How to make process model matching work better? An analysis of current similarity measures

Fakhra Jabeen[1], Henrik Leopold[1], and Hajo A. Reijers[1,2]

[1] VU University Amsterdam,
De Boelelaan 1081, 1081 HV Amsterdam, The Netherlands
f.jabeen|h.leopold|h.a.reijers@vu.nl
[2] Eindhoven University of Technology
PO Box 513, 5600MB Eindhoven, The Netherlands

Abstract. Process model matching techniques aim at automatically identifying activity correspondences between two process models that represent the same or similar behavior. By doing so, they provide essential input for many advanced process model analysis techniques such as process model search. Despite their importance, the performance of process model matching techniques is not yet convincing and several attempts to improve the performance have not been successful. This raises the question of whether it is really not possible to further improve the performance of process model matching techniques. In this paper, we aim to answer this question by conducting two consecutive analyses. First, we review existing process model matching techniques and give an overview of the specific technologies they use to identify similar activities. Second, we analyze the correspondences of the Process Model Matching Contest 2015 and reflect on the suitability of the identified technologies to identify the missing correspondences. As a result of these analyses, we present a list of three specific recommendations to improve the performance of process model matching techniques in the future.

Key words: Process Model Matching, Performance Improvement, Weakness Analysis, Activity Similarity

1 Introduction

Process model matching refers to the automatic identification of corresponding activities between two process models, i.e. activities that represent the same or similar behavior. By automatically producing such activity correspondences, process model matching techniques are a prerequisite for many advanced analysis techniques. Among others, the identification of activity correspondences is required for the harmonization of process model variants [1, 2], process model search [3, 4], and the detection of process model clones [5, 6]. Recognizing the importance of matching for the automated analysis of process models in general, researchers have defined a plethora of process model matching techniques (see e.g. [7–10]).

Despite the considerable attention that has been devoted to the problem of process model matching, the performance of existing matching techniques is not yet convincing. The results from the Process Model Matching Contests in 2013 and 2015 show that, depending on the data set, the best F-measures range between 0.45 and 0.67 [11, 12]. While the need for performance improvements is widely recognized, attempts to further improve the results of process model matching, for example through performance prediction, were not very fruitful [13]. This raises the question of how the performance of process model matching techniques can be further improved.

In this paper, we aim to answer this question by a) systemically analyzing the technological state of the art and b) by analyzing the missing capabilities of existing process model matching techniques. To this end, we first conduct a structured literature review on process model matching. We provide an overview of existing techniques and the specific technologies they use to identify similar activities. Then, we analyze the characteristics of the correspondences of the Process Model Matching Contest 2015 that the participating matching techniques failed to identify. In this way, we aim to develop an understanding to what extent current matching performance can be explained by a focus on a limited set of technologies and which directions might be promising to improve process model matching performance in the future.

The remainder of the paper is organized as follows. Section 2 introduces the problem of process model matching using a running example. Section 3 discusses the methodological details and the results of our literature review. Section 4 presents the analysis of the correspondences of the Process Model Matching Contest 2015. Section 5 elaborates on opportunities for improving the performance of process model matching and gives three specific recommendations. Section 6 concludes the paper.

2 The Problem of Process Model Matching

Process model matching techniques aim at automatically identifying activity correspondences that represent similar behavior in both models. Figure 1 illustrates the matching problem by showing the recruitment processes from two different companies. The grey shades highlight the correspondences between the two processes. For example, the activity "*Evaluate*" from company B corresponds to the activities "*Check grades*" and "*Examine employment references*" from company A. The correspondences show that the two models differ with respect to the terms they use (e.g. "*eligibility assessment*" versus "*aptitude test*") as well as their level of detail (e.g. "*Evaluate*" is described in more detail in the model from company A).

Given such differences, the proper recognition of the correspondences between two process models can become a complex and challenging task. The complexity of the matching task is also highlighted by the rather moderate performance of process model matching techniques. A recent comparative evaluation in the context of the Process Model Matching Contest (PMMC) 2015 showed that the F-measures lie between 0.45 and 0.67 for different data sets [12].

