Uncovering the Hidden Significance of Activities Location in Predictive Process Monitoring

Mozhgan Vazifehdoostirani, Mohsen Abbaspour Onari, Isel Grau, Laura Genga, and Remco Dijkman

Eindhoven University of Technology, Eindhoven, the Netherlands
{m.vazifehdoostirani, m.abbaspour.onari, i.d.c.grau.garcia, l.genga, r.m.dijkman}@tue.nl

Abstract. Predictive process monitoring methods predict ongoing case outcomes by analyzing historical process data. Recent studies highlighted the increasing need to enhance the interpretability of these prediction models. This is often achieved by exploiting post-hoc explainable methodologies to assess the importance of different process features on the predicted outcome. However, the significance of the *location* of process activities on prediction models remains unexplored. In several real-life contexts, there might be potential meaningful relations between the location of the activities and process outcome. This information facilitates insights into process management optimization and decision-making. This paper introduces a novel post-hoc explainable artificial intelligence technique inspired by permutation feature importance to assess the impact of activity locations in predictive models. The experimental results on real-life event logs validate the feasibility of the proposed method, showcasing the influence of the location of (group of) activities on outcome predictions.

Keywords: Predictive Process Monitoring · Explainbale AI · Feature Permutation Importance

1 Introduction

Predictive Process Monitoring (PPM) methods aim to predict the future status of ongoing cases by analyzing historical process data. In recent years, a plethora of Machine Learning (ML) approaches have been proposed to support PPM. Extensive studies and benchmarks have indicated the effectiveness of black-box models, such as XGBoost and Random Forest [18], in accurately predicting process outcomes across diverse domains. Recently, the adoption of Deep Learning approaches in PPM has been on the rise [7]. Although these black-box models demonstrate impressive predictive capabilities, they also bring complexity and limited interpretability as trade-offs.

eXplainable Artificial Intelligence (XAI) aims to address the lack of interpretability of black-box models by supporting process analysts in investigating how a given classifier made its decisions. A popular class of XAI methods is *post-hoc* methods, which explain decisions made by black-box models after they are

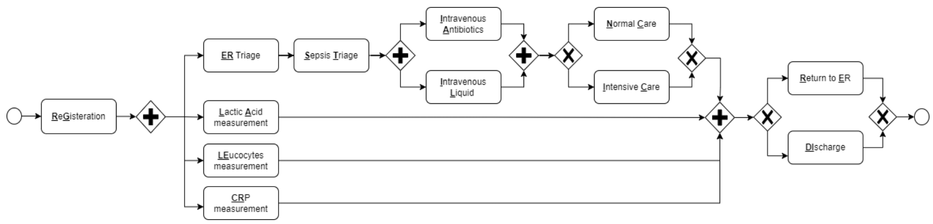

Fig. 1: A simplified version of the sepsis patient trajectory. [14]

Table 1: An example of a training event log from Sepsis.

ID	Trace	Label
1	RG, ***LA***, LE, CR, ER, ST, IL, IA, NC, ER, CR, LE, CR	Negative
2	RG, ER, ST, IL, LE, CR, IA, ***LA***, IC, DI	Positive
3	RG, ***LA***, ER, ST, LE, CR, IA, NC, CR, LE, CR, LE, CR	Negative
4	RG, LE, LE, ***LA***, ER, ST, CR, IA, IC, LE, ***LA***, NC, DI	Positive
5	RG, ER, ST, LE, CR, IL, IA, ***LA***, NC, CR, LE	Positive
6	RG, ***LA***, CR, ER, ST, LE, IL, IA, NC, DI	Negative

built. Earlier research in PPM has utilized techniques like SHAP [4, 23], LIME [16], and Permutation Feature Importance (PFI) [5] to obtain the importance of different process features at both local and global levels. Utilizing these existing XAI methods, prior research has addressed the impact of executing each activity on process outcome [4]; however, the importance of the location of executing activities through the process has not been discussed.

