



Article

A Multi-View Framework to Detect Redundant Activity Labels for More Representative Event Logs in Process Mining

Qifan Chen ^{1,*}, Yang Lu ¹, Charmaine S. Tam ² and Simon K. Poon ¹

¹ School of Computer Science, The University of Sydney, Sydney, NSW 2006, Australia; yalu8986@uni.sydney.edu.au (Y.L.); simon.poon@sydney.edu.au (S.K.P.)

² Centre for Translational Data Science and Northern Clinical School, The University of Sydney, Sydney, NSW 2006, Australia; charmaine.tam@sydney.edu.au

* Correspondence: qifan.chen@sydney.edu.au

Abstract: Process mining aims to gain knowledge of business processes via the discovery of process models from event logs generated by information systems. The insights revealed from process mining heavily rely on the quality of the event logs. Activities extracted from different data sources or the free-text nature within the same system may lead to inconsistent labels. Such inconsistency would then lead to redundancy in activity labels, which refer to labels that have different syntax but share the same behaviours. Redundant activity labels can introduce unnecessary complexities to the event logs. The identification of these labels from data-driven process discovery are difficult and rely heavily on human intervention. Neither existing process discovery algorithms nor event data preprocessing techniques can solve such redundancy efficiently. In this paper, we propose a multi-view approach to automatically detect redundant activity labels by using not only context-aware features such as control-flow relations and attribute values but also semantic features from the event logs. Our evaluation of several publicly available datasets and a real-life case study demonstrate that our approach can efficiently detect redundant activity labels even with low-occurrence frequencies. The proposed approach can add value to the preprocessing step to generate more representative event logs.

Keywords: process mining; activity label; process event log; data quality



Citation: Chen, Q.; Lu, Y.; Tam, C.S.; Poon, S.K. A Multi-View Framework to Detect Redundant Activity Labels for More Representative Event Logs in Process Mining. *Future Internet* **2022**, *14*, 181. <https://doi.org/10.3390/fi14060181>

Academic Editor: Filipe Portela

Received: 19 May 2022

Accepted: 8 June 2022

Published: 9 June 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Process mining combines traditional model-based process analysis and data-centric mining techniques [1]. It is a technology known to be useful for understanding business processes and constructing process models by using event logs captured in information systems [2]. Process mining includes process discovery, conformance checking, and enhancement [3]. Among them, process discovery is a paramount task that aims at automatically discovering process models from structured event logs to analyse and improve the internal business processes [4]. Process mining has shown promising potential in many aspects, such as discovering significant insights and improving process performances [5]. A typical event log refers to a collection of events, each with a timestamp that records the executed time. An event represents a unique execution of an activity, which is a well-defined step in the process, such as “doctor appointment”. Cases group these events, also called process instances. For example, a case could be a patient who follows a treatment process in a hospital.

In recent years, with an effort being made to discover accurate and comprehensible process models from structured and clean event logs, many advanced discovery algorithms have been proposed, such as the Heuristic Miner [6] and Split Miner [7]. However, like other data mining technologies, the quality of input data (event logs) has a great impact on the resulting model [2]. Moreover, real-life event logs can suffer various data-quality issues and lead to “spaghetti-like” business models, which refer to models that have

numerous activities and intricate relations. Such spaghetti-like' models are too complex to comprehend easily even for domain experts [8]. The process mining manifesto [9] has emphasised the importance of event log quality in process discovery. The first guideline for process mining is to treat event data as first-class citizens. Suriadi et al. [10] summarised 11 event log imperfection patterns and pointed out that data quality on activity labels in event logs is unique in process mining research, as the quality can be affected by integrating data sources or discrepancies of labelling in the same information system. One particular imperfection pattern is redundant activity labels, which refers to labels that have different syntax but share the same behaviours [11]. Two labels are recognised as sharing the same behaviours not only if they play identical roles in process models but also if they represent the same activity in reality. One contributing factor to such redundancy is data integration from separate systems, e.g., electronic medical records (EMR) from different hospitals, because multiple systems use different labels for the same activity. The other is the free-text input or human error in providing an initial suggestion in the same system [11]. Redundant activity labels are commonly observed in real-life event logs. For instance, nearly 30% of activity labels in observation tests are considered redundant in the publicly available MIMIC-III database [12]. Event logs with such redundant activity labels can suffer a significant quality loss in the discovered process model, because most existing process discovery algorithms assume that every activity label in the event log is meaningful and unique. Several methods are proposed in the field of process mining to detect redundant activity labels in event logs [11,13,14]. However, none of these methods can accurately detect redundant activity labels in event logs without domain knowledge, especially when the redundant activity labels are less frequent and contain numerical values as attributes.

In this paper, we propose a multi-view framework that efficiently incorporates the control-flow relations, attribute values, and label semantic information of each activity label to detect redundancy in event logs. A consensus is then guided by a majority voting mechanism to integrate results produced from multiple views.

1.1. Motivating Example

To demonstrate how redundant activity labels can introduce unnecessary complications to discovered process models and the motivations behind our proposed framework clearly, we describe a simple patient treatment process as an example. Assuming there are eight activities in the process; that is, (A) registration, (B) visiting the doctor, (C) performing colonoscopy, (D) performing a laboratory test, (E) performing an MRI, (F) performing surgery, (G) paying the bill, and (H) discharging the patients. Hypothetically, the clean event log contains eight traces denoted by $L_1 = \{ABDEFGH, ABCEFGH, ABCDEFGH, ABFGH, ACEFH, BADFGH, BACDEGH, ABCFGH\}$. The goal of automated process discovery algorithms is to construct a process model that can accurately describe the process behaviours [15]. For instance, if we apply the popular algorithm used by Disco, which is a tool widely used to generate visual and actionable insights for process mining [16], we can obtain the process model as shown in Figure 1. The numbers in each box indicate the case coverage in Disco. It is easy to interpret the process model: patients usually register first and then visit the doctor. After that, they may be asked to perform a colonoscopy, laboratory tests, or an MRI. A surgery follows depending on the patient's situation. Finally, they pay the bills and are discharged from the hospital. However, redundant activity labels may exist if the event log contains merged data from sources that do not share a common schema [10]. Therefore, the same real-world activity is recorded with different labels in each source. Suppose there are two activity labels, (B1) "DrSeen" and (B2) "Medical Assign", which both represent activity (B) in the event log L_2 [10]. Moreover, activity labels (H1) "Release C" and (H2) "Release D" represent the same method by which patients are discharged [17]. The event log is denoted by $L_2 = \{ABDEFGH, AB1CEFGH1, AB2CDEFGH2, ABFGH, ACEFH, B2ADFGH2, B1ACDEGH1, ABCFGH\}$. Consequently, we obtained another process model based on this redundant event log, which is shown in Figure 2. Comparing the two process models, it can be seen that such redundancy

brings unnecessary relations (red arcs) and unwanted activity labels (activity B1, B2, H1, and H2) to the process model in Figure 1. Such redundancy causes confusion for process analysts and has a negative impact on the simplicity and comprehension of the discovered process model.




Figure 1. Process model discovered from clean event log L_1 .

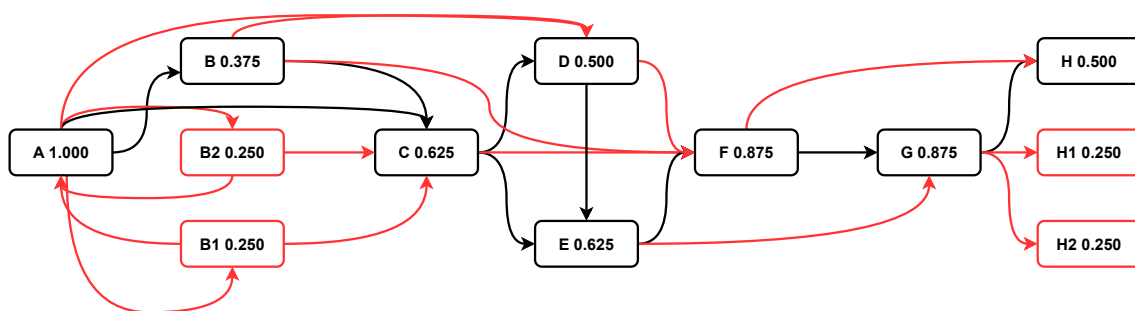


Figure 2. Process model discovered from redundant event log L_2 .

The goal of this paper is to propose a framework as a data preprocessing tool that can accurately detect the redundant activity labels to enhance the quality of event logs for better process analysis in process mining. The proposed approach has a wide range of application in various real-life process mining tasks, especially in the healthcare domain. For example, the proposed approach can be adopted to select representative activity labels when analysing complicated healthcare processes.

1.2. Contributions of This Paper

The contributions of this paper are as follows:

- For the purpose of improving the quality of event logs, a novel data preprocessing framework is proposed for process mining.
- A multi-view framework is proposed to detect redundant activity labels in event logs. In particular, our framework integrates control-flow relations, attribute values, and label semantic information in event logs. In terms of the control-flow relation (i.e., the ordering of activities), we adopt the Earth Mover's Distance (EMD) statistical method to compare the directly-follows and indirectly-follows relations of different activity labels. In terms of the attribute value (i.e., categorical or numerical values of recorded activity labels), activity labels are first clustered and followed by EMD to compare the value's distribution. We assess labels' semantic similarity by using the pre-trained NLP model as another view. A consensus is guided by a decision-making mechanism to integrate the results produced from multiple views.
- Experiments on publicly available datasets under various settings show that our framework can accurately detect redundant activity labels even when the redundant activity labels are infrequent and contain numerical values as attributes compared with the existing state-of-the-art approach.
- A case study in the healthcare domain using the 5-year EMR dataset collected from two local health districts (LHDs) in Sydney, Australia [18], further demonstrates that our framework can be used as a preprocessing tool in real-life event logs.

The paper is structured as follows. Section 2 discusses the background. Section 3 introduces the basic concepts used throughout the paper. Our proposed framework is presented in Section 4. Section 5 presents the experimental results. A real-life case study is explored in Section 6. The paper concludes with Section 7.

2. Related Work

2.1. Process Discovery Algorithms

Various process discovery algorithms have been proposed in the last decade to discover accurate process models from inputted event logs. The alpha miner [19] was the very first discovery algorithm that aims to automatically discover Petri nets based on clean and noise-free event logs. Later research papers extended the original alpha miner to discover invisible tasks [3] and non-free-choice behaviours [20]. However, the alpha algorithm suffers from poor quality in real-life event logs. The Heuristic Miner [6] was proposed to handle noisy event logs by calculating the relative frequency of each dependency from event logs and removing the dependency from the directly follows graph based on a predefined threshold. The Inductive Miner [21] adopts a divide-and-conquer approach to recursively filter out infrequent relations from discovered process trees and can discover sound process models. The BPMN Miner [22] can also handle noisy event logs by employing approximate dependency discovery techniques, which enables the algorithm to detect and remove infrequent relations even if not all instances are matched in event logs. The newly proposed split miner [7] utilises the breadth-first forward exploration to search for the best incoming and outgoing edges for each node on the directly-follows graph while maintaining its connectivity. The algorithm preserves the most frequent incoming and outgoing relations for each activity and filters out the rest. There are also other types of algorithms for discovering process models, such as the genetic algorithm [23] and declarative algorithms using the association rule mining [4].

