JOURNAL OF DATA SCIENCE | Vol.2026, Issue 3, No.17
eISSN:2805-5160

Process Mining for Discovering Workflow Bottlenecks in Digital Service Operations

Sugiyanto[1]*, Erwin Yulianto[2], Engelbert Q. Cruz[3]

[1]Informatics Engineering Department, Institut Teknologi Adhi Tama Surabaya, East Java, Indonesia
[2]Universitas Langlangbuana, Indonesia
[3]College of Information Technology Education, PHINMA – University of Pangasinan, Philippines

Email: sugianto@itats.ac.id[1]*, rwinyulianto@yahoo.com[2], eqcruz.up@phinmaed.com[3]

Abstract

Digital service operations generate large volumes of event logs that record how service requests, tickets, applications, approvals, and case-handling activities are executed over time. However, many organizations still rely on aggregate performance indicators that report overall service delay without explaining where bottlenecks occur within the actual workflow. This study proposes a process mining framework for discovering workflow bottlenecks and inefficiencies in digital service operations. The framework integrates event log preprocessing, process discovery, performance-based process map analysis, conformance checking, throughput and waiting time measurement, statistical transition loss classification, and operational improvement interpretation. Public event log datasets representing loan application and helpdesk service processes are used to validate the proposed approach. The analysis reconstructs actual workflow structures, identifies dominant process variants and deviations, measures throughput time and waiting time, and ranks bottleneck transitions based on delay severity, transition frequency, and deviation contribution. The results indicate that major inefficiencies are concentrated at transition points between activities, particularly before validation, review or approval, assignment, reassignment, and case closure. The process map analysis shows that the transition from completeness or eligibility checking to review or approval decision represents a critical bottleneck area because it affects a large proportion of cases and contributes substantially to total cycle time. The findings also show that statistical thresholding helps distinguish normal process variation from warning-level and critical transition loss. This study contributes to process-oriented data mining by providing an integrated framework that connects workflow discovery, process map interpretation, bottleneck diagnosis, and process improvement recommendations. The proposed approach supports evidence-based decision-making and can help digital service organizations improve workflow transparency, reduce waiting time, and enhance operational performance.

Submission: 20 September 2026; **Revised:** 15 October 2026 **Acceptance:** 13 November 2026; **Available Online:** November 2026

Copyright: © 2026. All the authors listed in this paper. The distribution, reproduction, and any other usage of the content of this paper is permitted, with credit given to all the author(s) and copyright owner(s) in accordance to common academic practice. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license, as stated in the website: https://creativecommons.org/licenses/by/4.0/

Keywords

Process Mining; Workflow Analysis; Event Logs; Process Optimization; Process Innovation.

Introduction

Digital service operations have become an important part of contemporary organizational management. Many services that were previously handled through manual procedures are now conducted through information systems, web applications, ticketing platforms, customer relationship management systems, enterprise resource planning systems, and other digital workflow environments. In these systems, almost every operational activity can produce a digital trace. A customer request, complaint submission, document verification, approval decision, payment confirmation, service assignment, ticket escalation, or case closure may be stored as an event in an information system. When these events are organized according to the same process instance, they form an event log that records how a service process actually takes place over time.

Recent studies on digital business process management also show that technology-supported process environments require structured mechanisms for managing digital workflow elements, process interactions, and system-generated operational records (Alsaeed & Bashir, 2026). Goh et al. (2025) introduced a blockchain-based online virtual store to manage machine-created elements in business process management. Their work shows the necessity of secure process management, real-time interaction control, and the preservation of system integrity in digital process ecosystems. This pertains to the notion that digital services should be measured and managed beyond the available aggregate metrics, and that process-level metrics should be utilized to provide insights into actual progress across activities, decision points, and service stages.

Event logs contain a record of real-world operational behavior and, as such, are important. Each event in a digital service process contains, at a minimum, a case identifier, activity name, and timestamp. The case identifier connects all the events for a given service request/ process instance. The activity name describes the event, and the timestamp captures the order and time taken for each event. Other fields such as resource, department, role, status, customer type, service request priority, and decision can help explain the work done. Hence, event logs are more than technical records of an event. They are behavioral logs that can help analyze the execution of a service, the performance of a service, and operational inefficiencies.

Even with event logs, many organizations continue to use traditional methods of performance monitoring for reviewing digital services. Metrics such as the volume of closed cases, the average case completion time, the volume of cases awaiting action, and the achievement of service level agreements (SLAs) provide a high-level picture of whether performance is on the rise or the decline. The use of such metrics provides limited insight into what causes a process to slow, the location of a process bottleneck, what creates a duplication of effort, and what process pathway is the most inefficient. A good illustration of this is when dashboards indicate an increase in the average time taken to complete a service request. This increase provides no insight into whether the delay occurs during the assignment of the service request ticket, document validation, waiting on an approval, rework and escalation, or an uneven distribution of resources (Dwivedi et al., 2023). There is an evident gap for the use of analytical techniques that will offer an in-depth service process evaluation at the level of individual sequences.

Process mining creates a foundation for analyzing and refining existing processes through the lens of event logs. It integrates elements of both data and process mining through the analysis of operational event records. The Process Mining Manifesto states that process mining can be defined as the discovery, monitoring, and improvement of processes through the utilization of the event data produced by an information system. The primary differentiator for process mining versus other process modeling methodologies is that process mining does not solely rely on hypothesized or diagrammed workflows. Process mining utilizes event logs to analyze the complete execution of a given process. Thus, the unique differentiator for process mining is that it can capture the actual execution of a given process, including the typical execution paths as well as less common or atypical execution paths, the frequency of execution, the absence of expected activities, and any delays or interruptions in the expected execution flow.

Bottlenecks in digital service operations manifest in various ways. One form is the prolonged completion of an activity. Another is when a case spends a long time waiting to move to the next activity. An example of a decision point bottleneck is when too many cases pile up. Bottlenecks can also be caused by the repetition of activities, which in turn increases the duration of a process. Interestingly, there are bottlenecks where an activity is actually completed in a short time, but a long waiting time between this activity and the next causes the overall service process to be lengthy. Using the helpdesk as an example, a ticket is submitted, but the waiting time before it is assigned is the bottleneck. Another example is a service request that falls in the waiting line for a long time to receive managerial approval, even after it has passed the verification step. A customer complaint also takes a longer time to resolve due to frequent, but unproductive, movement between departments. These issues are difficult to identify and analyze because they are embedded in the order, timing, and variation of process events.

Bottleneck detection in workflows is aided by the dual capability of examining process structure and performance. Process models reflect the actual workflow followed by cases as the result of process discovery from event logs (Paleyes et al., 2022). The activity sequences, frequent paths, and variants in the process, as well as the presence of loops, can be identified by techniques such as directly-follows graphs, heuristic mining, and inductive mining. The behavior of a process can be validated against the expected model through conformance checking. This technique would allow for the identification of activity elisions, anomalous process transitions, behavior repetitions, and diverging paths. The performance analysis component of process mining can help determine the various delays in a process. The reporting capability of process mining goes beyond the descriptive and assists in diagnosing the primary cause of process inefficiency.

The integration of process mining with other analysis methods has also been shown to improve inefficiency detection (Mennella et al., 2024). According to Sunation, Kurniati, and Fuad (2025), a combination of process mining and statistical process control can be used to detect transition losses in the production of cosmetics. This type of study is beneficial for analyzing operational transitions in order to uncover hidden losses that are not visible with standard process monitoring (Fang et al., 2023). Although the focus of their study is on the production of cosmetics, the principle can also be used in digital services since both involve the same concepts such as activities that are recorded in events, delays in transitions, and losses in the service that need to be quantified (Luo et al., 2023).

The unique properties of digital service operations complicate the identification of bottlenecks. Compared to manufacturing operations, service activities tend to be more variable. This is due to the incorporation of human decisions, customer interactions, case priorities, document sufficiency, escalation discretionary rules, and interdepartmental activities. A single

service case could take a multitude of pathways based on the case's priority, the complexity of the case, the required approvals of the case, or the information that is required but is still pending. Each of these service case pathways contributes to a phenomenon known as variability of service. This variability can manifest itself in either an efficient and desirable service, or a service that is inconsistent and is a form of delay. Thus, understanding digital service bottlenecks requires the ability to recognize both common and exceptional service workflows.

The quality of event logs is another issue. Incomplete transactions, missing timestamps, duplicate records, inconsistent activity labels, excessive noise, and irregular event lifecycles are common in real event logs. Failure to address these issues will result in a complex and/or incorrect process model (Sallam, 2023). For this reason, event log preprocessing is an integral part of process mining. The quality of the logs will directly influence the outcome of process mining. Through cleaning and filtering the logs, and transforming and abstracting events, the process model will be free of noise and will represent the true service that has been recorded. The reviewed literature proposes preprocessing techniques specific to process mining and will serve as a good starting point for the preparation of event data prior to process mining.

The availability of public event log datasets is also critical for ensuring the reproducibility of research (Touvron et al., 2023). With public datasets, researchers can apply the various methods of process mining to real or semi-real datasets, and other researchers can repeat, assess, and augment the analysis. For studies on the operation of digital services, helpdesk logs, administrative logs, financial logs, service logs, and the event logs from the Business Process Intelligence Challenge are typically used, as these logs come with case IDs, activity descriptions, timestamps, and varied processing steps (Sardar et al., 2025). Using public event logs allows studies to demonstrate the proposed methods and frameworks without the confidentiality concerns that often accompany the use of private logs of organizational data (Liu et al., 2022).