As an example, let us consider the process model in Figure 1, which is a simplified version of the *Sepsis* patient pathways presented in [14]. Table 1 shows an excerpt of an *event log* used to train a model to predict patients returning to the emergency room within 28 days of discharge. Let us consider the activity *Lactic Acid measurement* (***LA***). This activity can be executed at different moments (locations) in the process. From the table, no relation can be detected between the occurrence of this activity and the process outcome. However, the situation is different when its location is considered. The execution of activity ***LA*** in position 2 corresponds to negative process outcomes, whereas its occurrence in other locations correlates with positive outcomes. This intriguing observation might encourage a process analyst to assess the importance of the location of ***LA*** and compare this importance with other activities.

To assess the importance of activity locations with post-hoc methods, one needs to employ a trace encoding technique able to represent activity locations explicitly. To incorporate the location information in training an ML classifier, index-based encoding has been used in several studies [13, 18]. Existing post-hoc XAI approaches usually return two kinds of explanations using index-based encoding. The first one only measures the importance of each *index (location)* in the process. For example, we might understand that the third activity is more important than the second activity executed in the process, regardless of the type of activity. The second one, instead, considers both the kind of the activity and its location, generating explanations in the form "*Occurrence of activity $\mathcal{A}$ at index 3 in case 10 impacts the prediction.*" While providing useful insights, this form of explanation combines the effect of *executing* activity $\mathcal{A}$ and the *location* of its execution in the process. Therefore, the lingering question is whether executing activity $\mathcal{A}$ impacts the process outcome regardless of its occurrence in a particular location. As a result, uncertainties arise regarding the potential impact of relocating activities, like shifting activity $\mathcal{A}$ from timestamp 3 to timestamp 8, on the predictive model. This query can be extended to encompass *multiple* activities that occur collectively (e.g., activities belonging to the same *subprocess*). Prior research has unveiled distinct outcomes when activities are considered in groups as opposed to their individual effects [19]. However, the impact of their location on the outcome remains unexplored.

We argue that being able to provide explanations on which activities may affect the classifier's performance and whether their impact depends on their position provides the process manager with valuable insights about the process. In flexible processes, there is often little or no prior knowledge of how the decision to execute a (group of) activity(ies) at a given moment may affect the outcome. Furthermore, the extracted relations can support process managers in different tasks. For instance, insights on the importance of activity location in determining the process performance can be used in selecting suitable redesign heuristics [6] to explore alternative control-flow constraints during business process re-engineering efforts.

In this paper, we seek to address the following research question: "Given a predictive model trained on a set of process executions, how can we assess the importance of the location of a group of process activities on the classifier performance?" Our main contribution is a novel post-hoc model-agnostic method, inspired by the PFI method, designed to assess the importance of the location of executing activities on outcome prediction. Our experimental results on real-life event logs show the feasibility of the method and provide evidence of the importance of the location of activities for trained outcome prediction models.

The remainder of this paper is organized as follows. In Section 2, we present a review of the relevant related work. Next, in Section 3, we introduce our proposal. The experimental settings and results are discussed in Section 4. Finally, Section 5 concludes the research and outlines potential directions for the future.

2 Related work

XAI approaches proposed within the PPM domain can be broadly categorized into *factual* and *counterfactual* explanations [3]. The former aims to reveal the reasoning behind specific predictions, emphasizing the most influential features. On the other hand, counterfactual explanations provide insights into what changes are necessary for an input sample to achieve a desired prediction [3].

Significant efforts have been made to generate realistic counterfactual explanations for process prediction such as DiCE4EL [9], LORELEY [10], and

CREATED [11]. While counterfactual explanations offer valuable insights into hypothetical scenarios, this paper's primary focus remains on factual explanations to gain a deeper understanding of the functionality of black-box models in the context of PPM. Regarding factual explanations, two groups of methods have been exploited in PPM: *intrinsically interpretable* and *post-hoc* methods. Intrinsically interpretable methods seek to build interpretable models from scratch, such as rule-based classifiers [12], neuro-fuzzy networks [15], and linear regression [5]. However, intrinsically interpretable models often underperform compared to their black-box counterparts [1].

Post-hoc methods can be classified into two groups: *model-specific* and *model-agnostic* [1]. Model-specific methods are designed to work with specific prediction models, such as Gated Graph Neural Network [8], LSTM with Layer-wise Relevance Propagation [20], and LSTM with attention layers [21]. In contrast, model-agnostic methods compute explanations based on the inputs and their associated outputs, allowing the process analyzer to use various prediction models.