Following Gal [14], we can subdivide the matching process into first line matching and second line matching. A first line matcher takes the sets of activities A_1 and A_2 from the process models as input and produces a similarity matrix $M(A_1, A_2)$ with $|A_1|$ rows and $|A_2|$ columns. Among others, such a similarity matrix can be obtained by comparing the activity labels. A second line matcher takes one or more similarity matrices produced by first line matchers as input and turns them into a binary similarity matrix

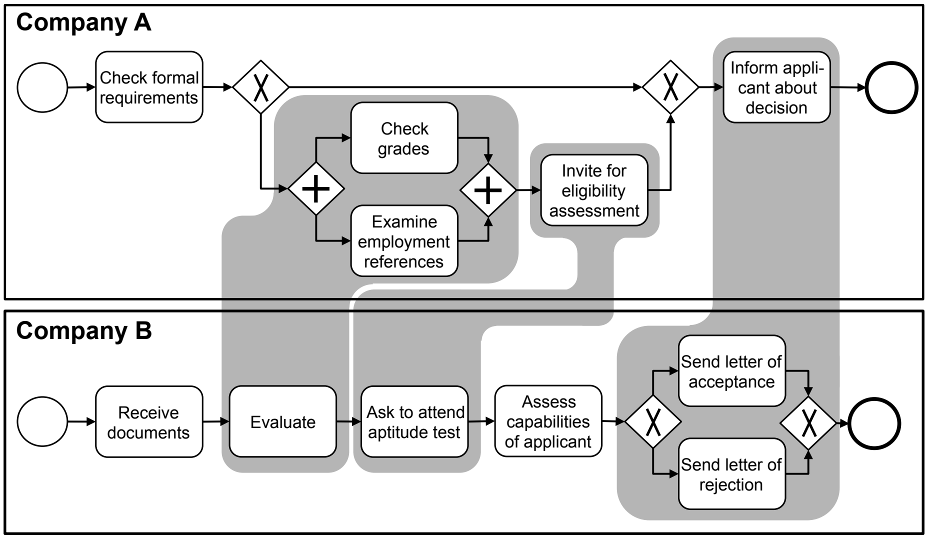

Fig. 1: Two business processes and their correspondences

$M(A_1, A_2)$ with entries of 0 or 1. The latter indicates a correspondence between two activities.

It is important to note that first line matching plays a particularly important role for the overall matching result. If a first line matcher computes a similarity value of zero for two activities, it is very unlikely that a second line matcher will include this particular activity pair in the final set of correspondences. In the next section, we therefore conduct a systematic literature review and analyze which technologies are employed for first line matching.

3 Review of Existing First Line Matching Measures

To gain insights into the state of the art of process model matching, we conducted a systematic literature review on the measures used in the context of first line matching. In Section 3.1, we describe our search strategy. In Section 3.2, we elaborate on the results of our review.

3.1 Search Strategy

We conducted a comprehensive *literature review* on process model matching. More specifically, we queried the ACM Digital Library, IEEEXplore Digital Library, Springer Link, and Science Direct for relevant conference and journal papers. As search terms we combined "*business process*", "*workflow*", and "*process model*" with "*matching*", "*similarity*", "*alignment*", and "*query*". Based on these search terms, we retrieved 5,862 papers in total, of which we selected 657 for further screening. We considered only those papers as relevant that proposed a measure taking two process models as input and producing a set of activity correspondences at some stage. In this way, after removing duplicates, we obtained a total of 30 papers, each describing and using at least one measure. We studied all measures in detail and analyzed their usage for identifying similar activities. Since some papers (i.e. [11] and [12]) described more than a single matching system, the total result of our literature study is a set of 35 process model matching systems, each employing one or more measures for first line matching.

3.2 Search Results

The result of our search is summarized in Table 1. Each row in the table lists a measure type that is used to identify activity correspondences. The *Total* column indicates the total number of matching systems using the respective measure type and the *Reference* column shows the papers discussing these matching systems. Overall, Table 1 shows that we identified a total of 10 measure types, which we categorized into *syntactic* and *semantic* measures.