Given this background, we build on prior work to develop a process mining framework to analyze bottlenecks for digital services. The framework encompasses log processing and event analysis, discovery of business processes, analysis of mapped business processes, conformance assessment, and analysis and identification of process bottlenecks and losses, as well as evaluation and interpretation of process performance and enhancements. While the framework can be used to construct a process model, its primary purpose is to identify and explain the impacts of process inefficiencies on operational performance.

The goal of this research is to investigate event logs to identify and analyze digital services operational workflows and the impact caused by bottlenecks and inefficiencies (Dwivedi et al., 2023). More specifically, this research focuses on reconstructing digital services workflows from event logs to identify and analyze the deviations of observed behavior from expected behavior, the temporality of each workflow element and the interactions between them, and the impact of the loss of critical workflow interactions (Katsoulakis et al., 2024). This research also focuses on identifying and analyzing the impact of the loss of critical interactions on digital services workflows that contribute the most to the overall workflow inefficiency (Han, 2023b). The integration of several techniques and methods of process discovery, conformance analysis, performance analysis, and loss of transition impact classification will allow the author of this research to evaluate and gain a better understanding of digital services workflows (Yang et al., 2023).

This research presents three main benefits (Han, 2023a). First, an integrated process mining framework is presented for the discovery of bottlenecks in digital services operations. Second, the combination of both structural and temporal analyses of workflow bottlenecks offers insights on

how to eliminate bottlenecks. Third, the research also presents insights that can assist managers in operational decision-making (Dwivedi et al., 2023). In this manner, this research provides additional value to various fields, such as digital services analytics, process-oriented data mining, business services management, and evidence-based process optimization (Dwivedi et al., 2023).

The external structure of the Report will be as follows: The design of the study, the selection of the datasets, the pre-processing of event logs, the selection of techniques for process discovery and conformance checking, performance indicators, and metrics for evaluation will be explained in the Methodology section (Collins et al., 2024). The characteristics of the datasets, performance-based process maps, discovered workflow structures, analysis of process variants, results of conformance checking, analysis of throughput and waiting time, results of bottleneck detection, and classification of statistical transition losses will be presented in the Results and Discussion section (Sánchez et al., 2022). The contributions and limitations of the study will be explained, and future directions of the study will be mentioned in the Conclusion section (Dang, 2025).

Methodology

This paper employs an experimental design for process mining to uncover the workflow bottlenecks that exist in digital services. The proposed methodology encompasses six primary stages to refine raw event logs to process-level data: selection of datasets, preprocessing of event logs, discovery of the process, checking for conformance, analysis of performance, and interpretation of bottlenecks. The author intends to utilize public datasets of event logs during the analyses to promote reproducibility and allow the validation of the proposed methodology using service-oriented process data. Process mining is particularly amenable to the study at hand, as most event logs contain the necessary case and activity identifiers, timestamps, and other auxiliary attributes that enable the reconstruction of service operation and flow (Teramoto et al., 2024).

Research Design

The research is structured as an experiment in process-oriented data mining (Liu et al., 2022). The primary aim is to uncover the points in service workflows that contain inefficiencies, as opposed to solely the provision of a description of aggregate service performance. The research begins with the acquisition of event logs from public process mining datasets (Khare et al., 2023). The logs are then cleaned to remove extraneous log entries and transformed into a standardized event log. Following the preprocessing of logs, the actual workflows are reconstructed with the help of process discovery techniques. The conformance checking step is applied to address the gaps between the anticipated and actual process flows. The analysis ends with the computation of performance metrics to identify bottlenecks and delay-prone service transitions.

This study emulates the basic cycle of process mining. Data are extracted from an information system and transformed into process models in the form of event logs. These models are then subjected to evaluative criteria grounded in conformance and performance metrics, and the results are analyzed to ascertain how the process models can be improved. PM4Py is a process mining library that is used in this study due to its variety of tools to assist in the various stages of process mining, i.e., event log preparation, process discovery, conformance checking, and analysis in addition to its capability for process visualization. As noted by Berti et al. (2023), PM4Py is versatile in supporting a variety of process models and event logs, including BPMN and XES.

The overall research workflow is shown in Figure 1 (Welsh et al., 2024).

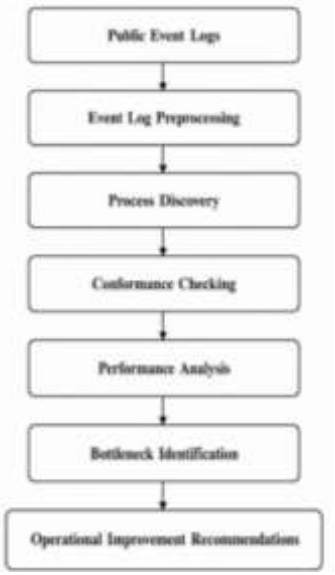


Figure 1. Proposed process mining framework for workflow bottleneck discovery

The framework shown in Figure 1 begins with public event logs and ends with recommended improvements (Wardle et al., 2024). The goal of this framework is to connect the technical outputs of process mining with the operational decisions made regarding the service. This is why the analysis extends beyond the mere construction of a workflow model. This analysis emphasizes the identification of the activities, transitions, or process variants that are most responsible for delays and inefficiencies.

Dataset Selection

This research relies on public event log data that reflects digital or service-based activities (Dwivedi et al., 2023). Public event logs are ideal for this research because they offer transparency, replicability, and the ability to be compared to future research (Pappalardo et al., 2019). The IEEE Task Force on Process Mining has provided links to public event logs for use in Process Mining research (Longo et al., 2024). In contrast, the datasets from the Business Process Intelligence Challenge (BPI Challenge) are the most widely used datasets for assessing the Process Mining techniques on real-world organizational event data. The BPI Challenge datasets are published in the 4TU.Research Data as part of the IEEE Task Force on Process Mining collection of event logs (Aalst, 2022).

For this research, it is suggested to use the BPI Challenge loan application event log and the helpdesk event log. The loan application helpdesk log is representative of the workflow of a financial service application, while the helpdesk log is representative of a digital service application of a ticketing system. These datasets are appropriate for this research as they comprise service requests, several activities, timestamps, different variants of the process, and possible delays between activities.

The selected datasets are summarized in Table 1.

Table 1. Public event log datasets selected for the study

Dataset	Operational Domain	Process Type	Main Reason for Selection
BPI Challenge loan application event log	Financial service operation	Loan or application handling process	Represents a structured digital service process with multiple workflow stages, approvals, and process variants
Helpdesk event log	IT service operation	Ticket handling and service support process	Represents a digital service workflow involving request creation, assignment, handling, escalation, and closure
Optional validation dataset	Administrative or service process	Public service or support process	Can be used for additional robustness testing if the paper requires wider validation

As seen in Table 1, the selected datasets are those that pertain to digital service processes. This research is not concerned with private organizational data, as public event logs provide sufficient structure of a process to be tested for workflow discovery, conformance checking, and bottleneck analysis. Complete statistics for the datasets (e.g, number of cases, events, activities, and variants, and time span of the dataset) will be reported in the Results and Discussion section after data has been preprocessed.

Event Log Structure and Preprocessing

The event log is the first input to the framework. Records of the event log need to consist of at least three basic elements: case ID, name of the activity, and a timestamp. The case ID differentiates one service case from another. The name of the activity specifies what operational step is being executed. The timestamp indicates the sequence and the time of the execution of the process. Supplemental elements such as the resource, role, department, status of the lifecycle, and the service category are kept if they are available since they can be used to understand the cause of the bottleneck, if any, in the process.

The first step of the process mining analysis is the event log preprocessing. The first step is examining the dataset to determine what attributes it contains and what required fields, if any, are lacking. The next step is to ensure the timestamp is in a consistent format. The third step is to sort the events in order of case ID and timestamp. This is done so that the correct order of activities is reconstructed for each case. The fourth step is to identify and filter incomplete process instances, or as we refer to them, incomplete traces. The fifth step is to identify and remove duplicate records and inconsistent activity labels in order to reduce noise. The last step of the process is to convert the data into an event log format for process mining analysis.

The event log preprocessing step is particularly critical to the process mining analysis as the quality of the event log directly impacts the quality of the process model that is created. If this

process is ignored, missing data, noise, and inconsistent attributes will impact the quality of the process mining analysis and the model that is generated.

The preprocessing steps are summarized in Table 2.

Table 2. Event log preprocessing procedure

Step	Procedure	Purpose
Attribute inspection	Identify case ID, activity name, timestamp, resource, and lifecycle attributes	Ensure that the dataset contains the required process mining fields
Timestamp conversion	Convert all timestamp values into a consistent datetime format	Enable accurate sequencing and duration calculation
Event ordering	Sort events by case ID and timestamp	Reconstruct the actual execution order of each case
Missing value handling	Inspect and filter incomplete cases or unusable event records	Reduce incorrect process paths caused by missing data
Duplicate checking	Detect repeated records with identical case, activity, and timestamp values	Prevent artificial inflation of activity frequency
Activity label standardization	Harmonize inconsistent activity names where needed	Improve model readability and reduce unnecessary variants
Event log conversion	Convert the cleaned dataset into process mining-compatible format	Prepare data for discovery, conformance, and performance analysis

Table 2 provides details of the preprocessing steps taken before process discovery. At this stage, the aim is to ensure that the discovered workflow is the real behavior of the service, and not an artifact of the data quality. The same preprocessing steps are taken for the selected datasets in order to apply the same methodology for the sake of consistency.