Regarding the model-agnostic post-hoc techniques, several studies employed LIME to address the problem of low accuracy by identifying the features that cause wrong predictions [16] and providing an explanation for various feature representations of the event log [17]. Various model-agnostic post-hoc methods, such as SHAP and FPI, are widely used in the literature to provide a local and global explanation [5, 23]. A recent study introduced an ML-based approach for generating multi-level explanations, employing logistic regression, attention-based LSTM, and the eXplainable Dual-learning Deep network [22]. However, it is notable that none of these studies have tackled the global importance of activity locations within a process. Hence, we are filling the mentioned gap by proposing a model-agnostic post-hoc explanation method.

3 Methodology

Figure 2 illustrates our proposed method for measuring the global importance of the location of one or a group of activities on the process outcome prediction. The steps depicted with hatched patterns represent this research's primary focus and contribution. Drawing inspiration from the PFI technique designed for tabular data, our method quantifies the reduction in model performance resulting from the random shuffling of a single feature value [2]. In our analysis, we are interested in analyzing the effect of changing activity locations in process executions. However, the direct application of conventional PFI to a tabular-encoded event log is not suitable for our analysis. Let us assume to encode the event log using index-based encoding. In this encoded log, the location of an activity does not map to a single column; instead, it encompasses all the columns related to the same activity. Therefore, the shuffling implemented by standard PFI techniques, which shuffles one column independently from the others, is inadequate for our analysis due to intricate interdependencies among features. Indeed, relocating an activity necessitates coordinated shuffling of all associated location-related columns to maintain distinct activity locations in generated traces. Further-

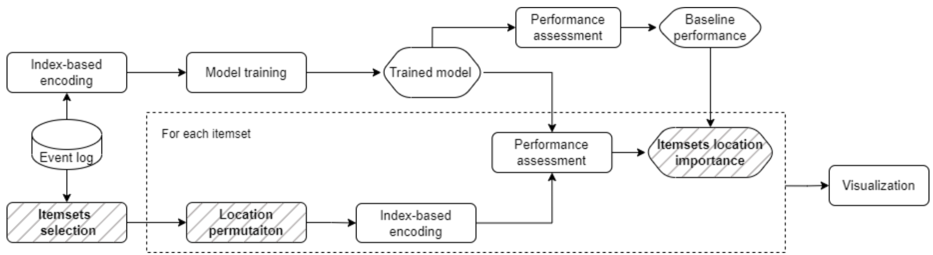

Fig. 2: Overview of the location permutation importance method.

more, shifting the location of an activity inevitably impacts also the location of other activities in the trace. Therefore, we introduced a tailored method that accommodates the sequential nature of processes.

The method takes as an input an *event log*, that is, a multiset of *traces*, each tracking the execution of a given process *instance* (or *case*). Each trace involves a sequence of *events* representing the execution of a given process *activity* at a given timestamp. Within the context of this work, we consider only the activity names as event attributes, assuming that their order in the trace is based on their execution timestamps. Table 1 shows an excerpt of an event log, where each execution corresponds to the treatment administered to a sepsis patient.

First, we *train* a black-box model using index-based encoding and evaluate its performance through k-fold cross-validation on the event log. This provides us with the *baseline performance*. Then, we identify a group of activities (aka. itemsets) whose location importance we want to measure. For each itemset, first, we use the location permutation module to shuffle the location of each item in the itemset across the existing traces. Then, we apply index-based encoding to the shuffled event traces. Next, we assess the model's performance on the permuted event logs using the original trained model through k-fold cross-validation. The difference between the baseline performance and the performance on the permuted event log indicates the importance of the itemset in question. These iterative steps are repeated for each individual itemset. Finally, we plot the importance of each itemset in one plot to facilitate the location importance comparison. In the remainder of this section, we delve into the itemsets selection (3.1) and location permutation (3.2) modules.

3.1 Itemsets selection

Depending on the purpose of the analysis, a process analyst might be interested in exploring the significance of the location of different groups of activities. To facilitate this exploration, we utilize an **I**temsets **E**xtractor function denoted as IE($\mathcal{L}$). Given an event log $\mathcal{L}$, this function generates itemsets of activities considered of interest for the analysis. This function can be customized according to the analyst's notion of interest and may encompass diverse methods. A common and straightforward approach is measuring an itemset's interest through frequency, often accomplished using the Apriori algorithm [24].

