g., performance DFG for quarterly board packs, Heuristic Miner PN for monthly SLA reviews, and inductive BPMN for model-risk deep analysis) supported by comprehension patterns.
3. Embed the scripts in a development pipeline so that freshly extracted event logs automatically generate fit-for-purpose diagrams for each audience.

In this scenario, research on informational needs may help avoid “one-size-fits-all” reporting and reduce the cycle time between data extraction, insight, and action. Over time, usage analytics (e.g., which diagrams are opened most frequently, where users zoom or hover) could feed back into the rubric, creating a virtuous loop between empirical evidence and visual analytics tooling.

6 Conclusion

This thesis was set to determine which process mining artefacts satisfy the information needs of stakeholders situated at different managerial tiers of Mintzberg’s organisational structure. Using the BPI Challenge 2017 event log, five process artefacts: Alpha Miner, Heuristic Miner, Inductive Miner (Petri-net form), and an Inductive-Miner BPMN export, together with a performance-annotated DFG, were generated and evaluated. Assessment relied on a piloted and revised question-driven rubric to recognise visual cues such as arc thickness and dotted edges. The refined rubric provided a yes/no test of whether each visualisation supported five stakeholder decisions (Q1–Q5).

Three headline findings emerged. First, the performance DFG satisfied the strategic-apex, middle-line, and operating-core questions (Q1–Q3), owing to its explicit depiction of cycle-time dispersion and bottleneck arcs (Table 8-1). Second, Petri-net models produced by Heuristic and Inductive Miners with the BPMN rendition consistently exposed rare-variant branches and exception transitions, meeting technostructure requirements (Q4). Third, none of the evaluated models visualised cross-functional hand-offs enough to satisfy the support-staff information criterion (Q5), indicating a persistent gap in outputs. Methodologically, the study extends cognitive-fit research (Vessey, 1991) by demonstrating that minimal visual cues, frequency weighting, and colored exception arcs can substitute for full KPI overlays when the evaluator’s rubric is aligned with stakeholder decision tasks. Therefore, the lightly revised rubric (Version 2) demonstrates a pragmatic balance between theory and tool realism, offering a template for future process mining evaluations.

For practitioners in banking and similar high-regulation environments, the

stakeholder-artefact table serves as a rapid guide: Boards can utilize performance-enhanced DFG dashboards during quarterly risk reviews. Divisional managers can monitor workload more effectively by using frequency-weighted Heuristic nets. Additionally, risk analysts can benefit from inductive models exported in BPMN, especially in cases where Petri net literacy is limited, to assist with variant-compliance diagnostics. The publicly available PM4Py scripts and parameter grids enable organisations to reproduce or extend the analysis with their event logs. This thesis aims to offer a practical foundation for role-sensitive, process-aware decision-making within the corporate hierarchy by exploring how process mining artefacts can be related to stakeholder decision tasks through a question-based approach. However, limitations of a single dataset and the absence of inter-rater validation may introduce bias in the results; therefore, this paper proposes future work to address the methodological drawbacks. Finally, this thesis concludes with the final message: *“Right artefact plus right manager leads to better informed decisions.”*