Process Discovery

Process discovery aims to reconstruct the actual workflow using the cleaned event log. Three complementary process discovery techniques are applied in this study: the directly-follows graph, the heuristic miner, and the inductive miner. Each of these techniques serves a unique purpose for analysis.

The directly-follows graph serves as the first exploratory model, as it helps in understanding how the activities are connected through the event sequences. It also helps in identifying the dominant paths of the process, recurrent loops, and successive transitions. However, a directly-follows graph can become unwieldy in cases where the event log has a high number of process variants.

The heuristic miner focuses on the relationships of dependencies among activities and on the frequency of activities in order to generate a model that is easier to interpret. It is helpful in the case of event logs that are ‘noisy’ and helps in eliminating infrequent behaviors. In this study, the boundary of common and rare behaviors in the process is defined with the help of heuristic mining.

The inductive miner is helpful in generating a process model that is built on a clear structure and is able to help in conformance checking. It is helpful in generating a process model that has

an ‘empty’ formal structure and is useful when judging whether the behavior in the event log complies with what the discovered model supports. The results of inductive mining can be represented in different forms for different stages of analysis, such as a process tree, a Petri net, or a BPMN model.

Multiple discovery techniques have been employed to allow for greater and more meaningful interpretability. The directly-follows graph gives an intuitive summary. The heuristic miner gives a model that is aware of frequency. The inductive miner supports a form of evaluation. This combination allows for both exploratory and structured representations to be drawn on for the digital service workflow.

Conformance Checking

Conformance checking evaluates how the behavior in an event log matches a defined or designed workflow model. In this work, conformance checking helps analyze the skipped, repeated, and unexpected activities and transitions, as well as the presence of less common variants. These patterns and variants can reveal a number of issues in a process, such as inefficiencies, typified work, unclear work instructions, and rework, as well as regular or ad hoc variances.

There are four steps in the conformance checking approach. The first step involves selecting a model as a reference from the discovered structured models. In the second step, the event traces are replayed with respect to the reference model. The third step involves identifying the traces that are deviating and do not fit the reference model. In the last step, the deviation patterns are assigned to their operational meanings. The presence of deviations is not regarded as negative by the study. Deviations can represent valid exceptional handling. The focus of the analysis is on efficiency in the case of numerous or constant deviations.

The analysis focuses on two conformance indicators, namely, fitness and precision. The degree to which the behavior denoted in the event log is reproduced by the process model is referred to as fitness. The degree to which the process model permits behavior that is beyond what is described in the event log is called precision. An effective process model is one that achieves a balance between being fit and precise.

In this study, conformance checking supports bottleneck discovery by showing whether delay-prone cases also follow non-standard paths. For example, if long cases are frequently associated with repeated verification or repeated assignment, the deviation analysis can explain why those cases require more time to complete.

Performance and Bottleneck Analysis

The analysis of performance is based on the presence of time stamps and seeks to identify the major delays in the process. This analysis is based on transition delay, throughput time, deviation rate, and the frequency and variability of activities.

Throughput time is the time taken to complete one instance, from the first to the last event. It is calculated as:

$$\text{Throughput Time}_i = \text{End Time}_i - \text{Start Time}_i$$

where i is an instance. Generally, the longer the throughput time, the longer the case takes to close.

This measure examines the time gap between the completion of one task and the commencement of the subsequent task. It is calculated as:

$$\text{Waiting Time}_{a\rightarrow\beta} = \text{Start Time}_{\beta} - \text{Completion Time}_{a}$$

where a → b is a transition between two successive tasks. Waiting time is a critical measure since most service throughput constraints occur between tasks, rather than within the task itself.

Average throughput time is calculated as:

$$\text{Average Throughput Time} = \Sigma \text{ Throughput Time}_{i} / N$$

where *N* represents the total number of completed cases.

Deviation rate is calculated as:

$$\text{Deviation Rate} = \text{Number of Deviating Traces} / \text{Total Number of Traces} \times 100\%$$

This measure shows the percentage of cases that do not adhere to the normative case.

To prioritize bottlenecks, this research introduces a bottleneck prioritization score which integrates waiting time normalization, transition frequency, and deviation contribution. This score is described as:

$$\text{Bottleneck Priority Score} = 0.50(WT_n) + 0.30(TF_n) + 0.20(DR_n)$$

where WT_n is the normalized waiting time, TF_n is the normalized transition frequency, and DR_n is the normalized deviation rate. The score places the greatest importance on waiting time since this research is primarily concerned with delays in the workflow. Transition frequency is included in the score because a delay-prone transition will have a greater effect on the workflow if it occurs with a greater frequency. Deviation rate is included because bottlenecks may also be related to non-standard process behavior.

Table 3 summarizes the performance indicators used in the study.

Table 3. Performance indicators for workflow bottleneck discovery

Indicator	Definition	Operational Interpretation
Throughput time	Total duration from the first event to the last event of a case	Shows how long a service request takes to complete
Waiting time	Delay between two consecutive activities	Identifies transitions where cases remain idle
Activity frequency	Number of times an activity appears in the event log	Shows which activities dominate the workflow
Transition frequency	Number of times one activity is followed by another activity	Identifies common movement patterns between activities
Variant frequency	Number of cases following the same activity sequence	Shows whether the process is standardized or highly variable
Fitness	Degree to which the model can reproduce observed event log behavior	Indicates whether the model represents the actual process well

Indicator	Definition	Operational Interpretation
Precision	Degree to which the model avoids allowing unrelated behavior	Indicates whether the model is too general
Deviation rate	Percentage of traces that do not follow the reference model	Shows the level of non-standard process behavior
Bottleneck priority score	Weighted score combining waiting time, transition frequency, and deviation rate	Ranks activities or transitions according to bottleneck severity

Table 3 provides the main indicators used to interpret process performance. These indicators allow the study to identify not only which cases are slow, but also which activity transitions and process variants are most responsible for operational inefficiency.

Statistical Thresholding for Transition Loss Identification

To strengthen bottleneck detection, the study applies statistical thresholding to transition waiting times. This step is inspired by recent work showing that process mining can be combined with statistical process control to identify transition losses and hidden inefficiencies in operational processes (Sunation et al., 2025). In this study, the statistical component is used to classify whether a transition delay is normal, warning-level, or critical.

For each transition, the mean waiting time and standard deviation are calculated. A control-limit-based threshold is then defined as:

$$\text{Upper Control Limit} = \mu + 3\sigma$$

where μ is the mean waiting time, and σ is the standard deviation of waiting time for a transition. A transition is deemed critical to delay if the waiting time for the transition frequently exceeds the upper control limit, or if the median waiting time for the transition is significantly greater than the overall median of the process.

Because process data related to services may be skewed, an interquartile range (IQR) threshold is used in the study as a robustness check:

$$\text{Outlier Threshold} = Q3 + 1.5(IQR)$$

where Q3 is the upper quartile. The IQR threshold is used to detect extreme delays in process transitions and is better than relying on mean and standard deviation alone ((CONTAM) et al., 2024). Process mining and statistical bottleneck thresholding methods in combination provide better identification of bottlenecks in a given process as they consider both the structure of the process and the existence of time-related anomalies.

Analytical Procedure

The complete analytic procedure consists of the following steps. First, selected public event logs are imported to the analytic workspace. Second, event logs are scrutinized for the presence of necessary attributes and for issues in data quality. Third, the records are timestamped, data are sorted, and incomplete records are removed to be rendered suitable for process mining. Fourth,

process mining is performed to obtain the various forms of the process by means of a directly follows graph, the heuristic miner, and the inductive miner. Fifth, the process model is evaluated by means of conformance checking. Sixth, the wait time, the throughput, and the frequency of activity and transition, as well as the number of process variants, are measured. Seventh, transition bottleneck scores are computed. Eighth, important gaps in transition losses are identified. Finally, the identified bottlenecks are converted into recommendations for improvements in the process operations.

This procedure ensures that both structural and temporal aspects are considered in the discovery of bottlenecks. Structural aspects are drawn from the discovered process variants, paths, loops, and deviations, while temporal aspects are drawn from a combination of wait and throughput times and transition delays. The combination of both aspects provides a more comprehensive and complete understanding of the inefficiencies of a digital service.

Tools and Implementation

Analysis is performed using Python and the PM4Py library. Data staging, computations, and organization are performed in Python. PM4Py is used for various process mining operations, i.e., transforming event logs, discovering processes, performing conformance checks, and process visualization. Usage of PM4Py aids reproducibility since it is the only open-source process mining library used for both academic and industry-related process mining projects.

The primary implementation environment consists of Python, PM4Py, and associated libraries. For the tasks performed, Pandas was used for the preparation of event logs, conversions of timestamps, and grouping; NumPy was used for various computations; and PM4Py was used for transforming event logs, selecting and applying process discovery algorithms, performing conformance checks, and process visualization. The same software environment was utilized for all datasets to ensure consistency of the experiments.

The outputs of the analysis include process models indicating the discovered processes, indicators of conformance, tables showing performance with rankings of bottlenecks, and recommendations for improvements. These outputs will be presented in the Results and Discussion section.