As an example, let us consider the event log depicted in Table 1 and let us assume to use an implementation of IE($\mathcal{L}$) which returns itemsets with a minimum frequency of 80%. An excerpt of the output of the itemset selection step would then be IE($\mathcal{L}$) = {{RG}, {ER, CR}, {RG, LE, CR}, ...}.

3.2 Location permutation

We perform the location permutation module for each discovered itemset to shuffle the location of activities in that itemset within all traces. However, randomly shuffling activities' locations has important drawbacks. First, it is likely to result in unrealistic traces. Second, when permuting itemsets involving more than one activity, a completely random shuffle may result in a permuted trace where the relative order of activities in the itemset has also been changed, thus introducing noise in assessing the importance of the itemset location. To mitigate these issues, we introduce the following two constraints:

- **Feasibility Constraint**: We restrict the location permutation of an activity to the observed locations of the occurrences of that activity throughout the event log.
- **Preserving Ordering Relation**: When shuffling itemset activities, we preserve their relative order within each trace.

These constraints allow us to balance between maintaining meaningful process behaviors in generated traces and introducing enough randomness and variability in itemset locations within the event log.

Keeping these constraints in mind, given an itemset $\mathcal{I}$ extracted by an event log $\mathcal{L}$, we permute the location of $\mathcal{I}$ in each trace in which it occurs by implementing the following steps. We first extract a set of **O**bserved **L**ocations for $\mathcal{I}$ throughout the event log. To uphold the sequential order of activities within an itemset, we define tuples encompassing the itemset activities. For instance, the **O**bserved **L**ocations for the itemset {ER, CR} in the event log $\mathcal{L}$ shown in Table 1 can be defined as OL($\mathcal{L}$, (ER, CR)) = {(5,4), (10, 11), (2,6), (3,6), (5,7), (2,5), (4,3)}. In certain locations, activity CR follows activity ER, whereas, in other instances, CR precedes ER. It's important to note that if a single activity from the itemset appears multiple times in a trace and results in an overlapped occurrence of the itemset, only the first complete occurrence is considered for collecting observed locations and subsequent permutation. For instance, in case number 1 in Table 1 we observe the occurrence of {CR, ER} twice, along with an additional CR. Since the final occurrence of CR lacks a distinct accompanying ER to form the itemset, its location is not included in the function's output.

Next, we randomly rearrange the positions of itemsets in traces using the pool of observed locations. We extract a random location from the set of observed locations for every activity of the itemset within a trace. If the chosen location conforms to the existing sequential order of activities of the itemset instance, it is retained. Otherwise, an alternative location is drawn to ensure that the activity relations remain consistent within that specific itemset instance. This operation is repeated for each occurrence of the itemset in the trace.

As an example, let us consider case number 1 from the running example: $\sigma_1 = \langle RG, LA, LE, \mathbf{CR}, \mathbf{ER}, ST, IL, IA, NC, \boldsymbol{ER}, \boldsymbol{CR}, LE, CR\rangle$. A possible permutation for itemset {ER, CR} generates the permuted trace $\sigma_1^{'} = \langle RG, LA, \mathbf{CR}, \mathbf{ER}, \boldsymbol{ER}, LE, \boldsymbol{CR}, ST, IL, IA, NC, LE, CR\rangle$.

4 Implementation and experiments

4.1 Settings

We utilized XGBoost, a top-performing ensemble model in outcome prediction according to [18], and evaluated it using a 5-fold cross-validation. To account for the variability introduced by the permutation process, we repeated the location permutation 10 times and utilized box plots to visualize the results. A decrease in the f1-score indicates the importance of each itemset under consideration.

We considered two analysis goals, i.e., deriving the location importance for the single and frequent group of activities. For the latter, we extracted the top 10 frequent itemsets of size greater than one using the Apriori algorithm [24].