To summarise, efforts have been made to propose algorithms that can accurately discover process models from event logs. However, most process discovery algorithms require high-quality event logs as input. The performance of existing process discovery algorithms can be affected if data-quality issues are presented in event logs, such as the redundancy of activity labels.

2.2. Event Log Quality

Event log quality has been identified as a critical issue that affects process mining results from many domains in both process discovery and enhancement [1,5]. The process mining manifesto [9] has emphasised the importance of event log quality. The first guideline for process mining is to treat event data as first-class citizens. Later on, Suriadi et al. [10] outlined 11 common event log imperfection patterns, such as incorrect timestamps and redundant labels. Fox et al. [24] and Mans et al. [25] suggested quality frameworks for assessing EMR data in the healthcare domain. Additionally, Bose et al. [26] and Aalst [27] raised concerns for event data quality in process mining. Therefore, apart from proposing advanced process discovery, it is also essential to address data quality as early as the event log level.

However, compared with process discovery, less effort has been made to improve event data quality. Event log quality can be improved by detecting erroneous data and potentially repairing it by relying on a reference model or observed correct data [28]. Conforti et al. [28] proposed to identify and repair events with the same timestamp by using information from correctly ordered events in the log. Rogge-Solti et al. designed an approach to identify and restore missing events by using a reference process model annotated with execution times. Similarly, Sim et al. [29] proposed likelihood-based multiple imputations by event chains to repair missing events in logs. Alharbi et al. [30] proposed an interval-based event selection method to reduce variations in event logs by targeting the behaviour of repeated activities.

In order to detect redundant activity labels, relevant works [11,13,14] suggest ways to address this issue at the event log level. Sadeghianasl et al. [11] proposed a contextual

approach that takes control–flow relations, resources, time, and data attributes into consideration. From the control–flow perspective, the method reports the similarity between rows of the footprint matrix, which may not well distinguish the frequency difference between two activity labels and suffers from noisy or infrequent relations. Thus, the method achieves relatively low accuracy with low-occurrence-frequency activity labels. Other methods largely adopt a probability density function to assess the value distributions between activity labels. This approach reports relatively poor results if there are numerical values as activity attributes, and this is a common phenomenon in a real-life event log. For instance, healthcare logs contain various laboratory tests and medications as activity labels. It relies on a weighted clustering method to combine the final results, which requires domain knowledge or ground truth to determine the best weight setting. The other method [13,14] collaboratively and interactively detects problematic activity labels by adopting a gamified crowdsourcing approach, which utilises gamification elements (e.g., badges) to encourage a large group of domain experts to identify and repair redundant activity labels.

The issue of activity labels has also been studied in process matching areas at the model level [31–34]. These approaches match two process models from different data sources with the aim to find similar structures and activity labels. It is difficult to address redundant labels within the same log because separated logs may have incomplete processes. Hence, they are more widely used in a process similarity comparison instead of solving problematic event logs. Other approaches [35,36] look at activity labels themselves while ignoring other information from logs, which may cause erroneous results.

Though previous studies have made efforts to address redundant activity labels [11,13,14], many of them have difficulties identifying activity labels with low-occurrence frequencies, or if invalid labels have been used. However, redundant labels usually occupy a small portion of event logs. Many of these approaches rely on event logs with categorical resources as attributes instead of numerical values [11]. Nevertheless, many activities have such attributes, especially in healthcare logs, such as laboratory tests and observations. Other approaches [13,14] require domain knowledge to improve the data quality. Hence, we propose a multi-view framework to consider contextual information from event logs and aim to accurately detect redundant activity labels without domain knowledge even when the redundant activity labels are infrequent and contain numerical values as attributes.

3. Preliminaries

Problem Definition

In this section, we introduce some basic concepts used in this paper. The full notations used in this paper is available in Table A1.

Definition 1 (Event Log, Trace, Activity, Event). *An event log L is a collection of traces. A trace $t \in L$ is a sequence of events. A is the set of activities, and an event e is an execution record of an activity $a \in A$. $\#_n(e)$ is a function that obtains attribute values recorded for an event e . For example, $\#_{activity}(e)$ obtains the activity name for an event, and $\#_{attribute}(e)$ obtains the attribute value for an event.*

For example, let $E = \{a, b, c, d\}$ be a set of activities. $t = \langle e_1, e_2 \rangle$ is a trace, where $\#_{activity}(e_1) = a$ and $\#_{activity}(e_2) = b$. $L = \{\langle e_1, e_2 \rangle, \langle e_3, e_4 \rangle\}$ is an event log, where each e_n represents a unique execution record of a specific activity. For the sake of understanding, $\#_{activity}(e_n)$ will be used for each event for the rest of the paper.

Definition 2 (Ordering Relation). *Let L be an event log and $t \in T$ be a trace. For $\forall a, b \in A$, the ordering relations between a and b are defined as follows:*

- *Directly-follows relation: $a >_W b$ holds if there is a trace $t = \langle e_1, e_2, e_3, e_4, \dots, e_n \rangle$ and $i \in \{1, 2, 3, \dots, n - 1\}$ such that $t \in L$ and $\#_{activity}(e_i) = a$ and $\#_{activity}(e_{i+1}) = b$.*
- *Indirectly-follows relation: $a >>_W b$ holds if there is a trace $t = \langle e_1, e_2, e_3, e_4, \dots, e_n \rangle$ and $i < j$ and $i, j \in \{1, 2, 3, \dots, n - 2\}$ such that $t \in L$ and $\#_{activity}(e_i) = a$ and $\#_{activity}(e_j) = b$.*

Because we are only interested in indirectly-follows relations with strong connections, we define the following measurement based on [6] to measure how reliable an indirectly-follows relation is.

Definition 3 (Long Distance Measure).

$$a \Rightarrow_W b = \left(\frac{2(|a \gg_W b|)}{|a| + |b| + 1} \right) - \left(\frac{2Abs(|a| - |b|)}{|a| + |b| + 1} \right) \tag{1}$$

If the first part of the equation is close to 1, then activity a is always followed by activity b . The second part of the equation measures the frequency distributions for activities a and b . A value close to 0 indicates that the frequency of activities a and b is about equal [6]. Therefore, a value of the long-distance measure close to 1 means that the activities a and b have a strong indirectly-follows relation. Based on this, we define a strong indirectly-follows relation between two activities.

Definition 4 (Strong Indirectly-Follows Relation). $a \gg\gg_W b$ holds between two activities $a, b \in A$ if $a \Rightarrow_W b$ is larger than a given threshold p . In this paper, p is set to 0.9, as recommended in [6].

For example, $L = \{ \langle a, b, c, d \rangle, \langle b, c, d \rangle^5 \}$ is an event log. The directly-follows relations are $a \gg_W b, b \gg_W c,$ and $c \gg_W d$. The indirectly-follows relations are $a \gg\gg_W c, a \gg\gg_W d,$ and $b \gg\gg_W d$. The long-distance measures for these three indirectly-follows relations are 0, 0, and 0.92. Thus, the strong indirectly-follows relation is $b \gg\gg_W d$.

Definition 5 (Directly-Follows Graph). A directly-follows graph is defined as $G = (A, K)$ where A is a finite set of activities in the event log (same as Definition 1), and $K \subseteq A \times A$ is a set of directed arcs, which represent directly-follows relations (i.e., $a \gg_W b$ exists if $(a, b) \in K$). An example is shown in Figure 3.

Definition 6 (Pre-Sets and Post-Sets). Let $G = (A, K)$ be a directly-follows graph. For $a \in A$, we have

$$a \bullet = \{ b | b \in A \wedge (a, b) \in K \}$$

and

$$\bullet a = \{ b | b \in A \wedge (b, a) \in K \},$$

where $a \bullet$ and $\bullet a$ are called a post-set and a pre-set of a , respectively. $a \bullet$ represents all the directly outgoing activities from a , e.g., $A \bullet = \{ H, B \}$, and $\bullet a$ represents all directly incoming activities to a , e.g., $\bullet C = \{ H, B \}$, as shown in Figure 3.

Definition 7 (Count Frequency). $|(a, b)|, (a, b) \in K$ counts how many times the relation $a \gg_W b$ occurs in G (e.g., $|(A, H)| = 50$ in Figure 3).




Figure 3. An example directly-follows graph.

4. A Multi-View Detecting Framework

This section describes our proposed framework, shown in Figure 4, to detect redundant activity labels. The underlying principle is that redundant activity labels share the same patterns on both control-flow relations and attribute values. We also include semantic similarity as an additional view. Thus, our framework assesses similarities from the above views by using an EMD statistical method and a pre-trained NLP model. To this end, we first introduce EMD to compare control-flow relations' probability distributions. We then demonstrate how to extend EMD to calculate attribute value similarity. We apply a powerful NLP model in semantic similarity. Finally, we briefly describe how to use the decision-making mechanism to combine results from different views to obtain the final output.