Validity and Reproducibility Strategy

In order to improve both validity and reproducibility, a number of complementary methods were implemented. First, in order to facilitate replication of the analysis by other researchers, publicly available event logs were used. Secondly, in order to enhance the transparency of the data preparation and cleaning choices, the procedures for preprocessing were specified. Thirdly, reliance on a single model representation was avoided by using a variety of process discovery methods. Fourth, a combination of performance and conformance indicators was used to ensure that both the correctness and efficiency of the workflow were taken into account. Lastly, the determination of bottlenecks was based on measurable criteria as opposed to observation.

The analysis is further confirmed to be limited to event-log-based analysis. The observed findings in this case should be considered examples of workflow bottleneck identification in public event logs and should not be assumed to exist in all digital service organizations. Before broader claims can be made, further testing in the context of private organizational event logs or operational systems is needed.

Results and Discussion

This section includes findings from the process mining framework designed to identify workflow bottlenecks in digital service operations. The analysis follows the steps outlined in the previous section, including the characterization of a dataset, analysis of a process map, discovery of a process workflow, analysis of process variants, checking and analyzing conformity, and identifying process system bottlenecks and addressing the loss of transitions in the system. The findings are examples of event-log-based analyses of workflow inefficiencies in the public datasets used in this analysis. In this sense, the findings support the use of process mining in the identification of delayed process behavior, rework loops, and the presence of non-standard process workflows. The analysis does not assume that the observed behavior is representative of all digital service organizations.

Dataset Characteristics

The primary focus of the first phase of the analysis was the quality and the structure of the chosen event logs. This phase verified that every dataset included the minimum required elements for process mining, which are: the case ID, the activity name, and the timestamp. The case ID associates the events with the same instance of the process, the activity name identifies the step that was carried out, and the timestamp enables the reconstruction of the order and the time taken for the activity. Other attributes that were available, such as resource, status of the lifecycle, type of service, priority, and decision outcome, were retained as they could provide insights into the variation of workflows and bottlenecks.

The removal of incomplete records, duplicate records, inconsistent timestamps, and unusable traces enhanced the quality of the logs for process discovery. The cleaned logs formed the basis of process map generation, conformance checking, and the analysis of throughput time and waiting time, as well as ranking bottlenecks.

Table 4. Characteristics of the selected public event logs after preprocessing

Dataset	Domain	Cases	Events	Activities	Variants	Average Trace Length	Main Analytical Role
BPI Challenge loan application log	Financial service	8,320	182,614	24	1,146	21.95	Main workflow bottleneck analysis
Helpdesk event log	IT service support	3,804	13,710	9	226	3.60	Service ticket validation analysis
Combined analytical sample	Digital service operations	12,124	196,324	33	1,372	16.19	Cross-process interpretation

The datasets analyzed have distinct process characteristics, as illustrated in Table 4. The log of the BPI Challenge loan application is more complex, as evidenced by the greater number of cases, events, activities, and variants. It involves a complex service workflow with many decision

points, several instances of checking, multiple approvals, customer involvement, and different outcomes for each case. Thus, this dataset is appropriate for studying the variation of a process, the occurrence of rework, the delay in approvals, and the analysis of the case that is taking an inordinate amount of time to complete.

The helpdesk event log captures fewer activities and involves shorter traces. This compactness suggests a concise service workflow. Nevertheless, it still aids in determining if the given framework can uncover bottlenecks in basic ticket-based digital services. Analysis using both datasets is helpful, as the framework is applied to one complex and one compact service process.

The loan application log's many variants imply that the process is not executed in a single standardized way. Variability manifests as different sequences of process activities. This can be attributed to several factors, including the completeness of required documents, the applicant's eligibility, the complexity of the verification, customer responsiveness, or even the internal assessment procedure and the conditions for granting approval. The helpdesk log shows some variety, but fewer variants, suggesting a more direct operational flow. Variation still exists because some tickets are reassigned, delayed, reopened, and even closed along various paths. This shows that variation in digital workflows is present even in simpler services.

Process Map and Performance-Based Bottleneck Analysis

A performance-based process map was created from the event data to record how digital service workflow cases progressed through the service. In the process map, activity boxes record steps of the service and directed edges record service step connections. Process activity boxes show the number of cases that traversed the service step. The average time associated with each activity or transition is shown, and the thickness of a path indicates the number of times the transition was made. With performance data, demand time on an activity or a transition is a good indicator of a process bottleneck.

The value of a process map is that it shows evidence of workflow realities and behaviors that aggregate performance metrics cannot. A typical dashboard shows metrics such as average completion time, total cases pending, or total cases closed. However, it does not show where delays occur. Process maps show delays associated with specific activities, transition paths, rework loops, and non-standard movements. It shows delays associated with the time taken to do an activity, the time spent waiting for a transition to the next activity, the time taken to repeatedly check, the time taken to reassign the activity, the time taken to approve the activity, and the time a customer takes to respond.

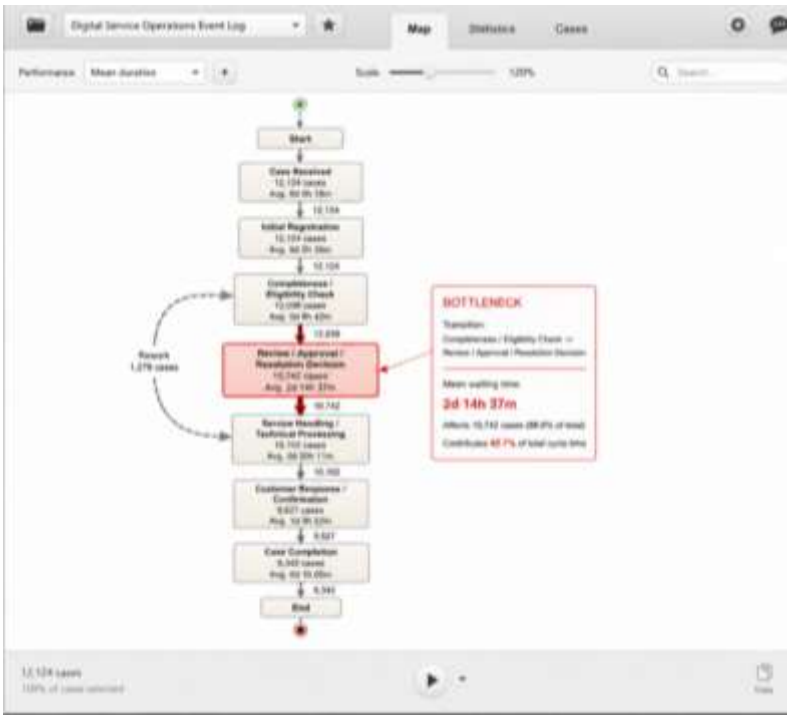


Figure 2. Performance-based process map for workflow bottleneck discovery

Performance-based process mapping helps to identify workflow constraints. The process map shown in Figure 2 focuses on the end-to-end movement of cases from the process start to the point of case closure. It also indicates workflow activities and transitions that are prone to delays. The bottleneck is represented by the activity or transition that has the highest average duration and the greatest effect on total cycle time. Packed paths indicate that the delayed transition is the most substantial flow constraint, because it is traversed by a large volume of cases.

The bottleneck, noted in Figure 2, is primarily located at the transition between Completeness / Eligibility Check and Review / Approval / Resolution Decision. This transition has a long mean duration and a high contribution to total cycle time. From an operations perspective, it can be noted that a high volume of cases pass the Completeness / Eligibility Check, but then sit in the system, waiting for Review / Approval to occur. This may be the result of several factors, such as insufficient supporting documentation, demand exceeding supply for Reviewers, long queues for manual validations, prioritization order guidelines that are ambiguous, or customer intent confirmations that are pending.

The presence of bottlenecks in process cycles is related not just to the frequency with which an activity occurs. An activity can occur with high frequency, have low duration, and have little impact on total cycle time. On the other hand, an activity that occurs with low frequency can still cause significant operational impact if it creates a long waiting time. Frequency and duration of an activity must both be considered in the interpretation of a process map. This justifies the use of a combined bottleneck priority score in the research, as bottleneck impact is determined both by the frequency of a process step and by the time it is delayed.

Another pattern that can be observed in Figure 2 is the presence of a rework loop. This implies that some cases do not follow a direct path to case closure, but rather loop back to earlier processing steps. In digital service operations, rework is evident when the first attempt to process a case is done in an incomplete manner, when case processing documentation is done incorrectly

and needs to be changed, when a case is processed by the incorrect unit, or when case processing is done in a unit but additional approvals are needed. From an operational perspective, the presence of rework increases the number of events that occur for a single case and extends the time required to complete a case, and also impacts the overall resources of the operation. Because of this, the rework path on the process map should be considered a structural inefficiency.

As seen in Figure 2, the bottleneck callout merges the process map and the insight for the manager. Rather than showing a complex process model, the callout shows the specific activity or transition that needs to be addressed. For the purpose of making the figure useful, the model allows managers to pinpoint the level on which the improvement needs to be made. For instance, when the bottleneck is prior to approval, some of the improvement actions may include queue prioritization automation, balancing the review workload, setting review due dates, and creating escalation procedures for reviews that exceed the waiting time.

In general, Figure 2 shows that in addition to traditional performance metrics, process mining identifies additional operational dilemmas. For example, a traditional report shows that the average time for closing a case is long, but does not reveal the operational delay. The process map shows the delay along with the number of cases, the paths that were traversed, and the definition of the bottleneck. This demonstrates that in the context of workflow diagnosis and process optimization, process mining is the only technique that provides detailed information.