In addition to assessing the importance of itemset locations, we have also employed the conventional PFI technique to assess the importance associated with the existence of itemsets. To achieve this, we utilized binary encoding, where each feature represented a distinct itemset. We trained an XGBoost model on the binary-encoded event log. It is worth mentioning that the difference between the importance of the location and the existence of the selected itemsets is also due to using different encoding methods. Our aim is not to establish an absolute comparison between the obtained values but, instead, to shed light on the insights gained from examining the location of activities as opposed to their existence. The implementation of the proposed method is available at GitHub[1].

We used public event logs widely used in the literature, such as *bpic2011*, *bpic2012*, and *Sepsis* event logs with the same labeling strategy as in [18].

4.2 Results and discussion

Single-activity itemsets. Figure 3 represents the top 20 most important single activities regarding their location in the process (in green) and their corresponding existence importance values (in blue). Activities that hold the greatest positional significance do not necessarily maintain equivalent existence importance. For instance, in the *bpic2012 accepted* event log, the location of the activity "W Nabellen Offer" stands out as an influential factor, leading to a decrease in performance between 12% and 14%. However, its existence does not yield substantial explanatory power for the predictive model. To delve more into this activity, since we lack expert access to validate explanations from the black box model, we have explored the event log to assess whether the data aligns with our findings. In particular, we have plotted the distribution of process outcomes associated with different locations of this activity in Figure 4a. According to the

[1] https://github.com/MozhganVD/PermutationLocationImportance

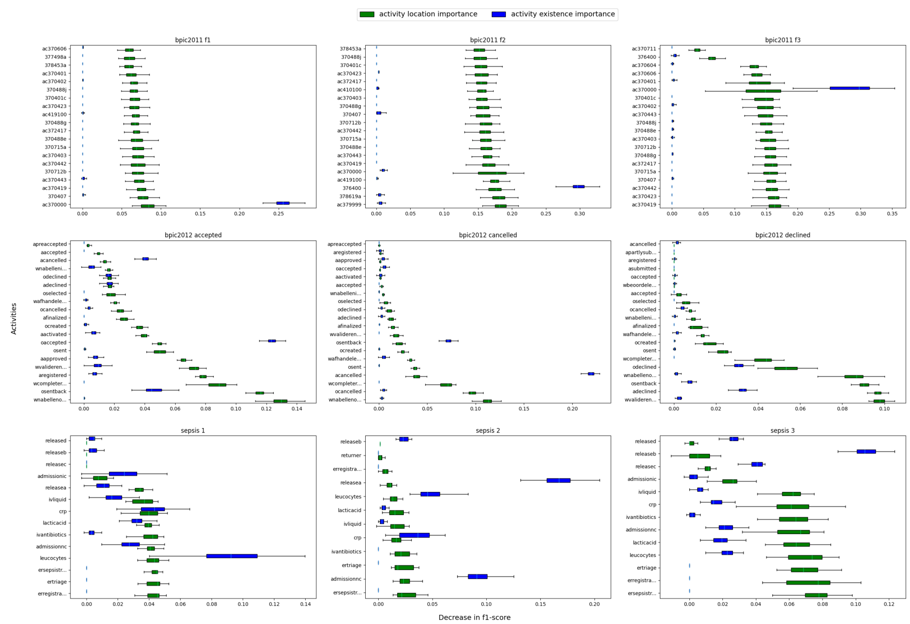

Fig. 3: Single activities location importance vs. their existence importance.

event log, the occurrence of this activity leads to the process being accepted in 61% of the cases; however, we observe variations in the acceptance percentage from 50% to 70% in different locations. While these observations cannot serve as direct validation for the presented results, they suggest that our findings in terms of location and existence importance are meaningful in the log under analysis.

Likewise, across all three labeled event logs extracted from the *bpic2011* dataset, most activities demonstrate high significance based on their location, while their existence is insignificant for the model. Conversely, we encounter activities showing a relatively strong influence due to their sheer presence. Take the activity "ac370000" within the *bpic2011 f1* event log; for instance, it demonstrates notably high importance in relation to its existence. A similar trend is observable across all event logs extracted from the *sepsis* dataset, where multiple activities such as "leucocytes" in *sepsis 1* (see Figure 4b), "release A" in *sepsis 2*, and "release B" in *sepsis 3* showcase the same behavioral patterns. As depicted in Figure 4b, the presence of the "leukocytes" activity strongly correlates with positive outcomes, only with minor variations in location.