Table 1: Measures used for first line matching

Measure Type	Total	References
Syntactic		
Distance-based	21	5 in [11], 2 in [12], [15–28]
Jaccard/Dice	5	[3, 12, 21, 24, 29]
Cosine similarity	5	3 in [12], [7, 26]
Substring	4	3 in [12], [24]
Jensen-Shannon distance	1	[12, 30]
Semantic		
Synonym-based	16	3 in [12], [3, 11, 15, 16, 19, 21, 24–26, 29, 31, 32]
Lin	12	3 in [12], [8, 10, 11, 21, 22, 22, 27, 33, 34]
Hypernym-based	5	[11], 2 in [10, 12]
Wu & Palmer	3	2 in [11], [12]
Lesk	2	[27, 34]

Syntactic Measures. Syntactic measures relate to simple string comparisons and do not take the meaning or context of words into account. The most prominently employed syntactic measures are *distance-based* measures such as the Levenshtein distance. Given two labels l_1 and l_2, the Levenshtein distance counts the number of edit operations (i.e. insertions, deletions, and substitutions) that are required to transform l_1 into l_2. Another distance-based measure is the Jaro-Winkler distance, which works in a similar way, but produces a value between 0 and 1.

Besides distance-based measures, many matching systems rely on plain word comparisons. Very common measures include the *Jaccard* and the *Dice* coefficient, which both compute the similarity between two activity labels based on the number of shared and non-shared words. An alternative approach based on word comparisons is the *cosine similarity*. To compute the cosine similarity, activity labels are transformed into vectors, typically by weighing words with their frequency of occurrence. The cosine similarity is then given by the cosine of the angle between two activity vectors. An alternative way of taking the word distribution into account is the *Jensen-Shannon* distance, which is a method for measuring the similarity of two probability distributions. However, so far, it has only been employed by the approach from Weidlich et al. [30].

A common pre-processing step is the consideration of *substring* relationships between activities. For instance, Dadashina et al. consider two activities labels l_1 and l_2 to be similar if l_1 is a substring of l_2 (or vice versa) [12]. Such labels are then removed from further similarity considerations and simply receive a similarity score of 1.

Semantic Measures. Semantic measures aim at taking the meaning of words into account. A very common strategy to do so is the identifciation of *synonyms* using the lexical database WordNet [35]. Typically, matching systems check for synonyms as part of a prepossessing step and then apply other, often also syntactic, similarity measures [12]. The most prominent semantic measure is the *Lin similarity*. The Lin similarity is a method to compute the semantic relatedness of words based on their information content according to the WordNet taxonomy. To use the Lin similarity for measuring the similarity between two activities (which mostly contain more than a single word), approaches typically combine the Lin similarity with the bag-of-words model. The bag-of-words model transforms an activity into a multiset of words, ignoring grammar and word order. The Lin similarity can then be obtained by identifying the word pairs from the two bags with the highest Lin score and by computing their average. Other measures based on the WordNet dictionary include *Wu & Palmer* and *Lesk*. The former computes the similarity between two words by considering the path length between these words in the WordNet taxonomy. The latter compares the WordNet dictionary definitions of the two words. Some approaches also directly check for *hypernym* relationships (a hypernym is a more common word). For instance, Hake et al. [12] consider "*car*" and "*vehicle*" as identical words since "*vehicle*" is a hypernym of "*car*".

Discussion. The findings of our literature highlight two important points. First, our review shows that *syntactic measures play a predominant role*. This means they strongly rely on the use of comparable vocabulary among the considered process models. Even common synonyms, such as "*assess*" and "*evaluate*" or "*conduct*" and "*perform*", cannot be detected by approaches relying on syntactic measures. Interestingly, 21 out of 35 systems even rely on the most basic syntactic measure: distance-based similarity. The disadvantage of edit-based distance measures is not only their inability to recognize synonymous terms, but also their tendency to consider unrelated words as similar. As an example, consider the unrelated words "*contract*" and "*contact*". The Levenshtein distance between these words is only 1, indicating a high similarity between the terms. Second, our review shows that the *employed semantic measures are very basic* and exclusively based on the WordNet dictionary. This represents a considerable problem since any WordNet-based measure returns a similarity score of zero if a term is not part of the WordNet dictionary. While the WordNet dictionary is quite extensive, it does not cover complex compound words (e.g. "*problem report*" or "*budget plan*"), which we often find in process models from industry.