Appendix

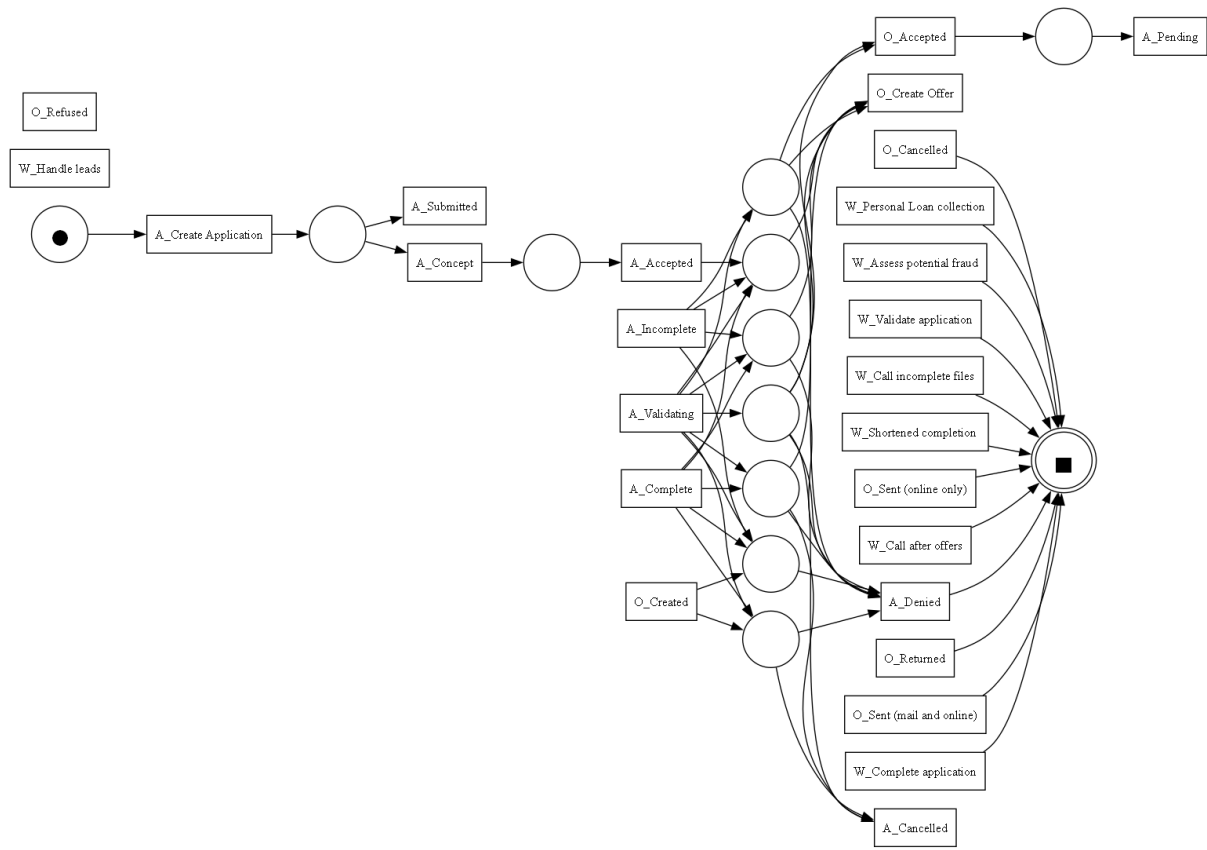


Figure 3: Alpha Miner Petri Net, Full Resolution Image Available at Github Repository (Griboed007, 2025)




Figure 4: Direct-Follows Graph, Full Resolution Image Available at Github Repository (Griboed007, 2025)

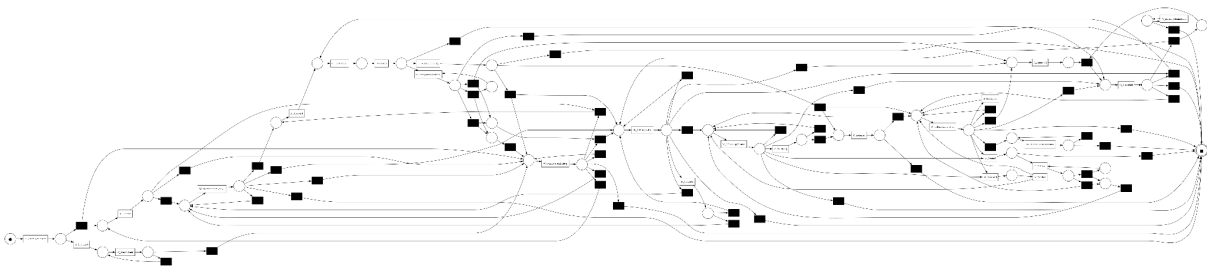


Figure 5: Heuristic Miner Petri Net, Full Resolution Image Available at Github Repository (Griboed007, 2025)




Figure 6: Inductive Miner BPMN, Full Resolution Image Available at Github Repository (Griboed007, 2025)




Figure 7: Inductive Miner Petri Net, Full Resolution Image Available at Github Repository (Griboed007, 2025)

Note on Using Generative AI

During the preparation of this work, the author utilized Grammarly AI to spell-check, reformat sentences without altering the underlying meaning, assess readability, and suggest reference style mistakes. ChatGPT was employed for understanding the dataset structure, troubleshooting scripts, explaining concepts of process mining, and assisting in searching for relevant literature and synonyms for words. After using these tools, the author reviewed and edited the content, reviewed and selected the literature as needed, and takes full responsibility for the content of the work.