Multiple-activities itemsets. Figure 5 represents the location and existence importance of the extracted itemsets. A first observation is that, in general, the existence of the frequent itemsets does not exhibit much importance for the prediction model. This observation could potentially stem from our selection of the most frequent itemsets, which might not necessarily be the most predictive ones.

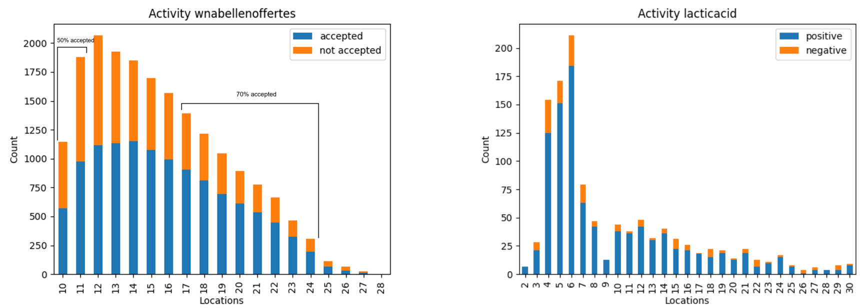

(a) Example from *bpic2012 accepted*. (b) Example from *sepsis 1*.

Fig. 4: Distribution of the locations of activities over different outcomes.

However, permuting their location could still affect the prediction model. Notably, within the *bpic2011 f1* log, except for the itemset {ac370000, ac370419} whose existence demonstrates a noticeable effect on the prediction model, the remainder of the frequent itemsets are primarily influential due to their location.

We have plotted the distribution of outcomes across diverse locations for itemset with the highest *existence* importance in *bpic2011 f1* (Figure 6b) and with the highest *location* importance in *bpic2011 f2* (Figure 6a). Note that when plotting itemset locations, we can easily have a lot of locations with low-frequency values. This would lead to many locations in the X-axis with many small bars, resulting in a very cluttered figure. To simplify the visualization, we aggregated all locations with the same starting point. For example, all locations of (1,3), (1,5), and (1,10) are shown in index 1 in the plot. While this representation is less accurate since only the starting position of each itemset is considered, it still provides a reasonable estimate of how the output distribution varies in relation to different locations. Notably, the itemset {ac379999, ac370000} exhibits significant variation in positive and negative class proportions across varying locations. There are also more possible locations for this itemset in general. In contrast, the itemset {ac370000, ac370419} maintains an almost consistent positive and negative fraction across different locations.

Discussion The obtained results confirm the feasibility and practicality of the proposed methods in measuring the importance of the location of activities. They also underscore the need to delve further into the significance of locations within process analysis, given the meaningful relationships between activity location and associated outcomes detected in real-life event logs.

Nonetheless, our approach presents some limitations. First, our method's focus is on classifier performance, which might not fully reveal the underlying relationships between activities and outcomes in the dataset. This prompts the need for expert validation of the insights gathered. Moreover, an activity location's importance is inherently intertwined with other activity locations. While