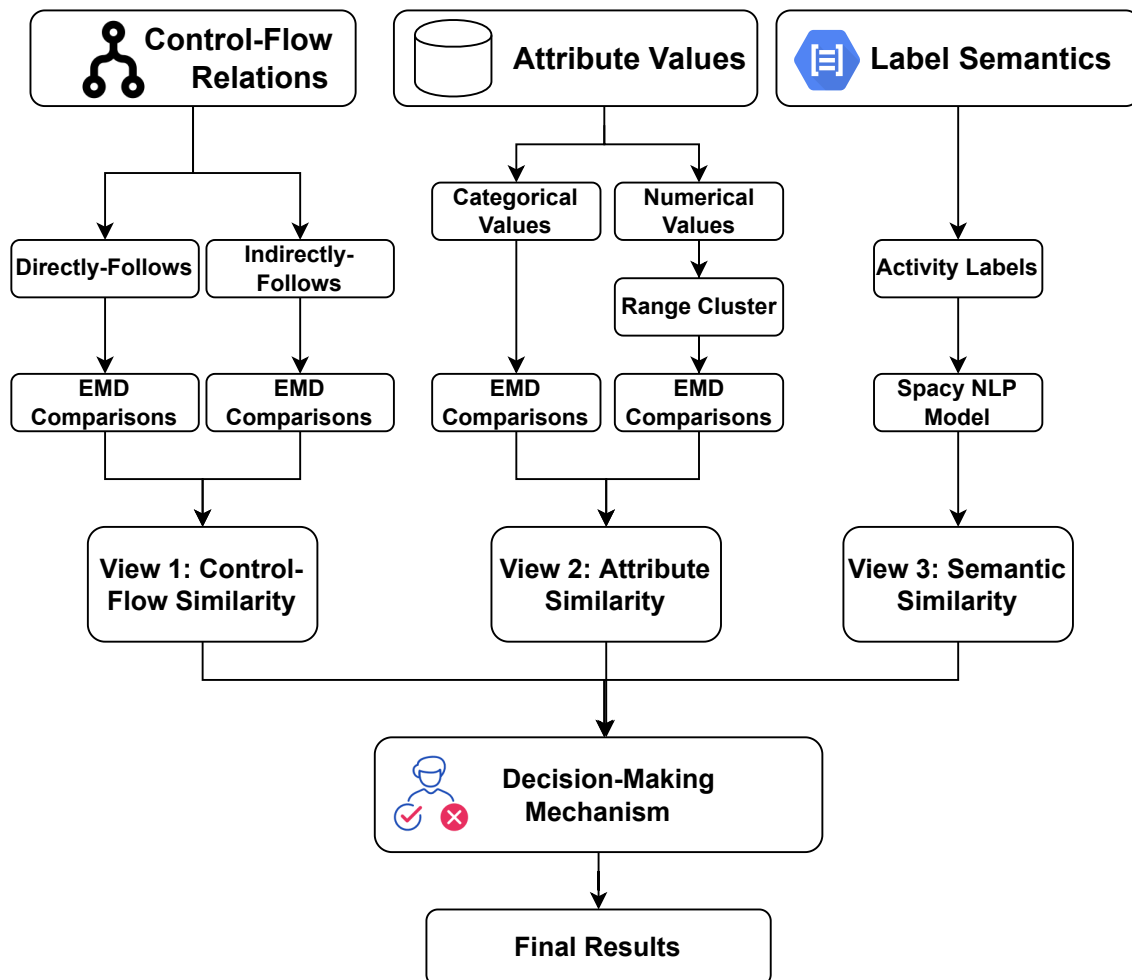


Figure 4. Overview of the proposed framework .

4.1. Earth Mover’s Distance

The EMD [37] is a method for comparing two multiple dimensional probability distributions over a region. It was first proposed as a matrix to retrieve images in the computer vision domain. However, it has been applied to many other fields [38,39]. The EMD calculates the lowest costs of transferring one distribution into another, given that two distributions indicate different ways of accumulating a certain amount of dirt in a region. A distance function defines the cost needed to move dirt between certain piles. It is frequency-aware, as it considers the magnitude of discovered differences, and the difference is determined by the ground distance function that can express different perceptions of similarity [40]. Below we formally introduce the EMD.

Let P be a probability distribution with $p_1, \dots, p_m \in P$ as different clusters and $w_{p_1}, \dots, w_{p_m} \in \mathbb{R}^+$ as the associated weight for these clusters. Another probability distribution Q has the same notations $(q_1, w_{q_1}), \dots, (q_n, w_{q_n})$. A ground distance $D = d(p_i, q_j)$ between cluster p_i and p_j is defined. We seek to find a flow $F = (f_{i,j}) \in \mathbb{R}^{m \times n}$ that minimises the overall costs to transfer P to Q . The following constraints should be followed:

- Non-negativity flow: $f_{i,j} \geq 0, \forall 1 \leq i \leq m, 1 \leq j \leq n$.
 - Sent and receive flow should not exceed weights in P and Q :
 - $\sum_{j=1}^n f_{i,j} \leq w_{p_i}, \forall 1 \leq i \leq m$;
 - $\sum_{i=1}^m f_{i,j} \leq w_{q_j}, \forall 1 \leq j \leq n$.
 - All weights possible have to be sent: $\sum_{i=1}^m \sum_{j=1}^n f_{i,j} = \min(\sum_{i=1}^m w_{p_i}, \sum_{j=1}^n w_{q_j})$
- The optimal flow F is defined as

$$EMD(P, Q) = \min \frac{\sum_{i=1}^m \sum_{j=1}^n f_{i,j} d(p_i, q_j)}{\sum_{i=1}^m \sum_{j=1}^n f_{i,j}} \tag{2}$$

4.2. View 1: Measuring Control–Flow Similarity

This section introduces the approach used to calculate similarity for the control–flow relations of activity labels. Redundant activity labels should share similar control–flow relations or ordering patterns. This similarity indicates not only identical control–flow relations but also closed distribution patterns. As shown in Figure 4, the overall idea behind a control–flow view is that, for each pair of activity labels $(a_i, a_j) \in A$, we adopt EMD to compare the directly-follows and the strong indirectly-follows relations along with their frequency distributions. Each directly-follows and strongly-indirectly follows comparison can be further divided into directly and indirectly outgoing (i.e., consequence) and incoming (i.e., precedent) relations. Thus, we obtain four different values, with the final similarity being the weighted average of these values.

The control–flow view is separated by directly-follows and strong indirectly-follows comparisons. We would like to place most of our effort on explaining the directly-follows comparison, because the strong indirectly-follows comparison follows the same algorithm, only the relations are strong indirectly follows. The reason we also consider strong indirectly-follows relations is to handle non-free-choice problems (i.e., whether we choose a task is dependent on what has been executed in the prior process [41]). For instance, both activities C and D in Figure 3 have identical directly-follows relations, but $D \gg \gg_W G$ (i.e., the dashed line) also exists. Thus, C and D should not be regarded as redundant.

Algorithm 1 presents our approach for calculating the directly-follows similarity. The starting point is to construct a directly-follows graph obtained from the event log (Line 1). For each activity label, we then calculate its outgoing and incoming activity sets (Line 2–4). By using Equation (3), the weights are calculated for each element in the activity set (Line 5–6), (e.g., $A \bullet_W = \{\frac{1}{2}, \frac{1}{2}\}$). Afterwards, for each pair of activity labels, we adopt EMD to calculate the similarity between incoming and outgoing activity sets by using the ground distance function D_{cf} from Equation (4) (Line 7–9). The activities in the sets (e.g., $a \bullet$) are clusters. The weights in the sets (e.g., $a \bullet_W$) are the associated weights for each cluster. For instance, suppose we would like to calculate the similarity between outgoing activity sets for H and B in Figure 3; the input signatures for EMD would be $P = \{(C, 0.46), (D, 0.48), (F, 0.02)\}$ and $Q = \{(C, 0.5), (D, 0.5)\}$. Lastly, the directly incoming and outgoing similarities are averaged to obtain the final directly-follows similarity for each pair of activity labels and added to the set S_d (Line 10–11).

The equation for calculating the weight of a single activity in incoming/outgoing activity sets is defined as:

$$W = \frac{|(b, a)|}{\sum\{|(\epsilon, a)| \epsilon \in \bullet a\}} \quad \text{or} \quad \frac{|(a, b)|}{\sum\{|(a, \epsilon)| \epsilon \in a \bullet\}} \tag{3}$$

Algorithm 1: Directly-Follows Similarity

Input: Event log L
Output: S_d : Set of Directly Similarities for All Pairs of Activities

- 1 $G \leftarrow \text{MakeDirectlyFollowsGraph}(L)$;
- 2 **foreach** $a \in A$ **do**
- 3 $a\bullet \leftarrow \text{OutGoing}(a)$;
- 4 $\bullet a \leftarrow \text{InComing}(a)$;
- 5 $a\bullet_W \leftarrow \text{CalculateWeight}(a\bullet)$;
- 6 $\bullet a_W \leftarrow \text{CalculateWeight}(\bullet a)$;
- 7 **foreach** $a, b \in A$ **do**
- 8 Outgoing Similarity = $\text{EMD}(a\bullet, a\bullet_W, b\bullet, b\bullet_W, D_{cf})$;
- 9 Incoming Similarity = $\text{EMD}(\bullet a, \bullet a_W, \bullet b, \bullet b_W, D_{cf})$;
- 10 Directly Follows Similarity = $\text{Average}(\text{Outgoing Similarity}, \text{Incoming Similarity})$;
- 11 $S_d \leftarrow \text{Directly Follows Similarity}$;
- 12 **return** S_d

The ground distance function D_{cf} for EMD between any two clusters p_i, q_j from activity sets is defined as:

$$D_{cf} = \begin{cases} 0 & \text{if } p_i = q_j \\ 1 & \text{otherwise} \end{cases} \quad (4)$$

Principle. The same activity label has no cost, and different ones have a unit cost. This cost function can be easily extended based on other matrices, e.g., global location for activity labels. Here, we just show the most basic version for better undesirability.

The same algorithm applies to the calculation of indirectly incoming and outgoing similarities. We construct a strong indirectly-follows graph from the event log. We have a set that contains strong indirectly-follows similarities as well. Thus, for each pair of activity labels, the overall control-flow similarity is the weighted average of directly- and strong indirectly-follows similarities, where there is a value between 0 and 1. The greater the value is, the more significant the effort needed to transfer one distribution to another is, which means that the two activity labels have less similarity with regard to the control-flow perspective. The combination of four different scores can be easily extended with other statistical or clustering algorithms, e.g., k-means clustering. We seek to show that our approach can achieve desirable results with the most fundamental method and requires no domain knowledge, e.g., the number of clusters, as the input.

4.3. View 2: Measuring Attribute Value Similarity

This section introduces the approach used to calculate similarity for the attribute values of activity labels. In an event log, attribute values can be the resources needed for executions (e.g., the person who performed the event) or the associated recorded values when executing the event (e.g., the result value for the event). Redundant activity labels share the same attribute values. Therefore, the proposed framework should incorporate attribute values (i.e., both categorical and numerical values) when detecting redundant activity labels.

The overall approach, shown in Figure 4, can be divided into two parts: activity labels with categorical values and numerical values. Activity labels with categorical values are relatively easy to calculate. We calculate the frequency distribution of the attribute values for each activity label and apply pairwise EMD between different activity labels. The ground distance function borrows the idea from Equation (4): the same attribute value has no cost, and different ones have a unit cost. For activity labels with numerical values, we firstly cluster each activity label into different clusters based on value percentiles. We then apply EMD to assess the data distributions of activity labels within each cluster. Clustering first ensures that only activity labels with the same data range are further

evaluated. Activities with different data ranges are unlikely to share similar data patterns, which do not need to be further assessed for data distributions.