Overall, our analysis suggests that current first line matching measures might be not good enough for recognizing the complex notion of similarity between some activities. In the next section, we will empirically investigate whether the choice of syntactic and basic semantic measures can indeed explain the low matching performance.

4 Analysis of the Results of the PMMC 2015

Our literature review in the last section revealed that matching techniques predominantly rely on syntactical and Wordnet-based semantic measures for first line matching. In this section, we investigate to what extend this insight allows us to explain the moderate performance of the matching systems in the Process Model Matching Contest 2015 [12].

To this end, we computed the similarity scores for all 1037 correspondences from the PMMC 2015 gold standard using the most prominently used syntactic and semantic measures, i.e. the Levenshtein distance and the bag-of-words-based Lin similarity. Figure 2 summarizes the results of our computation using box plots. It clusters the results based on the three datasets from the contest (referred to as "*Admission*", "*Birth*", and "*Asset*") and the number of matching systems that identified these correspondences. This means that the first column (0) from Figure 2 a) shows the distribution of the Levenshtein similarity of the correspondences that have not been identified by any matching system (separately for each dataset from the contest). The second column (1) respectively shows the distribution of the Levenshtein similarity of the correspondences that have been identified by exactly one system etc. Since a total of 12 matching systems participated in the PMMC 2015, each graph has 13 columns.

Analyzing the Levenshtein similarity distributions from Figure 2 a) shows that there is a clear relationship between the similarity score of the correspondences and the number of matching systems that successfully identified them. The higher the similarity score, the higher the number of matching systems identifying the correspondence. While this may not be completely surprising, it is striking that especially the correspondences that none of the matching systems identified have a median similarity score of below 0.1 among all datasets. This does not only emphasize how strongly existing matching systems rely on syntactic measures, but also that syntactic measures cannot represent a suitable option for identifying them. It is also interesting to note that the correspondences that were identified by 10 matching systems or more are mostly trivial correspondences, i.e. identical strings.

The bag-of-words-based Lin similarity distributions from Figure 2 b) show a less clear picture. While there is an overall tendency that a higher Lin similarity is associated with a higher number of matching systems identifying a correspondence, we also observe a significant number of deviations. For instance, the median Lin similarity of the correspondences from the Birth dataset that have not been identified by any matching system is already quite high (0.4). For the correspondences from the Admission dataset, we even observe an up and down movement from column 1 to 4. This is, among others, caused by words that are not part of WordNet and, thus, yield in a Lin similarity of 0. A notable analogy to the Levenshtein similarity is the median of 1.0 for correspondences that have been successfully identified by 10 matching systems or more. This can