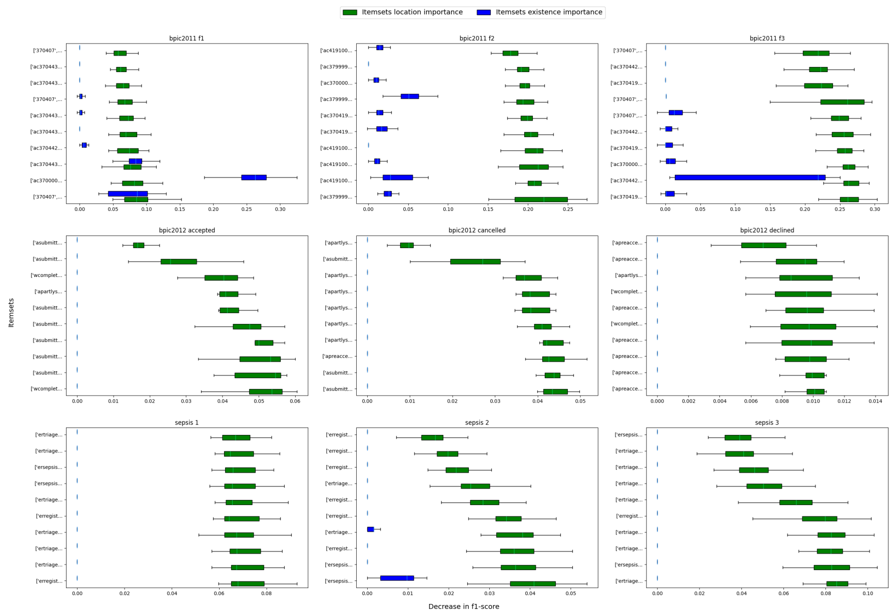

Fig. 5: Location vs. existence importance for 10 most frequent itemsets.

we have addressed this issue by implementing repeated permutations, quantifying the exact impact of these permutations remains a complex challenge. Also, despite imposing feasibility constraints, the risk of placing activities in prohibited locations due to dependencies persists, and more sophisticated methods are needed. Additionally, the effectiveness of our method is influenced by the design choices made during various methodological steps. This acknowledges the potential for refining strategies, such as employing predictive pattern detection methods, as previous studies discussed that frequent patterns are not necessarily the most predictive ones [19].

5 Conclusion and future work

Our exploration into XAI within PPM has shed light on the importance of understanding activities' location-based significance. Our novel post-hoc method, inspired by the Permutation Feature Importance, endeavors to address this gap by quantifying the impact of activities locations on predictive process outcomes. We performed several experiments on real-life event logs widely used in the PPM domain. Our experiments underscore the method's feasibility, as demonstrated through its successful application to real-life event logs. The results provide clear evidence regarding the latent importance of activities' location. This insight into the significance of location adds a new dimension to our understanding, one that is frequently overlooked by existing XAI techniques.

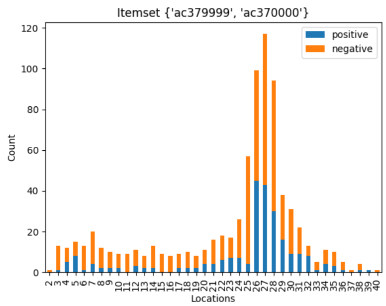

(a) Example from *bpic2011 f2*.

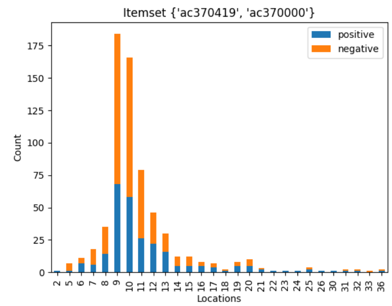

(b) Example from *bpic2011 f1*.

Fig. 6: Distribution of the locations of itemsets over different outcomes.

These considerations highlight the potential for further extension of this research, prompting us to investigate different approaches for itemset selection, seeking more meaningful itemsets rather than just being frequent. For instance, we intend to construct outcome-oriented patterns introduced in [19] rather than relying on frequent itemsets. Additionally, we are working on a methodology to evaluate the importance of the order of activities within each itemset, besides their location importance. Furthermore, our future research endeavors to develop more sophisticated approaches to ensure the feasibility of permuted traces.

References

1. Arrieta, A.B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., et al.: Explainable artificial intelligence: Concepts, taxonomies, opportunities and challenges toward responsible ai. Information fusion **58**, 82–115 (2020)
2. Breiman, L.: Random forests. Machine learning **45**, 5–32 (2001)
3. Buliga, A., Di Francescomarino, C., Ghidini, C., Maggi, F.M.: Counterfactuals and ways to build them: Evaluating approaches in predictive process monitoring. In: International Conference on Advanced Information Systems Engineering. pp. 558–574. Springer (2023)
4. El-khawaga, G., Abu-Elkheir, M., Reichert, M.: Xai in the context of predictive process monitoring: An empirical analysis framework. Algorithms **15**(6), 199 (2022)
5. Elkhawaga, G., Abu-Elkheir, M., Reichert, M.: Explainability of predictive process monitoring results: Can you see my data issues? Applied Sciences **12**(16) (2022)
6. Fehrer, T., Fischer, D.A., Leemans, S.J., Röglinger, M., Wynn, M.T.: An assisted approach to business process redesign. Decision Support Systems **156** (2022)
7. Harane, N., Rathi, S.: Comprehensive survey on deep learning approaches in predictive business process monitoring. Modern Approaches in Machine Learning and Cognitive Science pp. 115–128 (2020)
8. Harl, M., Weinzierl, S., Stierle, M., Matzner, M.: Explainable predictive business process monitoring using gated graph neural networks. Journal of Decision Systems **29**, 312–327 (2020)