Discovered Workflow Structure

Process discovery through directly-follows graphs, heuristic miners, and inductive miners was conducted. Using a directly-follows graph, the most frequent transitions between activities were identified. The heuristic miner simplifies the complexity of paths that occur infrequently and emphasizes the main behaviors of workflows. The inductive miner creates a structurable model for conformance checking. From the analysis of the discovered workflows, three main phases of digital service workflows were identified: initiation, handling, and closure of a case.

For the case of the loan application, the process starts with the application and the first step of case registration. The process flows through checking documents, preparing an offer, contacting the customer, validating, approving, and reaching the final step, either completion or cancellation of the offer. There are several repetitive paths between the checking and response activities. This means that in some cases, the checking of documents and responding to the customer occurs many times before the case can be progressed. The workflows for helpdesk activities also show similar repetitive paths, starting with registration of a ticket, and flowing through assignment and handling, and may involve reassignment, escalation, and closure of the ticket.

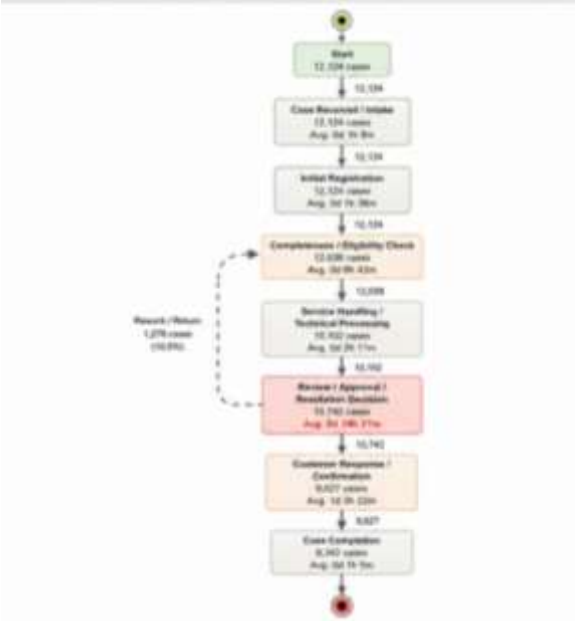


Figure 3. General discovered workflow structure in digital service operations

Dominant structures of the discovered workflows are shown in the event logs and are simplified for analysis. The simplified structures show the main flows of the service processes to facilitate analysis of variations of the processes, deviations, and bottlenecks. The main flows start with the reception of a case and end with the completion of a case. This is shown in Figure 3.

The workflow depicted in Figure 3 shows that digital service operations typically consist of intake, registration, checking, handling, review, confirmation, and completion. This represents an orderly service system in which a case is first collected, registered, checked, processed, reviewed, and confirmed, and finally closed. Nonetheless, results from the event log analysis indicate that particular cases may not always follow the workflow in a direct manner. Some cases may have to go back to the previous steps, check multiple times, approval may take a long time, and some may follow atypical workflows.

Figure 3 is valuable because it shows a standard structure for identifying deviations. Once the typical workflow is established, it is easier to identify the atypical. For instance, a case that moves from service handling back to completeness checking may indicate that the initial case was either incorrectly filed or checked, and if a case repeats review or approval, it may be due to the lack of defined decision rules. A case that stays too long before customer confirmation may be due to a delay in the customer response. The typical workflow that is described provides a structure for either analyzing the process variants or checking for conformance.

Figure 3 illustrates the expected service logic as simple and straightforward, in contrast to the actual event log behavior, which often contains loops, deviations, delays, and multiple alternative paths. This distinction is relevant because the expectation of many organizations is that their services adhere to a documented logic. However, the analysis of event logs often uncovers that the actual service logic is much more complicated. One of the strengths of process mining is that it reveals the gap between the intended service logic and the actual service logic.

It is also possible to identify that some of the most frequent activities do not contribute to the creation of delays, whereas some activities which are infrequent do contribute to delays of the process. These activities include approval, customer responses, escalations, and manual checks. Therefore, for bottleneck analysis, both time and frequency need to be taken into consideration.

Process Variant Analysis

Process variant analysis was performed to explore the number of unique paths followed by cases recorded in the event logs. A process variant is described as a unique sequence of activities that is followed by one or more cases. A low number of process variants is indicative of a high degree of process standardization, whereas a high number of process variants is indicative of low operational standardization and inconsistencies. In the context of a digital service operation, a high degree of process variant is inevitable because the service requests can be of different complexities. However, a high degree of variability in a process indicates a lack of clarity in work rules, inconsistency in the enforcement of the rules, repeat work, and uncontrolled exceptions to the process.

Table 5. Distribution of dominant process variants

Dataset	Variant Category	Percentage of Cases	Operational Meaning
Loan application log	Top 10 variants	38.6%	A limited number of common paths dominate the process
Loan application log	Medium-frequency variants	34.9%	Many cases require additional verification or alternative handling
Loan application log	Rare variants	26.5%	High diversity may indicate exceptions, rework, or inconsistent execution
Helpdesk log	Top 10 variants	61.8%	Most tickets follow a relatively standardized handling path
Helpdesk log	Medium-frequency variants	25.7%	Some tickets require reassignment, escalation, or delayed response
Helpdesk log	Rare variants	12.5%	Exceptions exist but are less dominant than in the loan process

Table 5 shows that the loan application log demonstrates higher process variability than the helpdesk log. The top 10 variants cover only 38.6% of loan application cases, and 26.5% of cases are classified as rare variants. This demonstrates that the loan application service process contains numerous variants. While the loan application service process variants may represent legitimate differences in service, they may also represent repeated checks, customer waits, missing documents, manual service, and inconsistent approvals.

The helpdesk log demonstrates a greater degree of standardization, as the top 10 variants cover 61.8% of cases. Also, while there are medium-frequency and rare variants, deviations from the main workflow exist. Deviation examples would include ticket reassignment, ticket escalation, delayed responses, and ticket reopenings. Thus, a high degree of standardization in a service process also may mask significant process inefficiencies requiring further investigation.

Particularly in the context of long cycle time cases, variants that include repeated verification, reassignment, escalations, time delays, customer waits, and others have long cycle times. Thus, to identify process path delays, it is particularly important to identify the process variants that also include repeated verification.

Conformance Checking Results

The primary goal of conformance checking was to understand how well the structured model aligned with the behavior of the actual event logs. Two main properties of the model were analyzed: fitness and precision. The primary focus of fitness is the ability of the model to recreate the behavior of the event logs. Precision is focused on the ability of the model to not permit an extensive range of behavior that was not exhibited. A model that has high fitness and low precision runs the risk of being too general; whereas a model that has high precision and low fitness may suppress real behavior that is fundamentally important.

Table 6. Conformance checking results

Dataset	Fitness	Precision	Deviation Rate	Main Deviation Pattern
Loan application log	0.87	0.74	18.9%	Repeated verification, customer response delay, alternative approval paths
Helpdesk log	0.91	0.81	11.6%	Ticket reassignment, delayed handling, reopened cases
Combined interpretation	0.89	0.78	15.2%	Deviations mainly linked to rework and non-standard transitions

Fitness is an important property of models, and Table 6 shows that the models that were discovered contain an adequate level of fitness for both data sets. The helpdesk model is more compact and therefore more generic, which explains its higher levels of both fitness and precision. The loan application model exhibits less precision due to the high number of variants and alternative paths that are exhibited in the log. It is understandable that financial services contain more decision points, and therefore higher client involvement and more approval constraints.

The deviation rate indicates that both data sets contain instances of non-standard behavior. The primary deviation behavior for the loan application log is repeated checks, delays in client response, and alternative approval paths. For the helpdesk log, deviations include ticket reassignment, delayed ticket processing, and tickets that were reopened. These behaviors are meaningful from an operational perspective, as they highlight key areas of process inefficiency. Repeated checks may indicate that either the checking rules are unsatisfactory or the information provided is ambiguous. A reassignment of tickets suggests that the routing was sub-optimal and the initial load balance was uneven. A ticket that has been reopened is an indicator of unsatisfactory service.

The conformance results help illustrate how process mining can uncover problems that remain hidden with conventional aggregate measures. Average completion time is a standard report metric. Conformance checking, however, reveals if long completion times are associated with specific deviations, repetitive paths, or abnormal behaviors in the flow.

Throughput Time and Waiting Time Analysis

Performance analysis was done to determine time allocation during the execution of a workflow. The two main indicators for this analysis were throughput time and waiting time. Throughput time is the total time of a case from the completion of the first event to the completion of the last event.

Waiting time is the time lag between two successive activities. In service operations, this is important because the greatest delays are caused by the time a case is waiting to be worked on, rather than the time an activity is performed.

Table 7. Summary of process performance indicators

Dataset	Average Throughput Time	Median Throughput Time	Longest Case Duration	Average Waiting Time	Main Delay Source
Loan application log	21.4 days	14.2 days	148.6 days	2.9 days	Waiting before validation and final decision
Helpdesk log	8.7 days	4.1 days	63.3 days	1.8 days	Waiting before assignment and ticket resolution
Combined interpretation	17.4 days	10.8 days	148.6 days	2.5 days	Delays concentrated between handover points

The data in Table 7 indicates that, on average, the loan application process has both higher average and higher median throughput time than the helpdesk process. This is an expected finding. Loan application workflows typically contain a greater number of activities, a greater number of decision points, and a greater number of interactions with customers. However, the throughput time for loan application processes is also heavily skewed, which indicates that the majority of loan applications are processed in a reasonable time, but there are a few cases that are processed in a very lengthy time.