We describe how our approach calculates attribute value similarity for activity labels with numerical values in Algorithm 2. For each activity, we first assess whether this activity has a numerical value attribute (Line 2). If not, it has minimal attribute value similarity with other activity labels (i.e., $AttributeValueSimilarity = 1$). If yes, Line 3 finds all events of the activity ($\#_{activity}(e) = a$) and obtains numerical values for that attribute ($\#_{attribute}(e)$) into a dataset (i.e., $DataSet_a$). Line 4 calculates the 25th and 75th percentiles for each dataset. We use the 25th and 75th percentiles as a 2-D vector and apply agglomerative hierarchical clustering [42] with a threshold θ_a for all datasets (Line 5). For example, if an activity label has the values 0.3 and 0.5 for its 25th and 75th data value percentiles, respectively, then the 2-D vector is (0.3, 0.5). We apply the Euclidean distance [43] as the distance measurement between two vectors. Activity labels that are not in the same cluster also have $AttributeValueSimilarity = 1$. There are many unique values in the attribute; it is hard to directly apply EMD because of the many different clusters in the distribution. As a result, we transfer each dataset to a histogram following Sturges' formula [44], where uniform maximum and minimum values are used to ensure two histograms have the same bin number and size when comparing activity label pairs within the same cluster (Line 11). We pick each interval's left boundary as cluster values (e.g., p_i, q_j), and Line 12–13 calculate the percentage of each bin as cluster weights (e.g., w_{p_i}, w_{q_j}). An example cluster is (10, 20%), (15, 30%), (20, 20%), and (25, 30%). In this way, we transfer each histogram as a cluster, and EMD is further used to compare two clusters using the distance function D_d in Equation (5) (Line 14). The $AttributeValueSimilarity$ is normalised to become a value between 0 and 1 and added to S_n . Similar to the control-flow perspective, in the attribute value perspective, the greater the value is, the less similarity they have.

Algorithm 2: Attribute Value Similarity

Input: Event log L , threshold θ_a
Output: S_n : Set of Attribute Value Similarities for all Pairs of Activities

```

1 foreach  $a \in A$  do
2   if  $HasNumericalValueAttribute(a)$  then
3      $DataSet_a \leftarrow ExtractAttribute(a)$ ;
4      $Q_{1a}, Q_{3a} \leftarrow CalculatePercentiles(a)$ ;
5    $Clusters \leftarrow AgglomerativeHierarchicalClustering(Q_{1a}, Q_{3a}, \theta_a)$ ;
6   foreach  $C \in Clusters$  do
7     if  $size(C) \leq 1$  then
8       continue;
9     else
10      foreach  $a, b \in C$  do
11         $H_a, H_b \leftarrow MakeHistograms(DataSet_a, DataSet_b)$ ;
12         $H_{aW} \leftarrow CalculateWeight(H_a)$ ;
13         $H_{bW} \leftarrow CalculateWeight(H_b)$ ;
14         $Attribute\ Value\ Similarity = EMD(H_a, H_{aW}, H_b, H_{bW}, D_d)$ ;
15         $S_n \leftarrow Attribute\ Value\ Similarity$ ;
16 return  $S_n$ 

```

The ground distance function D_d for EMD between any two attribute clusters p_i, q_j from histograms is defined as

$$D_d = |p_j - q_j|. \quad (5)$$

Principle. Because both p_i, q_j are numerical values, it takes less effort to transfer p_i to q_j if they are close to each other. We adopt the difference between p_i and q_j as the ground distance function.

4.4. View 3: Measuring Semantic Similarity

We are interested in measuring the semantic similarity between activity labels, because redundant activity labels may share close semantics. Although only looking at the semantics of activity labels may lead to false-positive results (i.e., labels that are incorrectly detected as redundant), it is still an important factor to consider when detecting redundant labels in the multi-view framework. There are many methods proposed for assessing semantic similarity between words, such as the string edit distance [45]. However, the string edit distance cannot easily handle synonymous activity labels with a different wording structure. Therefore, we integrate NLP into our multi-view framework. We apply a pre-trained NLP model Spacy, which is a well-known industrial-strength NLP tool that provides fast and accurate syntactic analyses, to assess the semantic similarity between every pair of activity labels [46]. The result between a pair of activity labels is a numerical number ranging from 0 to 1, where 1 means that they are identical and vice versa. However, in order to comply with the same rule setting in previous sections, we subtract the obtained results from 1. In this way, 0 means that they are identical in semantic similarity.

4.5. Decision-Making Mechanism: Majority Voting

For now, each pair of activity labels has three views of similarities: control–flow relations, attribute values, and label semantics. This section describes a decision-making mechanism to aggregate similarities from the above three views and generates final results, as shown in Figure 4. The decision-making mechanism is the majority voting, which is a widely adopted concept in ensemble learning [47]. For majority voting, we first need to determine the threshold for each view (i.e., θ_c , θ_d , and θ_s) to decide whether each pair of activity labels is similar in the corresponding view. Equation (6) describes the voting mechanism, where V_i represents the result from a particular view, and m represents the total number of views (i.e., $m = 3$ in this paper). The activity labels are detected as redundant pair if more than half of the views are regarded as similar and vice versa. Using the majority voting as the decision-making mechanism has the following advantages: (1) it is fast and easy to implement, and requires no domain knowledge as input, (2) it can be easily extended if more views are proposed for redundant activity label detection, and (3) it can be easily integrated with domain knowledge (e.g., the voting weight of each view).

$$Result = \begin{cases} \text{"Redundant"} & \text{if } |i : V_i = \text{"Similar"}| \geq m/2 \\ \text{"Non-redundant"} & \text{otherwise} \end{cases} \quad (6)$$

The time complexity of the framework depends on the time complexity of computing each view. Suppose n, m, k are the number of activities, events, and attributes values. $O(n \times (m + n))$, $O(n \times m \times k)$, and $O(n)$ for control–flow, attribute value and semantic views respectively. The final decision-making mechanism is also in linear time.

5. Evaluation

We conducted a large number of experiments to prove that our proposed framework can accurately detect redundant activity labels in event logs. Overall, two groups of experiments were performed. The first group of experiments compared the performance of our framework with the existing state-of-the-art method to detect redundant activity labels. The second group of experiments further analysed the effectiveness of our framework.

To evaluate the performance of our framework, we apply the same evaluation matrix as found in [11], which is the standard f-score metric. Detection results fall into one of the following four categories: (1) true-positive (TP), where positive outcomes are correctly detected (the detected redundant activity label is actual redundant), (2) false-positive (FP), where negative outcomes are detected as positive (the detected redundant activity label is actually not redundant), (3) true-negative (TN), where negative outcomes are correctly detected (the detected non-redundant activity label is actually not redundant), and (4) false-negative (FN), where negative outcomes are falsely detected (the actual redundant

activity label is detected as non-redundant). Based on these four indicators, we calculate the following evaluation scores. Precision defines how many positive classes are detected correctly out of all positive detection.

$$Precision = \frac{TP}{TP + FP} \quad (7)$$

Recall indicates how many from all positive classes are detected correctly.

$$Recall = \frac{TP}{TP + FN} \quad (8)$$

The f-score is the harmonic mean of precision and recall.

$$F\text{-score} = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (9)$$

Several publicly available datasets were utilised in the experiments. In total, our evaluation was based on four publicly available event logs:

- Hospital Billing log (<https://doi.org/10.4121/uuid:76c46b83-c930-4798-a1c9-4be94dfcb741> (accessed 7 April 2022)): An event log records processes related to billing medical services provided by a Dutch hospital.
- Sepsis log (<https://doi.org/10.4121/uuid:915d2bfb-7e84-49ad-a286-dc35f063a460> (accessed 7 April 2022)): An event log records treatment processes of sepsis patients from a Dutch hospital.
- Helpdesk (<https://doi.org/10.4121/uuid:0c60edf1-6f83-4e75-9367-4c63b3e9d5bb> (accessed 7 April 2022)): An event log contains the ticketing management process of the help desk in a software company in Italy.
- BPI Challenge 2012 (<https://doi.org/10.4121/uuid:3926db30-f712-4394-aebc-75976070e91f> (accessed 7 April 2022)): An event log of a loan application process in a Dutch financial institute.

The details of all used event logs are presented in Table 1. Following the data preparation methods in [11], we randomly selected a certain amount of activity labels and randomly varied the percentage (i.e., 1% to 30%) of its events to simulate activity labels with both low-occurrence and high-occurrence frequency. In total, seven different settings were used on four event logs. For each setting, 5 rounds were performed, and the average results are reported. It is worth mentioning that the Sepsis event log also contains different variants of discharging a patient, which are “Release C”, “Release D”, and “Release E”. They are regarded as redundant [11,17]. Thus, the ground truth not only contains activity labels that we manually changed, but also consists of any pair of these three activity labels.

Table 1. Characteristics of event logs used for evaluations.

Event Log	Number of Traces	Number of Trace Variants	Number of Events	Number of Attributes	Number of Activity Labels
Hospital Billing	100,000	1020	451,359	1105	18
Sepsis	1050	846	15,214	26	16
Helpdesk	4580	226	21,348	22	14
BPI Challenge 2012	13,087	4366	262,200	69	24

In real-life situations, activity labels that are similar in any two aspects will be regarded as redundant according to the majority voting mechanism. Because the data quality on the activity label is a unique problem in process mining, even though the redundant activity labels are similar in semantics, they still need to be regarded as similar in terms of the control-flow relation or the attribute value. Hence, we are more interested in evaluating whether our framework can detect redundant activity labels that are different in semantics, because the effectiveness of the first two features is critical in the proposed

framework. Therefore, we artificially renamed the activity labels arbitrarily following [11] and intentionally set the voting weight of semantic similarity to 0 when performing the experiments. Hence, activity labels that are similar in both two aspects were regarded as redundant in the experiments. We will further demonstrate the use of semantic similarity in the real-life case study proposed in Section 6. The proposed framework was implemented as a Python program for evaluations. We adopted $\theta_c = 0.25$ and $\theta_d = 0.1$ for each aspect.

5.1. Comparing with The Existing Method

Currently, three state-of-the-art approaches are proposed for detecting redundant activity labels in event logs [11,13,14]. However, the approaches in [13,14] are an interactive detection approach, which requires domain experts. Hence, for the baseline approach, we selected the SynonymousLabelRepair [11], which seems to be more advanced in handling redundant activity labels and requires less domain knowledge than the other methods. We used its default settings in evaluations, which are a 0.7 threshold and a uniform weight, as the optimal weight settings required extra domain knowledge.

We followed the evaluation pipeline proposed in [11] for the first group of experiments, which aimed at comparing the performance of our proposed framework with the existing method to detect redundant activity labels in event logs, as shown in Figure 5. The “Hospital Billing” and “Sepsis” event logs were evaluated by SynonymousLabelRepair in [11]. In total, 70 event logs were generated and evaluated in the experiment. Hence, in this section, we compare our framework with the SynonymousLabelRepair by using the above event logs.