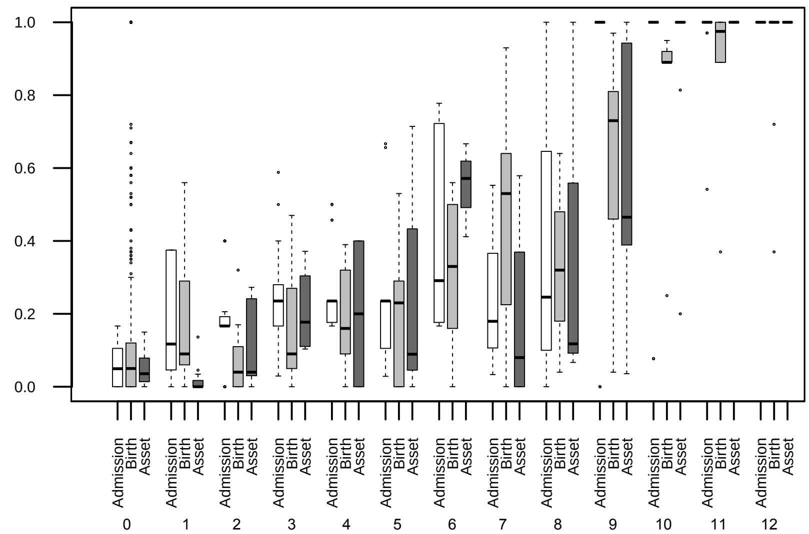

(a) Levenshtein distance

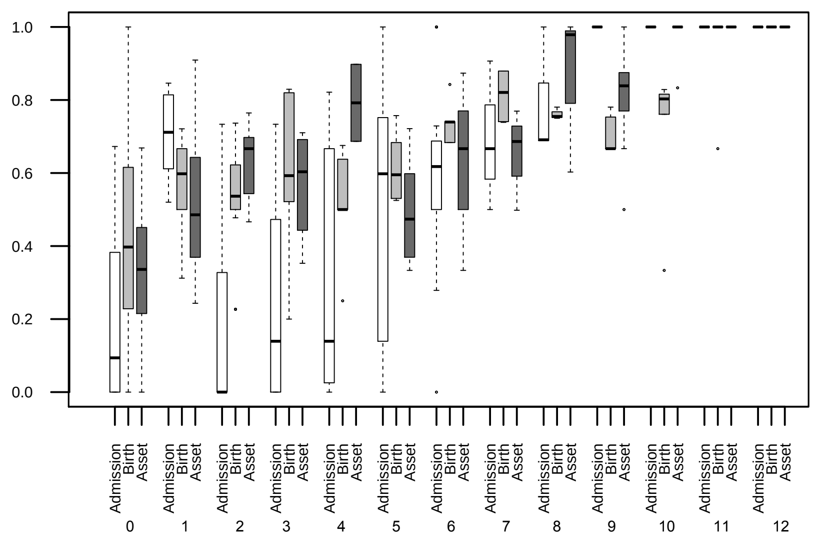

(b) Bag-of-word-based Lin similarity

Fig. 2: Relationship between the number of matching systems identifying the correspondences from the PMMC 2015 gold standard and the similarity score distribution of these correspondences

be easily explained by the fact that the Lin similarity equals the Levenshtein similarity for trivial correspondences. A general observation is that the Lin similarity tends to give higher similarity scores to correspondences than the Levenshtein distance. This is quite an expected outcome since the Lin similarity also semantically relates words. The disadvantage, however, is that the Lin similarity is computed on a word by word level. Consider, for example, the two activities "*Evaluate plan*" and "*Assess contract*". A bag-of-words-based Lin similarity would first identify the best word pairs (i.e. "*evaluate*" and "*assess*" as well as "*plan*" and "*contract*") and average the Lin similarity scores of these pairs. Since "*evaluate*" and "*assess*" are synonyms, their Lin similarity is 1. The Lin similarity between "*plan*" and "*contract*" is 0.42, resulting in an average of 0.71. The resulting Lin similarity between these activities is thus quite high, although they are actually not likely to be related. This example together with the numbers from Figure 2 b) highlights that the mere application of semantic technologies is not sufficient.

To better understand to what extend semantic technology currently contributes to the performance of the matching systems, we further analyzed the correctly identified correspondences by each matching system from the PMMC 2015. Figure 3 illustrates the ration between trivial (i.e. identical strings) and non-trivial correspondences.

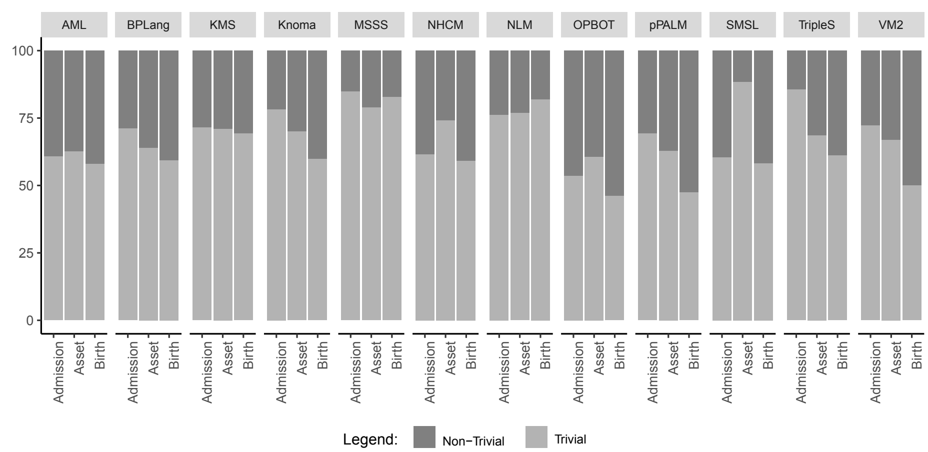

Fig. 3: Ratio between trivial and nontrivial correspondences that the matching systems of the PMMC 2015 correctly identified