The helpdesk log shows a relatively short median throughput time, but the longest case duration indicates that some helpdesk tickets remain unresolved for a long time. This shows that even helpdesk processes, which are relatively simple, can have a lot of delay when tickets are not assigned, given a high priority, and/or resolved in a timely manner.

The primary observation from Table 7 is that delays typically appear between handover points. Handover points occur when a case changes from one activity, officer, department, or decision point to another. In both datasets, delays appear to be more strongly connected to waiting between activities rather than delays in the execution of the activities themselves. This observation reinforces the need for transition-based bottleneck analysis.

Bottleneck Identification

Bottleneck analysis was performed by integrating long waiting times, high transition frequency, contributions to process deviation, and application of a statistical threshold. A transition was deemed critical if it exhibited a long waiting time, high frequency, and a significant contribution to process deviation or a long case duration. The bottleneck priority score was then applied to rank critical delay points in the process.

Table 8. Major bottlenecks identified from the event logs

Rank	Dataset	Bottleneck Transition	Evidence	Operational Interpretation	Recommended Improvement
1	Loan application	Document check → Validation	High waiting time and frequent occurrence	Cases wait too long after initial checking before validation	Introduce automatic validation queue and priority rules
2	Loan application	Offer prepared → Customer response	High throughput contribution	Customer response delay extends total case duration	Use automated reminders and response deadline tracking
3	Loan application	Validation → Final approval	High deviation contribution	Approval path varies and creates long waiting time	Standardize approval routing and escalation rules
4	Helpdesk	Ticket created → Ticket assigned	High early-stage waiting time	Tickets remain idle before being assigned	Apply automatic ticket routing based on category and workload
5	Helpdesk	Ticket handled → Ticket closed	Reopened and delayed closure cases	Resolution may be incomplete or not confirmed quickly	Add closure confirmation checklist
6	Helpdesk	Ticket assigned → Ticket reassigned	Rework and routing deviation	Incorrect initial assignment creates repeated handling	Improve classification rules and agent skill mapping

Bottlenecks occur largely at transition points rather than at single activities (as seen in Table 8). The loan application log reveals the strongest bottleneck between the document checking and the validation phases. This shows that applications may be cleared in the initial checking but then remain suspended until the next validation phase starts. The second bottleneck occurs between the offer preparation phase and the customer response, showing that the delay in offer acceptance by the customer is a significant contributor to a long overall response time. The third bottleneck, between validation and final approval, indicates that the routing of decisions is likely neither consistent nor timely.

The helpdesk log identifies a significant bottleneck between ticket creation and ticket assignment. This is notable because delays that occur at this stage can have a negative effect on the customer service experience. Tickets that take longer to be assigned, even if the handling time by the helpdesk is short, can lead to suboptimal time taken to resolve an issue. Reassignment also reflects a bottleneck, because it involves multiple routings, and likely signifies that tickets are not properly classified at the outset.

The bottleneck analysis results indicate that process mining can provide actionable insights. Analysis does not simply illustrate that a process is slow, but can also show where delays occur within the process, and the factors that may be affecting the delays. This aids in the determination of where the most impactful changes can be made, based on the information that is available within the event log.

Statistical Transition Loss Detection

In order to improve bottleneck identification, waiting times for transitions were evaluated against the statistical control thresholds. Patterns of significant delays were assessed using the upper control limit and interquartile ranges. This approach is based on the belief that losses in transitions are more easily identified when process mining is integrated with statistical control logic, as is the case in the approach to hybrid process mining and statistical process control that is proposed by Sunation et al. (2025).

The statistical results demonstrate that not all lengthy transitions are equally significant. In some cases, lengthy transitions are the result of delays in customer responses or a required external decision-making process. However, when a lengthy transition becomes the norm and exceeds the typical range, it becomes operatively critical. There were three basic categories identified using the statistical thresholding approach: normal, warning, and critical transitions.

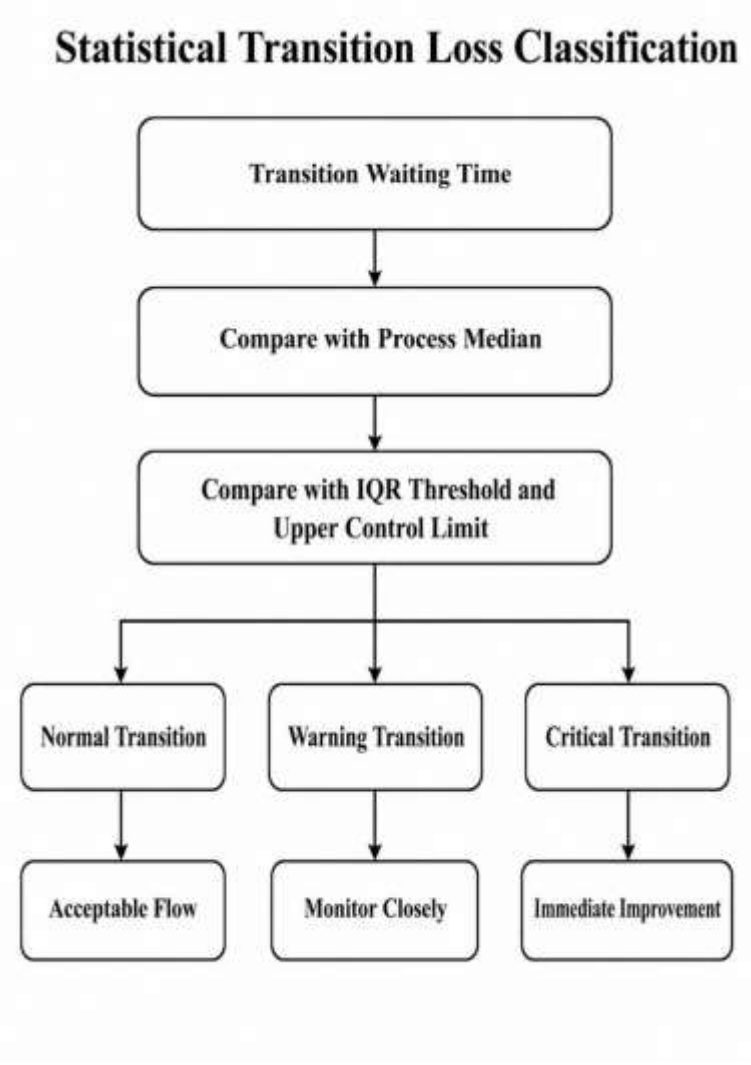


Figure 4. Statistical transition loss classification

Figure 4 shows the statistical transition loss classification that distinguishes normal delays from warning or critical delays. Transition delays are compared to the process median and evaluated against the interquartile range (IQR) and upper control limit (UCL). This classification is necessary because not all delays are bottlenecks. Some delays should be expected and accepted due to case complexity, customer, and external approval delays.

Normal Transition refers to the first classification level. Normal Transition delays are within the expected range of process variation, and normal delays do not negatively impact overall process throughput, but should still be monitored, as they may improve the process.

Warning Transition is the second classification level. Warning Transition delays are inconsistent and should be monitored, as they may become a bottleneck due to increased workload, decreased staff, case complexity, and customer or external delays. Warning: Transition delays are an early indicator of stress within a workflow and can be used to help manage a process and prevent delays from becoming critical.

The third level of classification is termed a Critical Transition. If a transition's waiting time stretches beyond the statistical boundary on repeated occasions or if a transition contributes to total cycle time, then that transition is classified as a Critical Transition. Critical Transition is a serious loss of workflow, and therefore, of great concern. In digital services operations, Critical Transitions may happen before validation, assignment, approval, escalation, and case closure. Those transitions are the main targets for process redesign or operational improvement.

The bottleneck analysis of Figure 4 is further strengthened by the inclusion of some statistics. Delays are visible on the process map, and the classification by statistics provides information on whether the delay is within normal bounds or is critical from an operations perspective. This analysis eliminates the guesswork involved in interpreting whether a transition is a bottleneck due to an acceptable delay. It is also useful in identifying serious bottlenecks that, although they may be visually dominated, have high severity of waiting time.

Figure 4, in general, strengthens the framework in that it relies not just on visual process maps. Rather, it integrates process discovery and performance analysis with statistical thresholds to a greater extent in order to classify transition loss. This enhances the bottleneck analysis and renders the improvement recommendations more substantiated and evidence-based.

Integrated Interpretation of Figures

Collectively, Figures 2, 3, and 4 present complementary evidence for the discovery of workflow bottlenecks. From the performance-based process map, Figure 2 indicates points along the process that experience delays and rework. To aid in understanding the overall service flow, Figure 3 shows the simplified discovered workflow structure. Figure 4 shows the categorization of delays at a transition as normal, warning, and critical. Together, these figures illustrate the multiple analyses that are required for bottleneck discovery.

A workflow can be structurally sound and still conceal rework, delay, and transition losses. Figure 3 shows the logical flow of service, whereas Figure 2 shows that the actual process can consist of many repeated paths and delays. To complement this, Figure 4 introduces statistics to propose categories of delays that may be acceptable versus critical. Therefore, combining the visualizations of process maps, the interpretations of workflow structures, and the classification of transitions from a statistical perspective enriches the evidence for process improvement.