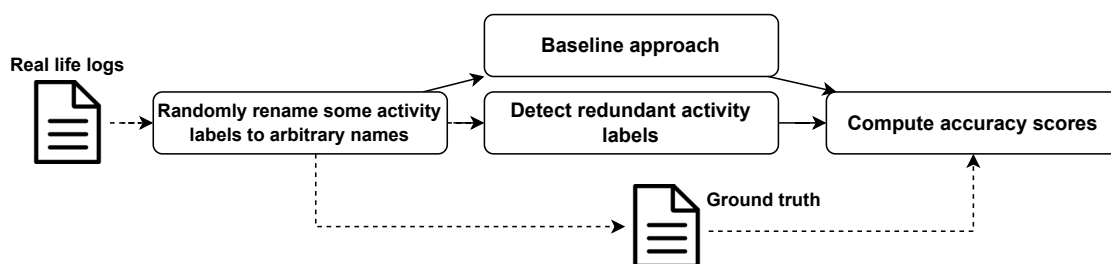


Figure 5. Overview of the evaluation pipeline.

Table 2 along with Figure 6 and 7 show the results for the evaluation. It is clear that our framework outperforms the baseline approach. The f-scores of our framework are all above 0.8 for the event logs under different settings, which indicates that almost all redundant activity labels can be successfully detected regardless of their occurrence frequencies in event logs. On the contrary, the average f-scores for the baseline were merely 0.33 and 0.43 in the experiments. Moreover, the baseline performs poorly when the redundant activity labels are less frequent. For instance, the baseline only achieves a 0.08 f-score when there are 1% redundant activity labels in the “Hospital Billing” event log. We also noticed that the baseline begins to catch up and even surpasses our framework in recall when redundant activity labels become more frequent in the “Hospital Billing” event log, which shows that both approaches can successfully detect most of the redundant activity labels. However, the low precision values indicate that the baseline constantly produces false-positive results (i.e., non-redundant activity labels are falsely detected as redundant), which further proves the necessity and effectiveness of our multi-view framework.

The baseline approach only compares directly-follows relations while ignoring their frequency distributions and indirectly-follows relations. However, low-frequency activity labels rarely contain all directly-follows relations while only maintaining the main one. In this case, frequency distributions of the control-flow relations are becoming important in the approach. We also notice that the control-flow view is limited when handling activity labels that share XOR relations, such that they have identical incoming and outgoing relations. This further illustrates the necessity to also consider other views when detecting the redundant activity labels. Furthermore, the cost function between different activity

labels in EMD can be better defined instead of simply adopting the unit cost if domain knowledge is available. Thus, more satisfying results can be potentially achieved by our framework. Additionally, the baseline relies on a probability density function to assess value distributions for activity attributes. However, distributions of attribute values are less structured when redundant activity labels are infrequent and mixed with categorical and numerical values.

Table 2. Comparison of our framework with SynonymousLabelRepair .

Event Log	Number of Redundant Activity Labels	Precision		Recall		F-score	
		Ours	Baseline	Ours	Baseline	Ours	Baseline
Hospital Billing	5254 (1%)	0.97	0.05	0.89	0.20	0.93	0.08
	21,693 (5%)	0.94	0.17	0.90	0.73	0.92	0.28
	44,890 (10%)	0.88	0.5	0.80	0.18	0.84	0.26
	66,368 (15%)	0.85	0.19	0.89	0.90	0.87	0.31
	90,273 (20%)	0.87	0.24	0.86	0.92	0.86	0.38
	112,840 (25%)	0.80	0.30	0.94	0.85	0.86	0.44
	135,480 (30%)	0.80	0.42	0.96	0.95	0.87	0.58
Sepsis	180 (1%)	0.76	0.39	0.90	0.23	0.82	0.29
	745 (5%)	0.75	0.47	0.89	0.42	0.81	0.44
	1569 (10%)	0.93	0.52	0.77	0.45	0.84	0.47
	2327 (15%)	0.93	0.33	0.71	0.25	0.81	0.29
	3086 (20%)	0.90	0.48	0.76	0.46	0.83	0.47
	3844 (25%)	0.80	0.55	0.85	0.49	0.82	0.51
	4605 (30%)	0.86	0.52	0.77	0.58	0.82	0.55

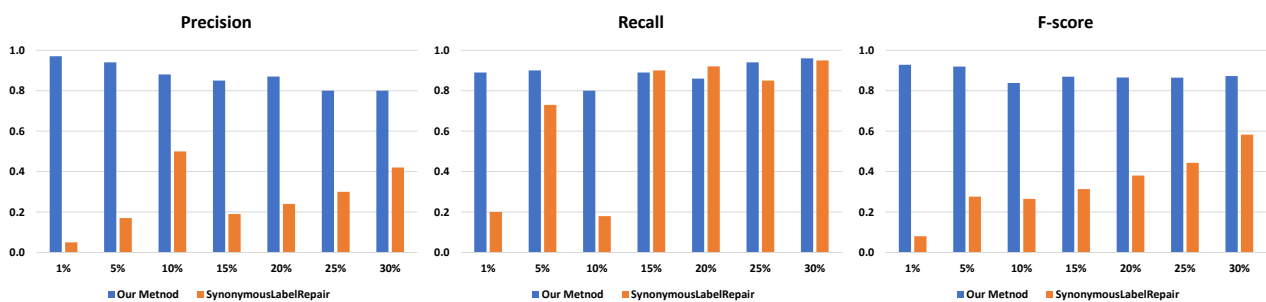


Figure 6. Comparison of our framework with the baseline using the Hospital Billing event log.

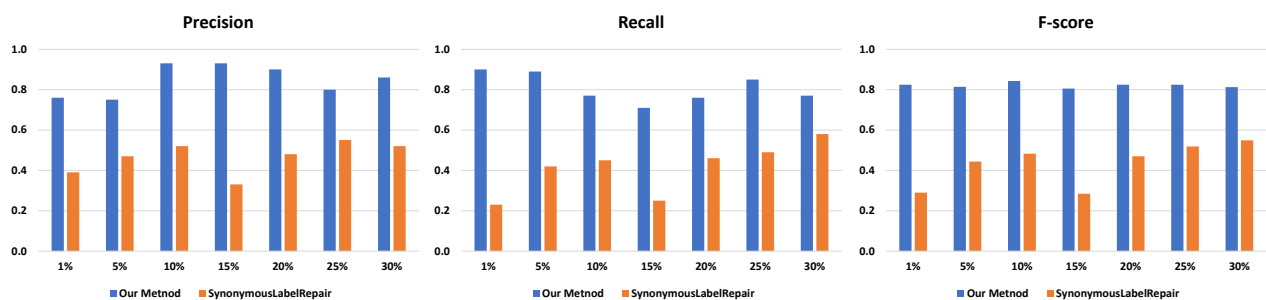


Figure 7. Comparison of our framework with the baseline using the Sepsis event log.

5.2. Further Analysis of Our Proposed Framework

To further analyze the performance of the proposed framework, more evaluations were performed with more event logs and redundant activity labels. Apart from running experiments on our multi-view framework, we also implemented two baselines to prove the effectiveness of our framework:

- Control–Flow Only: the baseline only relies on the control–flow similarity to detect redundant activity labels.
- Attribute Value Only: the baseline only relies on the attribute similarity to detect redundant activity labels.

In total, all four event logs are used to evaluate our framework. For each event log, we randomly rename 1%, 5%, 10%, 15%, 20%, 25%, and 30% of the activity labels. Table 3 and Figure 8 present the f-score comparisons between the baselines and our framework. All f-scores were averaged by 5 repeats. In total, 140 event logs were generated and evaluated in the experiment. The f-scores of our framework are much higher than the two baselines, which only detect from a single perspective. The results suggest that it is important to consider information from multiple views when detecting redundant activity labels. We also notice that, except for the BPI Challenge 2012 event log, the control–flow view usually plays a more important role than the attribute value view.

Table 3. Further analysis of our framework .

Event Log	Number of Redundant Activity Labels	Control–Flow Only	Attribute Value Only	Our Framework
Hospital Billing	5254 (1%)	0.72	0.71	0.93
	21,693 (5%)	0.74	0.69	0.92
	44,890 (10%)	0.70	0.66	0.84
	66,368 (15%)	0.78	0.67	0.87
	90,273 (20%)	0.71	0.66	0.86
	112,840 (25%)	0.69	0.65	0.86
	135,480 (30%)	0.67	0.63	0.87
Sepsis	180 (1%)	0.66	0.60	0.82
	745 (5%)	0.63	0.55	0.81
	1569 (10%)	0.70	0.67	0.84
	2327 (15%)	0.65	0.58	0.81
	3086 (20%)	0.71	0.58	0.83
	3844 (25%)	0.69	0.66	0.82
	4605 (30%)	0.72	0.61	0.82
Helpdesk	213 (1%)	0.73	0.66	0.92
	1067 (5%)	0.70	0.62	0.89
	2135 (10%)	0.77	0.58	0.90
	3202 (15%)	0.71	0.60	0.87
	4270 (20%)	0.71	0.62	0.88
	5337 (25%)	0.65	0.70	0.85
	6404 (30%)	0.68	0.65	0.84
BPI Challenge 2012	2622 (1%)	0.58	0.65	0.86
	13,110 (5%)	0.60	0.71	0.88
	26,220 (10%)	0.65	0.63	0.83
	39,330 (15%)	0.60	0.61	0.84
	52,440 (20%)	0.70	0.68	0.88
	65,550 (25%)	0.68	0.66	0.85
	78,660 (30%)	0.63	0.63	0.82

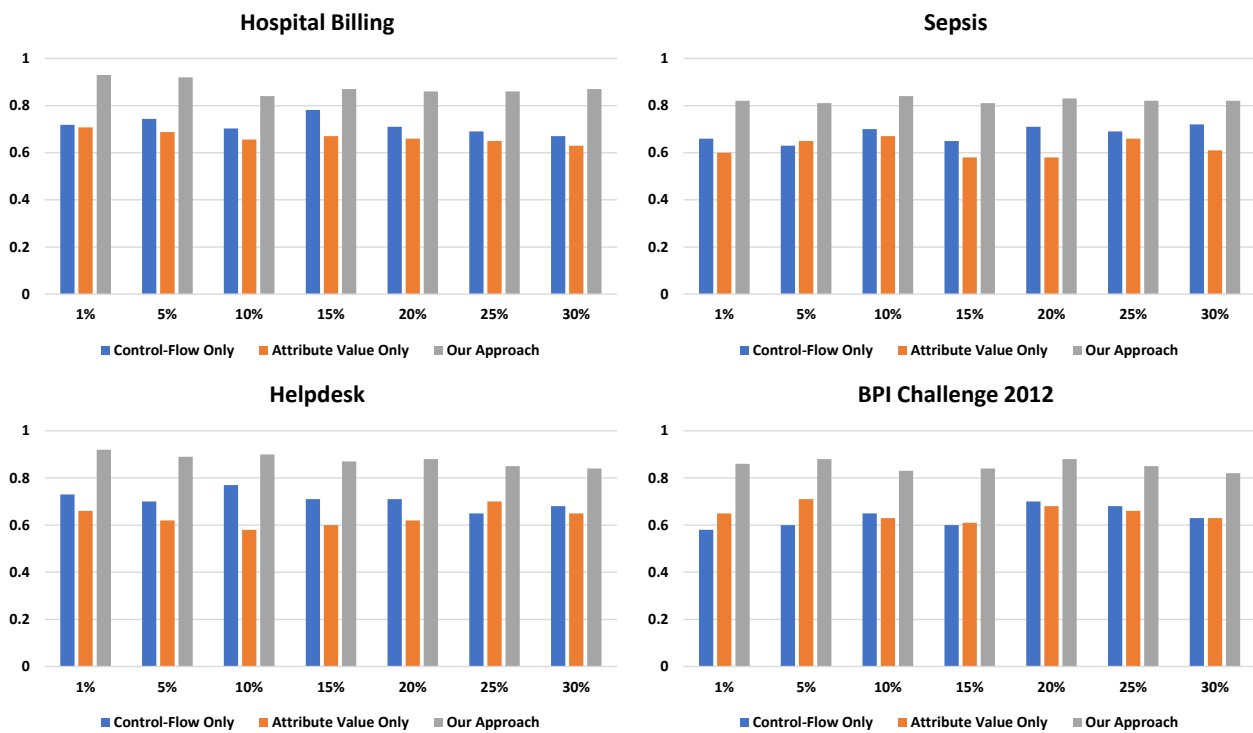


Figure 8. F-score comparisons of our framework with the baseline using four event logs.