The data from Figure 3 shows that the majority of the correspondences identified by the matching systems from the PMMC 2015 are actual trivial correspondences. Most systems identify between 20% and 50% trivial correspondences. It is interesting to note that the systems with the highest share of non-trivial correspondences (e.g. OP-BOT, AML, pPALM, NHCM) also have a particular good performance in the respective datasets. By contrast, the matching systems mainly identifying trivial correspondences (e.g. SMSL, Knoma, NLM, TripleS), also have a relatively bad performance. Looking into the specific techniques the more successful systems use, it is apparent that all of them at least partially build on semantic technology.

5 Recommendations for Improving Process Model Matching

Our literature review together with our empirical analysis of the correspondences from the PMMC 2015 revealed that the moderate performance of current matching systems can be well explained by their choices of first line matching measures. The employed measures are based on either syntactic or very basic semantic technology. As a result, a considerable number of complex correspondences cannot be successfully identified by existing matching systems. Based on our analysis, we derive the following recommendations for future work:

1. *Use syntactic technology for preprocessing only*: Syntactic technology is highly useful for recognizing trivial or almost trivial correspondences. We found that the best performing systems mainly use syntactic technology as a preprocessing step: They first match identical and almost identical labels and then apply semantic technology. The large-scale and sole application of distance-based measures, however, did not improve the results. They rather resulted in a high number of false positives, even with high cut-off values for, e.g. the Levenshtein distance.
2. *Apply of semantic technology beyond the word level*: Our analyses illustrated the importance of semantic technology for identifying non-trivial correspondences. However, it is also highlighted that semantic technology on the word level is not sufficient. Comparing activity labels by computing the semantic similarity between individual word pairs does not account for the cohesion and the complex relationships that exist between the words of activity labels. A first step would be the proper consideration of compound nouns such as "*customer complaint*" [36]. A more advanced step would be to also consider the relationship between nouns and verbs [37]. Only because two activities use the same verb (e.g. "*evaluate*") they are not necessarily related. Possible directions to account for these complex relationships are technologies such as distributional semantics [38]. They have been found to considerably improve matching results in other matching contexts [39].
3. *Use of domain-specific dictionaries*: Especially for the Birth dataset from the PMMC 2015, the recall values are particularly low. Analyzing the correspondences that the matching systems failed to identify, reveals that these correspondences often use domain-specific words or describe domain-specific procedures. Taking into account how current semantic technology is created and trained, it is not likely that there exists an off-the-shelf solution that is conducive for the identification of these correspondences. Hence, we recommend building on domain-specific dictionaries. They can be used for both inferring relationships between domain-specific words as well as for training statistical approaches, such as the previously mentioned distributional semantics methods. Existing methods for automatically extracting ontologies may represent a promising starting point here [40].

We believe that these three recommendations can appropriately address the weaknesses we identified in our analyses and hope that they provide valuable directions to further improve process model matching.

6 Conclusion

In this paper, we addressed the question of how to improve the performance of process model matching techniques. To this end, we conducted a literature review on existing process model matching systems and the specific technologies they use for identifying activity correspondences. Then, we analyzed the results from the Process Model Matching Contest 2015 in order to learn to what extend the employed technology for the identification of activity correspondences represents a reasonable choice.

Our literature review showed that all existing matching systems mainly rely on syntactic and simple, mostly WordNet-based semantic similarity measures. The analysis of the similarity values these basic measures produce for the correspondences from the Process Model Matching Contest 2015 further illustrated that these measures are not suitable for identifying the correspondences that could not been identified by any of the participating matching techniques. The main reason is that neither the employed syntactic nor the employed semantic similarity measures were able to detect the complex semantic relationships that exist between activity labels.