References

Anthony, R. N. (1965). *Planning and control systems: A framework for analysis*. Harvard Business School. Retrieved from: <https://archive.org/details/planningcontrols0000anth>

Basel Committee on Banking Supervision. (2015). *Corporate governance principles for banks*. Bank for International Settlements. Retrieved from: <https://www.bis.org/bcbs/publ/d328.pdf>

Beath, C., & Orlikowski, W. (1994). *The contradictory structure of systems-development methodologies: Deconstructing the IS–user relationship in information engineering*. *Information Systems Research*, 5(4), 350–377. Retrieved from: <https://doi.org/10.1287/isre.5.4.350>

Bemthuis, R. H., Govers, R. R., & Asadi, A. (2025). A CRISP-DM-based methodology for assessing agent-based simulation models using process mining. *Journal of Simulation*, 1–22. Retrieved from: <https://doi.org/10.1080/17477778.2025.2508245>

Berti, A., van Zelst, S. J., & van der Aalst, W. M. P. (2019). *Process Mining for Python (PM4Py): Bridging the Gap Between Process- and Data Science*. arXiv. Retrieved from: <https://doi.org/10.48550/arXiv.1905.06169>

- BPI Challenge. (2017). *BPI Challenge 2017 dataset*. Retrieved from: <https://doi.org/10.4121/uuid:3926db30-f712-4394-aebc-75976070e91f>
- Caire, P., Genon, N., Heymans, P., & Moody, D. L. (2013). *Visual notation design 2.0: Towards user comprehensible requirements engineering notations*. In 2013, 21st IEEE International Requirements Engineering Conference (RE) (pp. 115–124). IEEE. Retrieved from: <https://doi.org/10.1109/RE.2013.6636711>
- Cardoso, J. (2006). *Process control-flow complexity metric: An empirical validation*. In Proceedings of the IEEE International Conference on Services Computing (SCC '06) (pp. 167–173). IEEE. Retrieved from: <https://doi.org/10.1109/SCC.2006.82>
- Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., & Wirth, R. (2000). *CRISP-DM 1.0: Step-by-step data mining guide*. The CRISP-DM Consortium. Retrieved from: <https://www.kde.cs.uni-kassel.de/wp-content/uploads/lehre/ws2012-13/kdd/files/CRISPWP-0800.pdf>
- Committee of Sponsoring Organizations of the Treadway Commission. (2017). *Enterprise risk management—Integrating with strategy and performance*. COSO. Retrieved from: https://www.academia.edu/42188027/Enterprise_Risk_Management_Integrating_with_Strategy_and_Performance
- Davenport, T. H., & Harris, J. G. (2007). *Competing on analytics: The new science of winning*. Harvard Business Press. Retrieved from: <https://books.google.nl/books?id=n7Gp7Q84hcsC&printsec=frontcover&hl=ru#v=onepage&q&f=false>
- Dijkman, R., Dumas, M., & Ouyang, C. (2008). *Semantics and analysis of business-process models in BPMN*. *Information and Software Technology*, 50(12), 1281–1294. Retrieved from: <https://doi.org/10.1016/j.infsof.2008.02.006>
- Dumas, M., La Rosa, M., Mendling, J., & Reijers, H. A. (2018). *Fundamentals of Business Process Management* (2nd ed.). Springer. 189–193 & 440–442. Retrieved from: <https://link.springer.com/book/10.1007/978-3-662-56509-4>
- Franzoi, S., Hartl, S., Grisold, T., van der Aa, H., Mendling, J., & vom Brocke, J. (2025). Explaining process dynamics: A Process Mining Context Taxonomy for sense-making. *Process Science*, 2(2), Article 2. Retrieved from: <https://doi.org/10.1007/s44311-025-00008-6>.
- Goodhue, D. L., & Thompson, R. L. (1995). *Task-technology fit and individual performance*. *MIS Quarterly*, 19(2), 213–236. Retrieved from: <https://doi.org/10.2307/249689>
- Griboed007. (2025). *PMPerceptionThesis* [Source code]. GitHub. Retrieved from: <https://github.com/Griboed007/PMPerceptionThesis/tree/master>
- Haldane, A. G., & Madouros, V. (2012, August 31). *The dog and the frisbee*. *Speech presented at the Federal Reserve Bank of Kansas City's Economic Policy Symposium: "The Changing Policy Landscape," Jackson Hole, WY*. Bank for International Settlements. Retrieved from: <https://www.bis.org/review/r120905a.pdf>
- The Institute of Internal Auditors. (2024). *The IIA's Three Lines Model: An update of the Three Lines of Defense* (September 2024 update). Retrieved from:

<https://www.theiia.org/globalassets/documents/resources/the-iias-three-lines-model-an-update-of-the-three-lines-of-defense-july-2020/three-lines-model-updated-english.pdf>

Jensen, K. (2012, April 26). *CRISP-DM process diagram* [Image]. Wikimedia Commons. Retrieved from: https://commons.wikimedia.org/wiki/File:CRISP-DM_Process_Diagram.png

Kandel, S., Paepcke, A., Hellerstein, J. M., & Heer, J. (2012). Enterprise data analysis and visualization: An interview study. *IEEE Transactions on Visualization and Computer Graphics*, 18(12), 2917–2926. Retrieved from: <https://doi.org/10.1109/TVCG.2012.219>

Kukrety, V. (2010). *Data confusion: How to confuse yourself and others with data analysis* [Slide set]. SlideShare. Retrieved from: <https://image.slidesharecdn.com/dataconfusionhowtoconfuseyourselfandotherswithdataanalysis-101013115912-phpapp02/95/data-confusion-how-to-confuse-yourself-and-others-with-data-analysis-9-728.jpg>

Laudon, K. C., & Laudon, J. P. (2021). *Management information systems: Managing the digital firm (17th ed.)*. Pearson. Retrieved from: https://repository.dinus.ac.id/docs/ajar/Kenneth_C._Laudon,Jane_P._Laudon_-_Management_Information_System_12th_Edition_.pdf

Leemans, S. J. J., Fahland, D., & van der Aalst, W. M. P. (2013). *Discovering block-structured process models from event logs – A constructive approach*. In J. M. Colom & J. Desel (Eds.), *Application and Theory of Petri Nets and Concurrency: PETRI NETS 2013* (Lecture Notes in Computer Science, Vol. 7927, pp. 311–329). Springer. Retrieved from: https://doi.org/10.1007/978-3-642-38697-8_17

Mintzberg, H. (1979). *The structuring of organizations: A synthesis of the research*. Englewood Cliffs, NJ: Prentice-Hall. Retrieved from: <https://archive.org/details/structuringoforg0000mint>

Organisation for Economic Co-operation and Development. (2023). *G20/OECD principles of corporate governance 2023*. OECD Publishing. Retrieved from: https://www.oecd.org/en/publications/g20-oecd-principles-of-corporate-governance-2023_ed750b30-en.html

PM4Py. (n.d.). *PM4Py: Process mining for Python*. Retrieved from: <https://pm4py.fit.fraunhofer.de/>

Reijers, H. A., & Mendling, J. (2011). *Modularity in Process Models: A Review and Its Effects*. *Business & Information Systems Engineering*, 3(6), 381–391. Retrieved from: http://dx.doi.org/10.1007/978-3-540-85758-7_5

Simon, H. A. (1977). *The new science of management decision* (Rev. ed.). Prentice-Hall. Retrieved from: <https://archive.org/details/newscienceofmana0000simo/page/n5/mode/2up>

Sweller, J. (1988). *Cognitive load during problem solving: Effects on learning*. *Cognitive Science*, 12(2), 257–285. Retrieved from: https://doi.org/10.1207/s15516709cog1202_4

van der Aalst, W. M. P. (2012). *Process mining—Making the knowledge discovery process centric*. *ACM SIGKDD Explorations Newsletter*, 13(2), 45–49. Retrieved from: https://www.kdd.org/exploration_files/V13-02-07-vanderAalst.pdf

van der Aalst, W. M. P. (2016). *Process mining: Data science in action* (2nd ed.). Springer. Retrieved from: <https://doi.org/10.1007/978-3-662-49851-4>

Vessey, I. (1991). *Cognitive fit: A theory-based analysis of the graphs versus tables literature*. *Decision Sciences*, 22(2), 219–240. Retrieved from: <https://doi.org/10.1111/j.1540-5915.1991.tb00344.x>

Weijters, A. J. M. M., van Dongen, B. F., & van der Aalst, W. M. P. (2006). *Process mining with the Heuristic Miner algorithm*. In *Proceedings of the International Conference on Business Process Management* (pp. 66–78). Retrieved from: https://www.researchgate.net/publication/229124308_Process_Mining_with_the_Heuristics_Miner-algorithm