Finally, we also conducted experiments to show how redundant activity labels can impact the discovered process models and how our framework can be used as a preprocessing tool to repair event logs. We selected two event logs with the highest (Hospital Billing) and lowest (Sepsis Log) detection f-score for the experiments. A simple and fundamental repair mechanism was adopted. Redundant activity labels were replaced with the most similar non-redundant activity labels in event logs. A mode of advanced repair technology is left as future work. We aimed to compare the f-scores for discovered process models on original event logs (no redundant activity labels exist), problematic event logs (with redundant activity labels), and repaired event logs by using the proposed detection framework with the basic repair mechanism. We used the Inductive Miner infrequent [21] to mine process models. We conducted conformance checking between all discovered process models against the original event logs. The more conformance there was between the process model discovered from the repaired event log and the original event log, the higher the f-score achieved. We used alignment-based fitness and conformance-based precision tools in the PM4PY [48]. The f-scores for the original event logs are referenced from [15] directly. The results are presented in Table 4. We can see that the existence of redundant activity labels has a tremendous impact on the discovered process models. For instance, the f-scores drop significantly compared with the original f-scores. On the contrary, the f-scores on process models discovered from repaired event logs drop slower and remain closer to the original event log. The results indicate that our framework can be used as a preprocessing tool and help with improving the quality of process models when there are redundant activity labels in the event logs.

Table 4. F-score comparisons between process models discovered from the original, problematic, and repaired event logs.

Event Log	Number of Redundant Activity Labels	F-Score (the Original Log)	Average F-Score (the Logs with Redundant Activity Labels)	Average F-Score (the Repaired Logs)
Hospital Billing	5254 (1%)	0.75	0.69	0.73
	21,693 (5%)	0.75	0.64	0.72
	44,890 (10%)	0.75	0.62	0.70
	66,368 (15%)	0.75	0.59	0.69
	90,273 (20%)	0.75	0.55	0.67
	112,840 (25%)	0.75	0.53	0.67
	135,480 (30%)	0.75	0.51	0.66
Sepsis	180 (1%)	0.77	0.74	0.75
	745 (5%)	0.77	0.72	0.74
	1569 (10%)	0.77	0.66	0.72
	2327 (15%)	0.77	0.64	0.70
	3086 (20%)	0.77	0.61	0.70
	3844 (25%)	0.77	0.60	0.68
	4605 (30%)	0.77	0.57	0.67

6. Real-Life Case Study

We conducted a case study using the Speed-Extract EMR dataset to demonstrate that our framework can be used in real-life healthcare event logs. The Speed-Extract dataset comprises retrospective data from a historical dataset extracted between 2013 and 2018 from a single Cerner Millennium EMR domain in Sydney, Australia [18]. The Speed-Extract dataset comprises anonymised patients that presented with suspected acute coronary syndrome (ACS) to facilities in Northern Sydney LHD and Central Coast LHD [18]. We aimed to study the treatment process of ST-elevation myocardial infarction (STEMI) patients, which is a type of heart attack that mainly affects the patient’s heart’s lower chambers [18]. In this paper, we demonstrated how we extract the event log from the Speed-Extract dataset and applied the proposed framework to improve the quality of the event log. We verified and substituted the detected redundant activity labels by using the domain knowledge to obtain a more representative event log. In the end, we applied the existing tool to mine the process models from the two event logs. Comparisons between these two models demonstrate that our proposed framework can be used as a preprocessing method for the event log to obtain a more structured and easier-to-understand process model. The pipeline for the case study is shown in Figure 9.

The data we used in the Speed-Extract dataset includes the following tables: the *Patient-prepr* table includes patients’ demographics, such as age and gender; the *Medications-mapped* table records prescription orders for each patient; the *Diagnosis-prepr* table contains the diagnosis for each patient using the International Classification of Diseases (ICD)-10 codes.

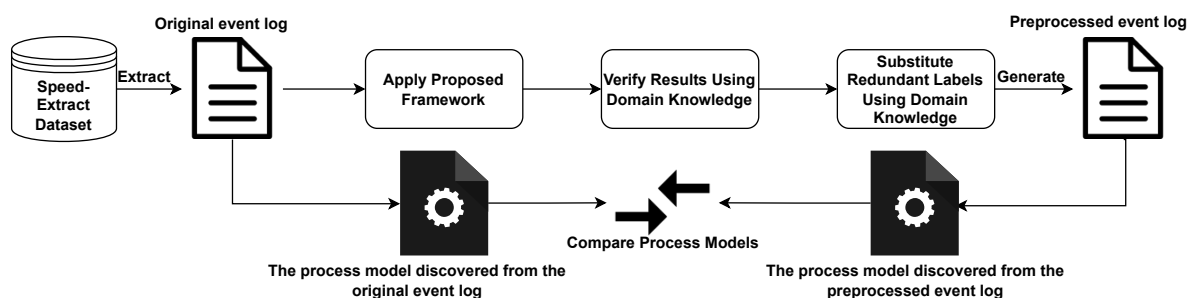


Figure 9. Overview of the case study pipeline.

6.1. Event Log Construction

STEMI patients were defined by using the ICD-10 codes, resulting in 5750 patients. Patients older than 85 or younger than 40 were excluded [49]. We treated patients as traces and their medications as activities to construct the event log. After extracting the initial event log, we filtered out traces where less than one activity was presented, because they exposed less useful information regarding control–flow similarity. As a result, 2141 traces with 615 activities are presented in the event log. The event log has 1363 trace variants due to the high number of activities. The primary reason for such an enormous number of activities is the existence of redundancy in the event log. When giving prescriptions to patients, the free-text nature causes redundancy. For example, when referring to the same medicine, some doctors prefer to write the exact medication name, while others prefer the brand name (e.g., *Telmisartan* is the medication name sold under the brand name *Micardis*). Furthermore, there are many different substitutable medications to choose from when addressing the same symptom, which is one of the other reasons for redundant activity labels in the event log. For instance, both *Candesartan* and *Eprosartan* can be used to treat high blood pressure. Moreover, some doctors prefer giving dose and frequency along with the medication name in one input field, such as *Aspirin one tablet per day*, which also introduces unnecessary redundancy to the event log.

6.2. Result and Discussion

Because the goal is to study the treatment process, medications with the same effect shared the same behaviour in the treatment process. Hence, we applied the proposed framework to detect redundancy in the event log and generate a more representative event log for further analysis. In this case study, the thresholds were set to 0.2 for the control–flow relation, 0.1 for the attribute values, and 0.1 for the label semantic information. As a result, nine groups of redundant activity labels were detected according to the transitive property of equality. The results were further evaluated by the domain knowledge to show that these are different therapies in the STEMI treatment [50]. Therefore, we adopted therapy names (e.g., *Beta Blockers*, *Anticoagulants*) to replace redundant activity labels in the original event log to generate the preprocessed event log. Compared with the original event log, the number of activity labels in the process was reduced from about 600 to 13. The number of distinct trace variants was reduced from about 1300 to 660, which is around half of the original event log.

We adopted Disco for process discovery on the event logs [16]. As a result, a spaghetti-like process model was discovered from the original event log due to the redundancy. Figure 10 shows a snapshot of the partial process model because the original process model is too complicated to display here. The original process model contains over 600 activity labels and countless relations between them. Such a process model cannot provide any useful insights toward understanding the treatment process for STEMI patients. On the contrary, Figure 11 shows the process model discovered by using Disco from the preprocessed event log after applying the proposed framework in the paper. The discovered process model in Figure 11 overall seems to be simple and insightful, representing only the important treatment behaviours of STEMI patients. We can observe that most patients were treated with *Aspirin*, *Beta blockers*, and *Statins* during their hospital stays, which is in line with the current treatment guideline [50]. Moreover, several control–flow relations are also presented in the process model, such as *Clopidogrel* $>_W$ *Digoxin* and *Insulin* $>_W$ *Statins*. Such a process model together with additional patients' records can be further used to study other related topics, such as how the treatment process influences a patient's survival time. This case study further shows that our framework can efficiently detect redundant activity labels in the event log. Moreover, the event log with redundant activity labels has a significant impact on the quality of the discovered process model.

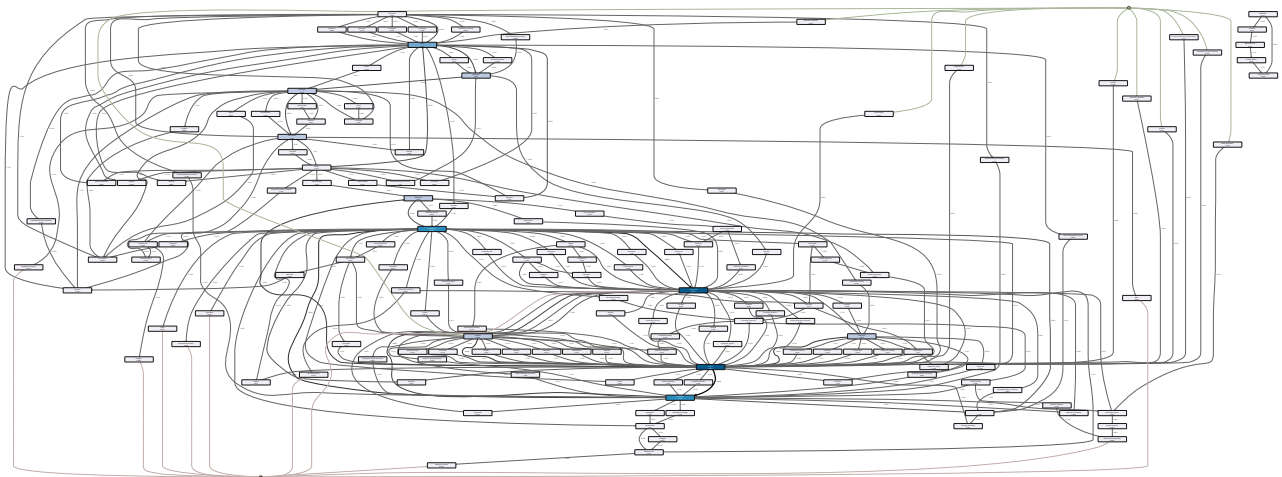


Figure 10. A snapshot of the partial process model discovered from the original event log.