To provide a basis for improving the performance of process model matching techniques in the future, we derived three specific recommendations. They address the main weaknesses we identified in the context of our analyses. First, we recommend using syntactic technology for preprocessing only. We found that syntactic measures can be useful for filtering highly identical labels but, beyond that, are often responsible for noise. Second, we recommend applying semantic technology beyond the word level. We observed that especially compound nouns and verb-noun combinations require specific attention. Third, we recommend using domain-specific dictionaries. Our analysis revealed that many of the missing correspondences contain words that are unlikely to be covered by general purpose resources such as WordNet.

We hope that the insights and recommendations we provide in this paper can represent valuable directions for future research on process model matching. We plan to build on the insights of this paper by developing a new matching technique that combines distributional similarity technology with a domain ontology.

References

1. La Rosa, M., Dumas, M., Uba, R., Dijkman, R.: Business process model merging: An approach to business process consolidation. ACM Transactions on Software Engineering and Methodology **22**(2) (2013) 11
2. Weidlich, M., Mendling, J., Weske, M.: A foundational approach for managing process variability. In: Advanced Information Systems Engineering, Springer (2011) 267–282
3. Jin, T., Wang, J., La Rosa, M., Ter Hofstede, A., Wen, L.: Efficient querying of large process model repositories. Computers in Industry **64**(1) (2013) 41–49
4. Awad, A., Polyvyanyy, A., Weske, M.: Semantic querying of business process models. In: Enterprise Distributed Object Computing Conference, 2008. EDOC'08. 12th International IEEE, IEEE (2008) 85–94
5. Dumas, M., GarcíA-BañUelos, L., La Rosa, M., Uba, R.: Fast detection of exact clones in business process model repositories. Information Systems **38**(4) (2013) 619–633

6. Uba, R., Dumas, M., García-Bañuelos, L., La Rosa, M.: Clone detection in repositories of business process models. In: Business Process Management. Springer (2011) 248–264
7. Weidlich, M., Dijkman, R., Mendling, J.: The ICoP Framework: Identification of Correspondences between Process Models. Advanced Information Systems Engineering (2010) 483–498
8. Leopold, H., Niepert, M., Weidlich, M., Mendling, J., Dijkman, R., Stuckenschmidt, H.: Probabilistic optimization of semantic process model matching. In: International Conference on Business Process Management, Springer (2012) 319–334
9. 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 (2013) 211–218
10. Sonntag, A., Hake, P., Fettke, P., Loos, P.: An Approach For Semantic Business Process Model Matching Using Supervised Machine Learning. In: Resarch in Progress Papers. 47. (2016)
11. Cayoglu, U., Dijkman, R., Dumas, M., Fettke, P., García-Bañuelos, L., Hake, P., Klinkmüller, C., Leopold, H., Ludwig, A., Loos, P., Mendling, J., Oberweis, A., Schoknecht, A., Sheetrit, E., Thaler, T., Ullrich, M., Weber, I., Weidlich, M. In: Report: The Process Model Matching Contest 2013. Springer (2013) 442–463
12. Antunes, G., Bakhshandeh, M., Borbinha, J., Cardoso, J., Dadashnia, S., Di Francescomarino, C., Dragoni, M., Fettke, P., Gal, A., Ghidini, C., et al.: The process model matching contest 2015. In: Proceedings of the 6th International Workshop on Enterprise Modelling and Information Systems Architectures. (2015)
13. Weidlich, M., Sagi, T., Leopold, H., Gal, A., Mendling, J.: Predicting the quality of process model matching. In: Business Process Management. Springer (2013) 203–210
14. Gal, A.: Uncertain schema matching. Synthesis Lectures on Data Management **3**(1) (2011) 1–97
15. Ehrig, M., Koschmider, A., Oberweis, A.: Measuring similarity between semantic business process models. In: Proceedings of the Fourth Asia-Pacific Conference on Conceptual Modelling. (2007) 71–80
16. van Dongen, B., Dijkman, R., Mendling, J.: Measuring similarity between business process models. In: Advanced Information Systems Engineering, Springer (2008) 450–464
17. Dijkman, R., Dumas, M., García-Bañuelos, L.: Graph matching algorithms for business process model similarity search. In: Business Process Management. (2009) 48–63
18. Dijkman, R., Dumas, M., Garcia-Banuelos, L., Kaarik, R.: Aligning business process models. In: Enterprise Distributed Object Computing Conference, IEEE (2009) 45–53
19. Dijkman, R., Dumas, M., Van Dongen, B., Krik, R., Mendling, J.: Similarity of business process models: Metrics and evaluation. Information Systems **36**(2) (2011) 498–516
20. Yan, Z., Dijkman, R., Grefen, P.: Fast business process similarity search with feature-based similarity estimation. In: On the Move to Meaningful Internet Systems, Springer (2010) 60–77
21. Niemann, M., Siebenhaar, M., Schulte, S., Steinmetz, R.: Comparison and retrieval of process models using related cluster pairs. Computers in Industry **63**(2) (2012) 168–180
22. Klinkmüller, C., Leopold, H., Weber, I., Mendling, J., Ludwig, A.: Listen to me: Improving process model matching through user feedback. In: Business Process Management, Springer (2014) 84–100
23. Liu, K., Yan, Z., Wang, Y., Wen, L., Wang, J.: Efficient syntactic process difference detection using flexible feature matching. Asia-Pacific Conference on Business Process Management. Springer International Publishing (2014) 103–116
24. Fengel, J.: Semantic technologies for aligning heterogeneous business process models. Business Process Management Journal **20**(4) (2014) 549–570