9. Hsieh, C., Moreira, C., Ouyang, C.: Dice4el: interpreting process predictions using a milestone-aware counterfactual approach. In: International Conference on Process Mining. pp. 88–95. IEEE (2021)
10. Huang, T.H., Metzger, A., Pohl, K.: Counterfactual explanations for predictive business process monitoring. In: European, Mediterranean, and Middle Eastern Conference on Information Systems. pp. 399–413. Springer (2021)
11. Hundogan, O., Lu, X., Du, Y., Reijers, H.A.: Created: Generating viable counterfactual sequences for predictive process analytics. In: International Conference on Advanced Information Systems Engineering. pp. 541–557. Springer (2023)
12. Lee, S., Comuzzi, M., Kwon, N.: Exploring the suitability of rule-based classification to provide interpretability in outcome-based process predictive monitoring. Algorithms **15**(6), 187 (2022)
13. Leontjeva, A., Conforti, R., Di Francescomarino, C., Dumas, M., Maggi, F.M.: Complex symbolic sequence encodings for predictive monitoring of business processes. In: Business Process Management: 13th International Conference, BPM. pp. 297–313. Springer (2015)
14. Munoz-Gama, J., Martin, N., Fernandez-Llatas, C., Johnson, O.A., Sepúlveda, M., Helm, E., Galvez-Yanjari, V., Rojas, E., Martinez-Millana, A., Aloini, D., et al.: Process mining for healthcare: Characteristics and challenges. Journal of Biomedical Informatics **127**, 103994 (2022)
15. Pasquadibisceglie, V., Castellano, G., Appice, A., Malerba, D.: Fox: a neuro-fuzzy model for process outcome prediction and explanation. In: International Conference on Process Mining. pp. 112–119. IEEE (2021)
16. Rizzi, W., Di Francescomarino, C., Maggi, F.M.: Explainability in predictive process monitoring: when understanding helps improving. In: International Conference on Business Process Management. pp. 141–158. Springer (2020)
17. Sindhgatta, R., Moreira, C., Ouyang, C., Barros, A.: Exploring interpretable predictive models for business processes. In: Business Process Management: 18th International Conference. pp. 257–272. Springer (2020)
18. Teinemaa, I., Dumas, M., Rosa, M.L., Maggi, F.M.: Outcome-oriented predictive process monitoring: Review and benchmark. ACM Transactions on Knowledge Discovery from Data **13**(2) (2019)
19. Vazifehdoostirani, M., Genga, L., Lu, X., Verhoeven, R., van Laarhoven, H., Dijkman, R.: Interactive multi-interest process pattern discovery. In: International Conference on Business Process Management. pp. 303–319. Springer (2023)
20. Weinzierl, S., Zilker, S., Brunk, J., Revoredo, K., Matzner, M., Becker, J.: Xnap: making lstm-based next activity predictions explainable by using lrp. In: International Conference on Business Process Management. pp. 129–141. Springer (2020)
21. Wickramanayake, B., He, Z., Ouyang, C., Moreira, C., Xu, Y., Sindhgatta, R.: Building interpretable models for business process prediction using shared and specialised attention mechanisms. Knowledge-Based Systems **248**, 108773 (2022)
22. Wickramanayake, B., Ouyang, C., Xu, Y., Moreira, C.: Generating multi-level explanations for process outcome predictions. Engineering Applications of Artificial Intelligence **125**, 106678 (2023)
23. Yang, M., Moon, J., Yang, S., Oh, H., Lee, S., Kim, Y., Jeong, J.: Design and implementation of an explainable bidirectional lstm model based on transition system approach for cooperative ai-workers. Applied Sciences **12**(13), 6390 (2022)
24. Yuan, X.: An improved apriori algorithm for mining association rules. In: AIP conference proceedings. vol. 1820. AIP Publishing (2017)