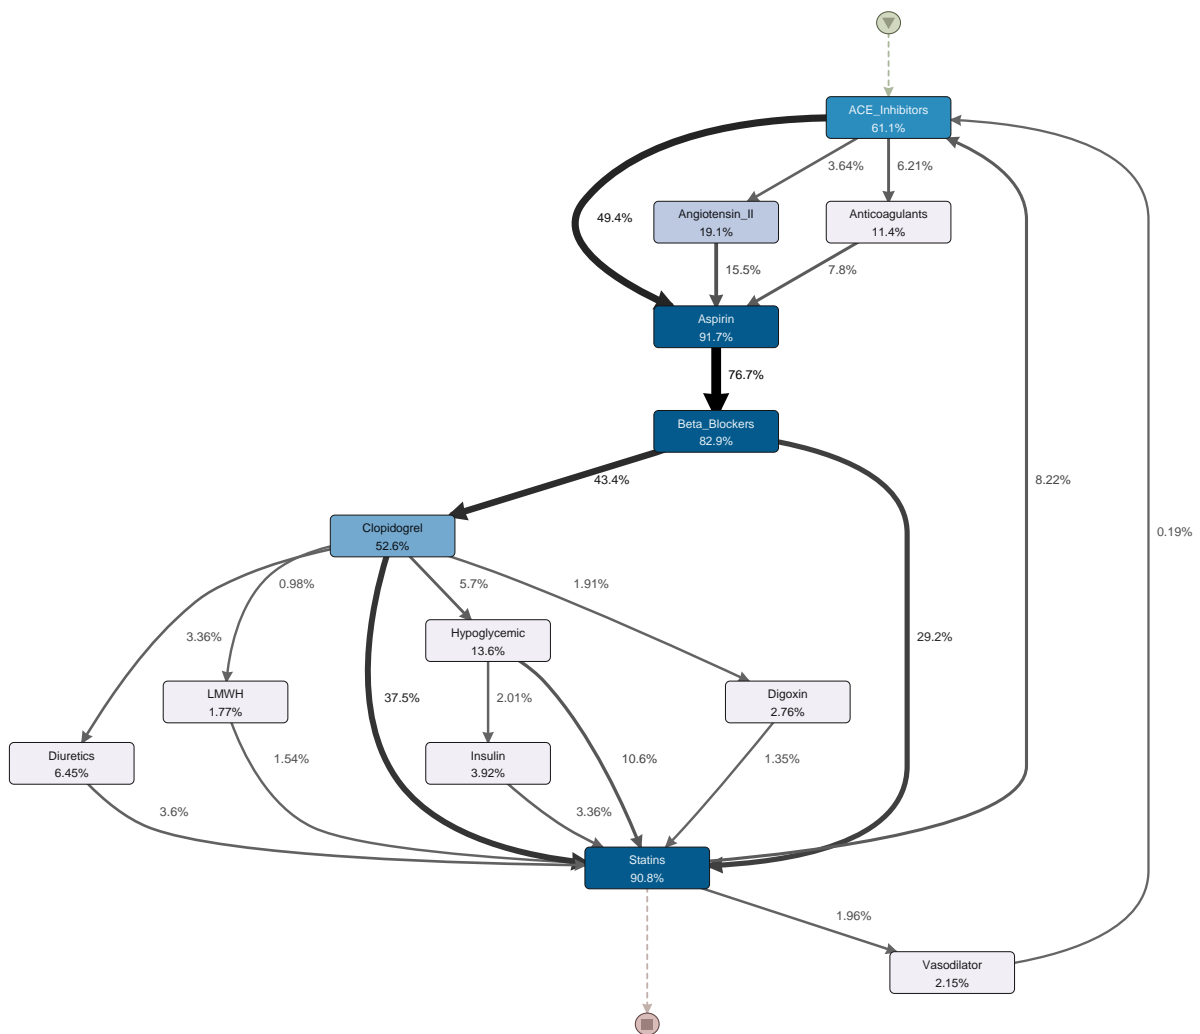


Figure 11. The process model discovered from the preprocessed event log.

7. Conclusions and Future Work

Existing process discovery algorithms assume that each activity label in the event log is unique and meaningful. The existence of redundant activity labels introduces unnecessary

complexities to the discovered process models, which leads to spaghetti-like models. Hence, this paper proposes a multi-view framework to accurately detect redundant activity labels and produce more representative event logs for process mining. Our framework considers information from different views (i.e., control-flow relations, attribute values, and label semantic information) when detecting redundant activity labels. A consensus is made through the majority voting mechanism. The results are superior to those of the existing method in terms of detecting redundant activity labels. The usability of our framework is further demonstrated by using a real-life case study based on Speed-Extract EMR dataset.

Of course, the study described in this paper also comes with its very own limitations. It has to be noted that, like other data preprocessing approaches, the detection results still vary under different parameter settings in different logs. The adopted NLP model is not designed explicitly for process mining tasks, which may impact the performance of the framework.

In future work, first, we plan to investigate how different parameter settings can impact the framework and to develop a method to automatically determine thresholds for different views. Secondly, we also aim to incorporate the NLP technique to automatically repair redundant activity labels by preserving the same contexts and by categorising differences according to their closest synonyms. Finally, we plan to investigate the feasibility of applying our framework in other domains.

Author Contributions: Conceptualisation and methodology, Q.C., Y.L., C.S.T., and S.K.P.; development, Q.C.; validation, Q.C., Y.L., C.S.T., and S.K.P.; writing—original draft preparation, Q.C.; writing—review and editing, Q.C., Y.L., C.S.T., and S.K.P.; supervision, C.S.T. and S.K.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Ethics approval for the study (2019/ETH09692) was provided by the Northern Sydney Local Health District (NSLHD) Human Research Ethics Committee (HREC). Governance approvals were provided by the NSLHD and Central Coast Local Health District HRECs.

Informed Consent Statement: We received a waiver of consent for the patients in this study, which was approved by the NSLHD Human Research Ethics Committee. Only de-identified data were used by the researchers for this analysis, so a waiver of consent was appropriate.

Data Availability Statement: All datasets used in Section 5 to evaluate the proposed framework are publicly available. Please refer to notes for links to access the datasets. The datasets generated and/or analysed in Section 6 are not publicly available, as they are owned by the Chief Executives of the LHDs and not by the researchers who performed this study.

Acknowledgments: The authors thank the members of the Speed-Extract research team and the technical assistance provided by the Sydney Informatics Hub, a Core Research Facility of the University of Sydney. The authors also acknowledge ICT services at NSLHD. The authors also thank the Agency for Clinical Innovation, the NSW Ministry of Health, the Sydney Health Partners, and Health Innovation Funding for funding the Speed-Extract study.

Conflicts of Interest: The authors declare that there are no conflict of interest.

Abbreviations

The following abbreviations are used in this paper:

EMR	Electronic Medical Records
MIMIC-III	Medical Information Mart for Intensive Care-III
EMD	Earth Mover's Distance
LHD	Local Health District
NLP	Natural Language Processing
ACS	Acute Coronary Syndrome
STEMI	ST-elevation Myocardial Infarction
ICD	International Classification of Diseases

Appendix A. Notations

Table A1. Used notations in the paper.

Symbol	Description
L	the event log
t	the trace
A	set of activities
e	the event
$\#_n(e)$	the function that obtains attribute values recorded for an event e
$>_W$	the directly follows relation
$>>_W$	the indirectly follows relation
\Rightarrow_W	the long distance measure
$>>>_W$	the strong indirectly follows relation
G	the directly follows graph
$a \bullet$	the post-set
$\bullet a$	the pre-set
P	the probability distribution
D	the ground distance between clusters
$EMD(P, Q)$	the EMD between two probability distributions

References

1. Van Der Aalst, W. Data science in action. In *Process Mining*; Springer: Berlin/Heidelberg, Germany, 2016; pp. 3–23.
2. Marin-Castro, H.M.; Tello-Leal, E. Event Log Preprocessing for Process Mining: A Review. *Appl. Sci.* **2021**, *11*, 10556. [[CrossRef](#)]
3. Wen, L.; Wang, J.; van der Aalst, W.M.; Huang, B.; Sun, J. Mining process models with prime invisible tasks. *Data Knowl. Eng.* **2010**, *69*, 999–1021. [[CrossRef](#)]
4. Maggi, F.M.; Bose, R.; van der Aalst, W.M. Efficient discovery of understandable declarative process models from event logs. In *Proceedings of the International Conference on Advanced Information Systems Engineering*, Gdansk, Poland, 25–29 June 2012; Springer: Berlin/Heidelberg, Germany, 2012; pp. 270–285.
5. Mans, R.S.; Van der Aalst, W.M.; Vanwersch, R.J. *Process Mining in Healthcare: Evaluating and Exploiting Operational Healthcare Processes*; Springer: Berlin/Heidelberg, Germany, 2015.
6. Weijters, A.; Ribeiro, J. Flexible heuristics miner (FHM). In *Proceedings of the 2011 IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, Paris, France, 11–15 April 2011; pp. 310–317.
7. Augusto, A.; Conforti, R.; Dumas, M.; La Rosa, M.; Polyvyanyy, A. Split miner: Automated discovery of accurate and simple business process models from event logs. *Knowl. Inf. Syst.* **2019**, *59*, 251–284. [[CrossRef](#)]
8. Chen, Q.; Lu, Y.; Tam, C.; Poon, S. Process Mining to Discover and Preserve Infrequent Relations in Event Logs: An Application to Understand the Laboratory Test Ordering Process Using the MIMIC-III Dataset. In *Proceedings of the Australasian Conference on Information Systems (ACIS)*, Sydney, Australia, 6–10 December 2021; pp. 30–41.
9. Van Der Aalst, W.; Adriansyah, A.; De Medeiros, A.K.A.; Arcieri, F.; Baier, T.; Blicke, T.; Bose, J.C.; Van Den Brand, P.; Brandtjen, R.; Buijs, J.; et al. Process mining manifesto. In *Proceedings of the International Conference on Business Process Management*, Clermont-Ferrand, France, 30 August–2 September 2011; Springer: Berlin/Heidelberg, Germany, 2011; pp. 169–194.
10. Suriadi, S.; Andrews, R.; ter Hofstede, A.H.; Wynn, M.T. Event log imperfection patterns for process mining: Towards a systematic approach to cleaning event logs. *Inf. Syst.* **2017**, *64*, 132–150. [[CrossRef](#)]
11. Sadeghianasl, S.; ter Hofstede, A.H.; Wynn, M.T.; Suriadi, S. A contextual approach to detecting synonymous and polluted activity labels in process event logs. In *Proceedings of the OTM Confederated International Conferences On the Move to Meaningful Internet Systems*, Rome, Italy, 10–14 September 2012; Springer: Berlin/Heidelberg, Germany, 2019; pp. 76–94.