25. Ling, J., Zhang, L., Feng, Q.: Business Process Model Alignment: An Approach to Support Fast Discovering Complex Matches. Proceedings of the I-ESA Conferences (2014) 41–51
26. Makni, L., Haddar, N.Z., Ben-Abdallah, H.: Business process model matching: An approach based on semantics and structure. 12th International Joint Conference on e-Business and Telecommunications (ICETE) **2** (2015) 64–71
27. Sebu, M.L.: Merging business processes for a common workflow in an organizational collaborative scenario. System Theory, Control and Computing (2015) 134–139
28. Belhoul, Y., Haddad, M., Duchêne, E., Kheddouci, H.: String comparators based algorithms for process model matchmaking. In: Ninth International Conference on Services Computing, IEEE (2012) 649–656
29. Humm, B.G., Fengel, J.: Semantics-based business process model similarity. Business Information Systems (2012) 36–47
30. Weidlich, M., Sheetrit, E., Branco, M.C., Gal, A.: Matching business process models using positional passage-based language models. Proceedings of the 32nd International Conference on Conceptual Modeling (2013) 130–137
31. Corrales, J., Grigori, D., Bouzeghoub, M.: BPEL Processes Matchmaking for Service Discovery. In: On the Move to Meaningful Internet Systems. (2006) 237–254
32. Koschmider, A., Oberweis, A.: How to detect semantic business process model variants? In: Proceedings of the 2007 ACM Symposium on Applied Computing. (2007) 1263–1264
33. Pittke, F., Leopold, H., Mendling, J., Tamm, G.: Enabling reuse of process models through the detection of similar process parts. In: Business Process Management Workshops. (2013) 586–597
34. Sebu, M.L.: Similarity of Business Process Models in a Modular Design. Applied Computational Intelligence and Informatics, IEEE (2016) 31–36
35. Miller, G.A.: Wordnet: A lexical database for english. Communications of the ACM **38**(11) (1995) 39–41
36. Kim, S.N., Baldwin, T.: Automatic interpretation of noun compounds using wordnet similarity. In: International Conference on Natural Language Processing, Springer (2005) 945–956
37. Kim, S.N., Baldwin, T.: Interpreting semantic relations in noun compounds via verb semantics. In: Proceedings of the COLING/ACL on Main conference poster sessions, Association for Computational Linguistics (2006) 491–498
38. Bruni, E., Tran, N.K., Baroni, M.: Multimodal distributional semantics. J. Artif. Intell. Res.(JAIR) **49**(2014) (2014) 1–47
39. Leopold, H., Meilicke, C., Fellmann, M., Pittke, F., Stuckenschmidt, H., Mendling, J.: Towards the automated annotation of process models. In: Advanced Information Systems Engineering, Springer (2015) 401–416
40. Kietz, J.U., Volz, R., Maedche, A.: Extracting a domain-specific ontology from a corporate intranet. In: Proceedings of the 2nd workshop on Learning language in logic and the 4th conference on Computational natural language learning-Volume 7, Association for Computational Linguistics (2000) 167–175