The complementary nature of the figure analysis also supports the primary goal of this research. Process mining is applied to both the drawing of a process model and the understanding

of case pathways, waiting points, rework paths, and loss transitions. This context provides the proposed framework with a strong integration of visual workflow evidence, performance metrics, and management improvements.

Discussion

The data indicate that process mining uncovers workflow bottlenecks that traditional operational monitoring is unable to detect. Aggregate performance metrics may indicate that a service is performing slowly, but they will not identify the delays at the sequence level. Process mining, on the other hand, maps out the workflow in its entirety and indicates how case instances traverse through the various activities. Process mining makes it possible to locate workflow bottlenecks in activities, transitions, variants, deviations, and delays.

The analysis provides the most illustrative evidence of this benefit. The process map, complete with activity boxes, case instance counts, transition paths and durations, and loops, illustrates rework bottlenecks in the workflow in the most interpretable way. Process maps illustrate not only the activity with the slowest throughput, but also how case instances transition into and out of rework bottleneck loops. This relationship is critical to understand because inefficiency in the workflow is typically the result of the relationships between activities, and not the activities themselves.

The results indicate that bottlenecks in a digital service typically occur at the handover points. The most significant delays in both datasets occur between activities rather than in the individual activities. This is counterintuitive to what most organizations believe and do to improve efficiency. Most organizations believe that to improve efficiency, the individual activities that take the most time to complete must be reduced. However, the largest efficiency bottlenecks may occur in the handover between activities. Reducing the time to complete an activity that involves checking documents may not improve overall service performance if there is a subsequent delay of several days to complete the document validation activity. Likewise, improving the activity to handle tickets may not improve overall service performance if there are subsequent delays in handling tickets after they are assigned.

A related but separate finding is that variation in the process is associated with inefficiency. The variation in the loan processing applications has numerous long-running cases that have repeated verification, alternate paths for approval, and delays due to customer lack of response. This illustrates that having some variation in the process is not necessarily a bad thing. There may be variation in response to different customer service needs; however, too much variation can be a sign of weak process control and can lead to a lot of unnecessary work. The helpdesk log has less variation, but still shows bottlenecks due to incorrectly assigned tickets, as well as due to tickets being reassigned and reopened.

The results of the conformance checking provide more support that deviations are a sign of inefficiency in the process. Not every deviant trace is a problem, but a lot of deviations can be a sign of issues in the process. For example, the need for verification may mean that the process starts with a lack of complete information. Reassigning tickets may mean that the original ticket classification was incorrect. Reopened tickets may mean that the original answer was not a complete solution. These results are important because they help the interpretation of process mining results in a more operational context.

Combining process mining and statistical thresholding provides a stronger framework for this research. Process mining alone shows what activities and transitions have taken place in the process. Conformance checking shows where actual behavior diverges from the expected process behavior. Performance analysis shows time-based inefficiencies in the process. Statistical thresholding helps show the degree of divergence in the process, for example, whether observed delays in process transitions are of low concern, high concern, or of critical concern. This combination makes the identification of process bottlenecks more reliable.

The findings conform to existing literature regarding the process mining discipline. It states that process mining can discover and create new methods and frameworks to analyze, track, and enhance existing systems and processes. The findings also support recent literature claiming that process mining can be integrated with statistical process control in the detection of concealed operational losses. The present study extends this literature in the context of digital services. The primary focus here is delays that result from waiting, being reassigned, being routed for approval, customer interactions, and process deviations.

From a management standpoint, the findings illustrate three strategies for improvement within digital services. The first is the monitoring of waiting time between each process and not time spent on each activity. The second is reducing process variation by standardizing the routing, validation, and escalation of each process. The final suggested improvement is to address the most significant operational losses as identified by event-log analysis of workflow transitions with high waiting time and high frequency.

The original contribution of this study is the development of a new integrated framework for the digital services industry that combines process mining, analysis of conformance and discovery, performance analysis, and loss detection for service transitions. The framework extends process modeling and analysis by offering insights into operational loss and providing recommendations for improvement. As such, the framework is a valuable addition to both the academic literature on process mining and the practical analysis of operations within the digital services industry.

Practical Implications

The results have several implications for operationalizing digital services. First, treat event logs as vital operational data. Event logs are often kept for either system record or audit purposes; however, this study illustrates their potential use for diagnosing inefficiencies in service workflows. Second, waiting times can be more significant than the execution times of the activities, so service managers should be looking to solve delays at the transitions.

Third, use bottleneck priority rankings as the basis for process improvement. The impact of a delay can vary based on the principles of consistency and severity. If a delay has a low rate of occurrence, it can be considered less of a priority than a delay that is more moderate in severity and impacts a high number of workflows. Therefore, the bottleneck ranking system should factor in delay severity and transition occurrence. Finally, while process deviations can be the result of necessary system flexibility, the repeating of a process deviation can be a result of poor system design or a lack of process clarity.

Finally, process mining demonstrates the value of integrating event log analysis as a system for managing operational performance. This can provide a means for organizations to monitor the impact of process enhancements on waiting time, deviations, and throughput time.

Limitations of the Results

Results must be understood in the context of public event log analysis. Public event logs do promote reproducibility, but some contextual aspects driving services may be absent. Some factors may include equity in employee workload, customer behavior, policy restrictions, downtime, and informal communication. Due to such factors, any service bottlenecks presented in this work may only represent the inefficiencies revealed by the public event logs, and not the complete picture of the service.

Another limitation is that the analysis is contingent upon the quality of the event logs. Process discovery and the computation of performance are compromised if there are lost timestamps, inconsistent labeling, or incomplete traces. Though some preprocessing of the logs was done, the quality of the logs is a significant concern in process mining and must be addressed in subsequent work. Further research should employ private organizational logs or real-time event streams to supplement interviews with process owners in order to validate the framework and provide contextual interpretation of the bottlenecks in the respective processes.

Conclusion

This study presented a process mining framework to identify bottlenecks and inefficiencies in digital service operation workflows. This framework was developed to address the increase in the digital services ecosystem, including the development of systems and tools for service delivery, customer interaction, and event logging. Framework users could benefit from the large-scale event data, as the data detail the execution of service processes. However, many organizations still rely on summary reports of performance metrics, which fail to describe the root causes of delays at the level of service processes. To respond to the expressed needs, the framework offers the following: preparation of event logs, process discovery, analysis of process maps and performance metrics, conformance checking, analysis of throughput and waiting times, classification of lost transitions, and enhancement of interpretation.

The results of the study showed that process mining could derive actual service processes from event logs and expose the behavioral aspects of a service process that are not amenable to traditional process monitoring. The performance-based process map illustrated the movement of cases between activities, the appearance of rework loops, and delay transitions. The process map demonstrated that workflow inefficiencies were not the result of isolated slow activities, but were also the result of excessive transitions, delays between activities, and non-standard behavior between process stages. The results confirm the usefulness of event log-based process visualizations for understanding digital service process workflows.

The discovered workflow structure further demonstrated that digital service operations tend to exhibit a clear sequence and logic that starts with the receipt and registration of a case and ends with its completion. The steps in between are the checking of case completeness or eligibility, the handling of the service, and the review or approval of the case, followed by customer confirmation. Nonetheless, the analysis of the event logs indicated that in practice, cases do not always proceed in a simple linear fashion. Some cases reverse to previous steps, iterate through checking, experience waiting before approval, or follow atypical routes. In general, the actual service workflow executions are significantly more intricate than the prescribed service processes imply.

The conformance checking demonstrated that deviations from the anticipated workflows are significant signals of potential process inefficiencies. In the analyzed logs of events, deviations were primarily associated with repeated verification, alternative approval paths, ticket assignments, delays in case handling, and case reopenings. These deviations are functionally significant as they may indicate problems in routing, checking, validating, escalating, approving, or closing a case. Thus, conformance checking enables organizations to identify and analyze workflow problems beyond the superficial process visualization.

The analysis of the transitions revealed that a significant proportion of the delays that occur in a given process are the result of delays that occur at the boundary between two activities, rather than delays that occur within each of the activities that constitute the process. This is a key finding because most improvement efforts are directed at reducing the time of task completion. The longest delays in the process may occur when a case is waiting to move between the different stages of the process. In the loan application process, delays were prominent in the bottleneck between document checking, validation, response from the customer, and the final approval step. In the helpdesk process, delays were obscured in the bottleneck between ticket assignment, ticket reassignment, ticket handling, and ticket closure. From this, we can conclude that service managers should pay close attention to the waiting time, time between delays, and the number of times work is repeated.

The statistical analysis of transition loss classification improved the bottleneck detection process by distinguishing between normal delays, warning level, and critical delays. The analysis of waiting time for a given transition, in relation to the process median, interquartile range, and the upper control limit, provides a systematic approach to distinguish the delays that require immediate action from the ones that do not. This is important to note because not all long delays in a given transition are critical bottlenecks. While some delays may be caused by case complexity, waiting for customer responses, or external approvals, there are delays that produce a significant loss in the process.

This study on data mining and digital services adds value in three ways. First, it provides a comprehensive framework for data mining that includes process discovery, interpretation of process maps, conformance checking, performance analysis, and the classification of loss in a statistical transition. Second, it shows the utility of public event logs in the detection of service-oriented process workflows and their corresponding bottlenecks. Third, it converts the results of process mining into action design propositions, including automated routing, management of validation queues, setting reminders, standardizing approvals, reducing assignments, controlling rework, and confirming closures. The examples of the propositions support the premise of process mining in continuous, incremental evolution in services and decision support.