12. Johnson, A.E.; Pollard, T.J.; Shen, L.; Li-Wei, H.L.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Celi, L.A.; Mark, R.G. MIMIC-III, a freely accessible critical care database. *Sci. Data* **2016**, *3*, 1–9. [[CrossRef](#)]
13. Sadeghianasl, S.; ter Hofstede, A.H.; Suriadi, S.; Turkay, S. Collaborative and interactive detection and repair of activity labels in process event logs. In Proceedings of the 2020 2nd International Conference on Process Mining (ICPM), Padua, Italy, 5–8 October 2020; pp. 41–48.
14. Sadeghianasl, S.; Ter Hofstede, A.H.; Wynn, M.T.; Turkay, S.; Myers, T. Process Activity Ontology Learning From Event Logs Through Gamification. *IEEE Access* **2021**, *9*, 165865–165880. [[CrossRef](#)]
15. Lu, Y.; Chen, Q.; Poon, S.K. A Deep Learning Approach for Repairing Missing Activity Labels in Event Logs for Process Mining. *Information* **2022**, *13*, 234. [[CrossRef](#)]
16. Günther, C.W.; Rozinat, A. Disco: Discover Your Processes. *BPM (Demos)* **2012**, *940*, 40–44.
17. Mannhardt, F.; Blinde, D. *Analyzing the Trajectories of Patients with Sepsis Using Process Mining*; RADAR+ EMISA@ CAiSE: 2017 Springer: Berlin/Heidelberg, Germany, 2017; pp. 72–80.
18. Tam, C.S.; Gullick, J.; Saavedra, A.; Vernon, S.T.; Figtree, G.A.; Chow, C.K.; Cretikos, M.; Morris, R.W.; William, M.; Morris, J.; et al. Combining structured and unstructured data in EMRs to create clinically-defined EMR-derived cohorts. *BMC Med Inform. Decis. Mak.* **2021**, *21*, 91. [[CrossRef](#)]
19. Van der Aalst, W.; Weijters, T.; Maruster, L. Workflow mining: Discovering process models from event logs. *IEEE Trans. Knowl. Data Eng.* **2004**, *16*, 1128–1142. [[CrossRef](#)]
20. Wen, L.; Van Der Aalst, W.M.; Wang, J.; Sun, J. Mining process models with non-free-choice constructs. *Data Min. Knowl. Discov.* **2007**, *15*, 145–180. [[CrossRef](#)]
21. Leemans, S.J.; Fahland, D.; van der Aalst, W.M. Discovering block-structured process models from event logs containing infrequent behaviour. In Proceedings of the International Conference on Business Process Management, Beijing, China, 26–30 August 2013; Springer: Berlin/Heidelberg, Germany, 2013; pp. 66–78.
22. Conforti, R.; Dumas, M.; García-Bañuelos, L.; La Rosa, M. BPMN Miner: Automated discovery of BPMN process models with hierarchical structure. *Inf. Syst.* **2016**, *56*, 284–303. [[CrossRef](#)]
23. Buijs, J.C.; Van Dongen, B.F.; van Der Aalst, W.M. On the role of fitness, precision, generalization and simplicity in process discovery. In Proceedings of the OTM Confederated International Conferences On the Move to Meaningful Internet Systems, Rome, Italy, 10–14 September 2012; Springer: Berlin/Heidelberg, Germany, 2012; pp. 305–322.
24. Fox, F.; Aggarwal, V.R.; Whelton, H.; Johnson, O. A data quality framework for process mining of electronic health record data. In Proceedings of the 2018 IEEE International Conference on Healthcare Informatics (ICHI), New York, NY, USA, 4–7 June 2018; pp. 12–21.
25. Mans, R.S.; van der Aalst, W.M.; Vanwersch, R.J.; Moleman, A.J. Process mining in healthcare: Data challenges when answering frequently posed questions. In *Process Support and Knowledge Representation in Health Care*; Springer: Berlin/Heidelberg, Germany, 2012; pp. 140–153.
26. Bose, R.J.C.; Mans, R.S.; van der Aalst, W.M. Wanna improve process mining results? In Proceedings of the 2013 IEEE Symposium on Computational Intelligence and Data Mining (CIDM), Singapore, 16–19 April 2013; pp. 127–134.
27. Van Der Aalst, W. Process mining: Overview and opportunities. *ACM Trans. Manag. Inf. Syst. (TMIS)* **2012**, *3*, 1–17. [[CrossRef](#)]
28. Conforti, R.; La Rosa, M.; Ter Hofstede, A.H.; Augusto, A. Automatic repair of same-timestamp errors in business process event logs. In Proceedings of the International Conference on Business Process Management, Seville, Spain, 13–18 September 2020; Springer: Berlin/Heidelberg, Germany, 2020; pp. 327–345.
29. Sim, S.; Bae, H.; Choi, Y. Likelihood-based multiple imputation by event chain methodology for repair of imperfect event logs with missing data. In Proceedings of the 2019 International Conference on Process Mining (ICPM), Aachen, Germany, 24–26 June 2019; pp. 9–16.
30. Alharbi, A.; Bulpitt, A.; Johnson, O. Improving pattern detection in healthcare process mining using an interval-based event selection method. In Proceedings of the International Conference on Business Process Management, Barcelona, Spain, 10–15 September 2017 Springer: Berlin/Heidelberg, Germany, 2017; pp. 88–105.
31. van der Aa, H.; Gal, A.; Leopold, H.; Reijers, H.A.; Sagi, T.; Shraga, R. Instance-based process matching using event-log information. In Proceedings of the International Conference on Advanced Information Systems Engineering, Essen, Germany, 12–16 June 2017; Springer: Berlin/Heidelberg, Germany, 2017; pp. 283–297.
32. Klinkmüller, C.; Weber, I.; Mendling, J.; Leopold, H.; Ludwig, A. Increasing recall of process model matching by improved activity label matching. In *Business Process Management*; Springer: Berlin/Heidelberg, Germany, 2013; pp. 211–218.
33. Dijkman, R.; Dumas, M.; Van Dongen, B.; Käärik, R.; Mendling, J. Similarity of business process models: Metrics and evaluation. *Inf. Syst.* **2011**, *36*, 498–516. [[CrossRef](#)]
34. Richter, F.; Zellner, L.; Azaiz, I.; Winkel, D.; Seidl, T. LIProMa: Label-independent process matching. In Proceedings of the International Conference on Business Process Management, Vienna, Austria, 1–6 September 2019; Springer: Berlin/Heidelberg, Germany, 2019; pp. 186–198.
35. Koschmider, A.; Ullrich, M.; Heine, A.; Oberweis, A. *Revising the Vocabulary of Business Process Element Labels*. *International Conference on Advanced Information Systems Engineering*; Springer: Berlin/Heidelberg, Germany, 2015; pp. 69–83.
36. Mendling, J.; Reijers, H.A.; Recker, J. Activity labeling in process modeling: Empirical insights and recommendations. *Inf. Syst.* **2010**, *35*, 467–482. [[CrossRef](#)]

37. Rubner, Y.; Tomasi, C.; Guibas, L.J. The earth mover's distance as a metric for image retrieval. *Int. J. Comput. Vis.* **2000**, *40*, 99–121. [[CrossRef](#)]
38. Assent, I.; Wenning, A.; Seidl, T. Approximation techniques for indexing the earth mover's distance in multimedia databases. In Proceedings of the 22nd International Conference on Data Engineering (ICDE'06), Atlanta, GA, USA, 3–7 April 2006; p. 11.
39. Zhang, M.; Liu, Y.; Luan, H.; Sun, M.; Izuha, T.; Hao, J. Building earth mover's distance on bilingual word embeddings for machine translation. In Proceedings of the AAAI Conference on Artificial Intelligence, Phoenix, AZ, USA, 12–17 February 2016; Volume 30.
40. Brockhoff, T.; Uysal, M.S.; van der Aalst, W.M. Time-aware Concept Drift Detection Using the Earth Mover's Distance. In Proceedings of the 2020 2nd International Conference on Process Mining (ICPM), Padua, Italy, 5–8 October 2020; pp. 33–40.
41. Guo, Q.; Wen, L.; Wang, J.; Yan, Z.; Philip, S.Y. Mining invisible tasks in non-free-choice constructs. In Proceedings of the International Conference on Business Process Management, Rio de Janeiro, Brazil, 18–22 September 2016; Springer: Berlin/Heidelberg, Germany, 2016; pp. 109–125.
42. Johnson, S.C. Hierarchical clustering schemes. *Psychometrika* **1967**, *32*, 241–254. [[CrossRef](#)]
43. Faloutsos, C.; Ranganathan, M.; Manolopoulos, Y. Fast subsequence matching in time-series databases. *Acm Sigmod Rec.* **1994**, *23*, 419–429. [[CrossRef](#)]
44. Sturges, H.A. The choice of a class interval. *J. Am. Stat. Assoc.* **1926**, *21*, 65–66. [[CrossRef](#)]
45. Levenshtein, V.I. Binary codes capable of correcting deletions, insertions, and reversals. *Sov. Phys. Dokl.* **1966**, *10*, 707–710. Doklady Akademii Nauk SSSR, V163 No4 845–848 1965.
46. Honnibal, M.; Montani, I. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. *Appear* **2017**, *7*, 411–420.
47. Ruta, D.; Gabrys, B. Classifier selection for majority voting. *Inf. Fusion* **2005**, *6*, 63–81. [[CrossRef](#)]
48. Berti, A.; Van Zelst, S.J.; van der Aalst, W. Process mining for python (PM4Py): Bridging the gap between process-and data science. *arXiv* **2019**, arXiv:1905.06169.
49. Switaj, T.L.; Christensen, S.; Brewer, D.M. Acute coronary syndrome: Current treatment. *Am. Fam. Physician* **2017**, *95*, 232–240. [[PubMed](#)]
50. Chew, D.P.; Scott, I.A.; Cullen, L.; French, J.K.; Briffa, T.G.; Tideman, P.A.; Woodruffe, S.; Kerr, A.; Branagan, M.; Aylward, P.E. National Heart Foundation of Australia and Cardiac Society of Australia and New Zealand: Australian clinical guidelines for the management of acute coronary syndromes 2016. *Med. J. Aust.* **2016**, *205*, 128–133. [[CrossRef](#)]