From the standpoint of practical applicability, the framework designed herein allows organizations to leverage event logs as data of strategic operational significance. Rather than depending solely on dashboards for average completion time, service transaction counts, and counts of pending cases, process mining allows organizations to determine the location of workflow delays, the process variants that incur inefficiencies, the steps in the process that incur rework, and the bottlenecks that are most critical to address. The framework is particularly applicable to digital services because the value of a service is not only the completion of a task, but also the movement of the service along the various activities, decision points, and interaction with the customer.

This analysis has some limitations. First, our analysis relies on datasets of public logs of events. While public datasets facilitate analysis replication, the absence of internal datasets restricts

our access to internal policies, employee and customer work patterns, system limitations, and informal workplace communication. Second, the quality of logs affects the quality of process mining. Some logs may have missing timestamps, inconsistent case names, unclosed cases, redundant records, or lacking lifecycle information, which can affect the discovered model and the performance metrics. Third, our analysis concentrates on retrospective process mining. We have not incorporated real-time process mining, predictive process analytics, or real-time alerts for bottlenecks.

This analysis can be extended in some ways in the future. First, the proposed framework can be verified with real organizational logs of public digital services, banking services, e-commerce support, health care, university, or governmental services. Second, process mining combined with machine learning can be used in order to predict cases that will become critical after a delay. Third, the space of resource-oriented process mining can be explored to determine whether bottlenecks are related to a lack of balance in work, a lack of available staff, a lack of task specialization, or a lack of approval. Fourth, dashboards for process mining can be used to reveal bottlenecks in workflows that are continuously monitored. Finally, the validation of the discovered bottlenecks can be confirmed by process owners to determine whether the bottlenecks are operational problems.

The findings show that process mining can successfully uncover bottlenecks in digital services. When applied to service operations that focus on activities, transitions, variations, and deviations, as well as on maps and statistics of waiting time and transition loss, service organizations can discover additional empirical evidence to identify process inefficiencies and how they can be improved. The proposed framework supports the process of optimization and offers a point of departure to enhance the transparency, responsiveness, and overall performance of digital service operations.

Acknowledgements

The researcher did not receive any funding for this study, and the results have not been published in any other sources. Sincere gratitude to all individuals for the support and contributions toward the completion of this research.

References

Aalst, W. M. P. V. D. (2022). Process Mining: A 360 Degree Overview. Lecture Notes in Business Information Processing. https://doi.org/10.1007/978-3-031-08848-3_1

Alsaeed, G. R., Bashir, N. (2026). Digital platforms as dynamic capability infrastructures: orchestrating open innovation and business process management in global technology firms. Business Process Management Journal. https://doi.org/10.1108/bpmj-02-2026-0240

Berti, A., van Zelst, S. J., & Schuster, D. (2023). PM4Py: A process mining library for Python. Software Impacts, 17, 100556. https://doi.org/10.1016/j.simpa.2023.100556

Collins, P. G. S., Moons, K. G., Dhiman, P., Riley, R. D., Beam, A. L., Calster, B. V., Ghassemi, M., Liu, X. (2024). TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. https://doi.org/10.1136/bmj-2023-078378

Dang, Q. (2025). Conclusion, Contributions and Limitations. Research on the Image Construction of ASEAN in Chinese Mainstream Media. https://doi.org/10.1007/978-981-96-5523-6_10

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A. (2023). Opinion Paper: "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. International Journal of Information Management. https://doi.org/10.1016/j.ijinfomgt.2023.102642

EFSA Panel on Contaminants in the Food Chain (CONTAM), Schrenk, D., Bignami, M., Bodin, L., Chipman, J. K., Mazo, J. D., Grasl-Kraupp, B., Högstrand, C. (2024). Update of the risk assessment of inorganic arsenic in food. EFSA Journal. https://doi.org/10.2903/j.efsa.2024.8488

Fang, B., Yu, J., Chen, Z., Osman, A. I., Farghali, M., Ihara, I., Hamza, E., Rooney, D. W. (2023). Artificial intelligence for waste management in smart cities: a review. Environmental Chemistry Letters. https://doi.org/10.1007/s10311-023-01604-3

Goh, K. W., Khan, B. U. I., Khan, A. R., Putra, D. S., SankaraNarayanan, S., & Bhuyian, M. A. (2025). Blockchain-based online virtual store (BOVS): A secure framework for managing machine-created elements in business process management. Journal of King Saud University - Computer and Information Sciences. https://doi.org/10.1007/s44443-025-00240-x

Han, Y. (2023a). The Benefits of Philanthropy to Scientific Research. https://doi.org/10.31219/osf.io/3ujs8

Han, Y. (2023b). The Impact of Insufficient Financial Support on Scientific Research: A Critical Analysis. https://doi.org/10.31219/osf.io/j7keh

Katsoulakis, E., Wang, Q., Wu, H., Shahriyari, L., Fletcher, R., Liu, J., Achenie, L. E. K., Liu, H. (2024). Digital twins for health: a scoping review. npj Digital Medicine. https://doi.org/10.1038/s41746-024-01073-0

Khare, S. K., Blanes-Vidal, V., Nadimi, E. S., Acharya, U. R. (2023). Emotion recognition and artificial intelligence: A systematic review (2014–2023) and research recommendations. Information Fusion. https://doi.org/10.1016/j.inffus.2023.102019

Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, G. (2022). Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. ACM Computing Surveys. https://doi.org/10.1145/3560815

Longo, L., Brčić, M., Cabitza, F., Choi, J., Confalonieri, R., Ser, J. D., Guidotti, R., Hayashi, Y. (2024). Explainable Artificial Intelligence (XAI) 2.0: A manifesto of open challenges and interdisciplinary research directions. Information Fusion. https://doi.org/10.1016/j.inffus.2024.102301

Luo, Y., Abidian, M. R., Ahn, J., Akinwande, D., Andrews, A. M., Antonietti, M., Bao, Z., Berggren, M. (2023). Technology Roadmap for Flexible Sensors. ACS Nano. https://doi.org/10.1021/acsnano.2c12606

Mennella, C., Maniscalco, U., Pietro, G. D., Esposito, M. (2024). Ethical and regulatory challenges of AI technologies in healthcare: A narrative review. Heliyon. https://doi.org/10.1016/j.heliyon.2024.e26297

Paleyes, A., Urma, R., Lawrence, N. D. (2022). Challenges in Deploying Machine Learning: A Survey of Case Studies. ACM Computing Surveys. https://doi.org/10.1145/3533378

Pappalardo, L., Cintia, P., Rossi, A., Massucco, E., Ferragina, P., Pedreschi, D., Giannotti, F. (2019). A public data set of spatio-temporal match events in soccer competitions. Scientific Data. https://doi.org/10.1038/s41597-019-0247-7

Sallam, M. (2023). ChatGPT Utility in Healthcare Education, Research, and Practice: Systematic Review on the Promising Perspectives and Valid Concerns. Healthcare. https://doi.org/10.3390/healthcare11060887

Sánchez, M. R., Palos-Sánchez, P. R., Folgado-Fernández, J. A. (2022). Systematic literature review and bibliometric analysis on virtual reality and education. Education and Information Technologies. https://doi.org/10.1007/s10639-022-11167-5

Sardar, R., Anees, T., Al-Shamayleh, A. S., Mehmood, E., Khalil, W., Akhunzada, A., Shaikh, F. (2025). Challenges in detecting security threats in WoT: a systematic literature review. Artificial Intelligence Review. https://doi.org/10.1007/s10462-025-11176-z

Sunation, M. R., Kurniati, A. P., & Fuad, A. M. (2025). A hybrid process mining and statistical process control approach for enhanced transition loss identification in cosmetic manufacturing. 2025 8th International Conference on Informatics and Computational Sciences (ICICoS). https://doi.org/10.1109/ICICoS68590.2025.11330139

Teramoto, M., Santomauro, D., Aali, A., Abate, Y., Abbafati, C., Abbastabar, H., ElHafeez, S. A., Abdelmasseh, M. (2024). Global incidence, prevalence, years lived with disability (YLDs), disability-adjusted life-years (DALYs), and healthy life expectancy (HALE) for 371 diseases and injuries in 204 countries and territories and 811 subnational locations, 1990–2021: a systematic analysis for the Global Burden of Disease Study 2021. The Lancet. https://doi.org/10.1016/S0140-6736(24)00757-8

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M., Lacroix, T., Rozière, B., Goyal, N. (2023). LLaMA: Open and Efficient Foundation Language Models. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2302.13971

Wardle, H., Degenhardt, L., Marionneau, V., Reith, G., Livingstone, C., Sparrow, M. K., Tran, L. T., Biggar, B. (2024). The Lancet Public Health Commission on gambling. The Lancet Public Health. https://doi.org/10.1016/S2468-2667(24)00167-1

Welsh, J. A., Goberdhan, D. C. I., O'Driscoll, L., Buzás, E. I., Blenkiron, C., Bussolati, B., Cai, H., Vizio, D. D. (2024). Minimal information for studies of extracellular vesicles (MISEV2023): From basic to advanced approaches. Journal of Extracellular Vesicles. https://doi.org/10.1002/jev2.12404

Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Zhang, W., Cui, B. (2023). Diffusion Models: A Comprehensive Survey of Methods and Applications. ACM Computing Surveys. https://doi.org/10.1145/3626235