Automatic Business Process Model Translation with BPMT

Kimon Batoulis[1], Rami-Habib Eid-Sabbagh[1], Henrik Leopold[2], Mathias Weske[1], and Jan Mendling[3]

[1] Hasso-Plattner-Institute, University Potsdam, Germany
kimon.batoulis@student.hpi.uni-potsdam.de
rami.eidsabbagh|mathias.weske@hpi.uni-potsdam.de
[2] Humboldt-Universität zu Berlin, Berlin, Germany
henrik.leopold@wiwi.hu-berlin.de
[3] Wirtschaftsuniversität Wien, Vienna, Austria
jan.mendling@wu.ac.at

Summary. Nowadays, many enterprises use business process models for documenting and supporting their operations. As many enterprises have branches in several countries and provide similar services throughout the globe, there is high potential for re-using these process models. However, the language barrier is a major obstacle for the successful re-use of process models, especially in multi-national companies. In this paper, we address this problem by presenting the Business Process Model Translator (BPMT), a technique for the automated translation of business process models that eases the re-use of business process models and reduces redundant work in multi-national companies. It builds upon the state-of-the-art machine translation system Moses and extends it with word and translation disambiguation considering the context of the domain. As a result, the BPMT can successfully deal with the compact and special language fragments that are typically found in business process models. A two-fold evaluation with the BLEU metric and an expert survey showed improvements of our approach over Moses.

Key words: business process models, statistical machine translation, word sense disambiguation, translation disambiguation

1 Introduction

Business process model collections are important assets of companies. Large enterprises operate with a network of branches in multiple countries with different official languages. Delivering the same services in new locations, the transfer, adaptation, and re-use of their business processes represents a major advantage [1]. However, language barriers impede a straightforward adaptation and consequently the re-use of existing process models.

As an example, consider the following scenario. An employee of a new subsidiary may want to define the new branch's order process as this does not yet exist in the new location. Hence, he will search in the company's global process repository for a relevant business process. This imposes several problems. First, the employee might not be able to find a specific model in the global repository as the processes are modeled in a foreign language and the search in his local language will not deliver any search results. Second, even if the employee speaks other languages and finds an according process model, he will not be able to use it as the majority of his colleagues may only speak the local language. In order to adapt the process model for re-use in the whole branch, a manual translation of all process labels would be required. For a large process collection this results in a massive amount of work which is associated with considerable costs. Accordingly, a technique that facilitates the automated translation of business process models could significantly increase the re-use of process models and hence reduce redundant work efforts as well as costs. As a result, multinational communication and collaboration among employees could be strengthened and eased.

Nevertheless, prior research highlighted the challenges that are associated with the automatic analysis of natural language in process models [2]. Since process model labels do not contain full and grammatically correct sentences, the application of standard tools for natural language processing such as parsers turned out to be hardly possible. Some authors even recommended to avoid the application of natural language processing in process models because of these issues [3]. Accordingly, the straightforward application of machine translation techniques for the translation of process models does not represent a promising strategy.

In this paper, we present the Business Process Model Translator (BPMT), a technique for the automated translation of business process model labels. In order to deal with the specific challenges of the natural language in business process models, we extend the state-of-the-art machine translation system Moses. By introducing word and translation disambiguation our technique includes the context of the domain and hence yields more stable results than a naive application of Moses. We use activity labels as use case for translation as they constitute an important part of business process models. While the presented approach is language-independent, we demonstrate the applicability of the technique by translating process models from English to German.

The remainder of this paper is structured as follows. Section 2 gives a brief introduction to machine translation and the Moses translation system.

Section 3 describes problems concerning process model label translations and presents our Moses extension, the Business Process Model Translator (BPMT). Section 4 presents a two-fold evaluation of our translation approach, followed by related work in Section 5. Section 6 concludes the paper.

2 Background

The general objective of machine translation is the automatic mapping of text from a source language (like English) to a target language (like German). As it is the case for all tasks of natural language processing, this problem can be tackled by using linguistic or statistical information about the languages in question. Of course, it is also possible to use a combination of both [4, 5]. In this paper, however, we concentrate on statistical methods. These methods are based on the basic model of statistical machine translation consisting of two components, a *translation model* and a *language model* [4]. The former gives information about the probability of the translation of a string from a source to a target language and the latter the general probability of the target language string. The target language string that maximizes the product of those two probabilities is then chosen as the source language's translation.

Our approach, the Business Process Model Translator (BPMT), is based on the open source statistical machine translation toolkit *Moses* [6]. Moses builds on the idea of phrase-based translation, which means that the probabilities of the translation model are not given for individual words or entire sentences but for phrases, a syntactical group consisting of one or more words. Currently, this procedure represents the state-of-the-art in modern machine translation [4]. As an example, consider the translation of the German sentence *natürlich hat john spass am spiel* to English, where the overall translation consists of the combination of translations of individual phrases as illustrated in Fig. 1. Note that it is possible to reorder phrases during translation. Moses makes it possible

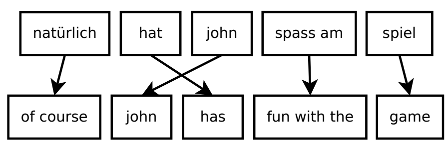

Fig. 1: Phrasal translation from German to English [6, p. 184].

to extend the basic model of statistical machine translation mentioned above with weighted *feature functions* to improve translation quality. The underlying equation is as follows:

$$\arg\max_T P(T|S) = \arg\max_T \prod_i h_i(T, S)^{\lambda_i} \qquad (1)$$

where S is the source language, T the target language, h_i the feature functions (e.g. the translation model), and λ_i the corresponding weights.

The translation system created for the BPMT was trained according to the description of building a "baseline system" on the Moses website[1]. This system's model consists of six overall features [7, 8]. Both the phrase translation model and the language model were trained using the *News Commentary* corpus[2] as there is no corpus for the business process management domain.

[1] `http://www.statmt.org/moses/?n=Moses.Baseline`.
[2] `http://www.statmt.org/wmt12/translation-task.html`.

3 Business Process Model Translator

When using a translation such as Moses and applying it for the translation of activity labels in business process models, the quality of the translation will certainly not be satisfying. This has one major reason, namely the fact that the language style used in business process models is not representative of the language style of the corpora on which the translation system was trained. On the one hand, business process models use very specific vocabulary that has to be interpreted in the context of the process model's domain. On the other hand, the language is very compact and uses *recurring patterns* of sentence fragments only. As an example, consider the label *creation of master record for tangible assets*. It contains phrases like *master record* and *tangible assets* whose translations are even hard for humans to obtain without complete knowledge of the business context. Moreover, these phrases *must* be translated as a whole since *master* as an individual word has a very different translation than the phrase *master record*. The trained translation system of the previous section would translate this label to German as *Schaffung von Herrn für Sachanlagen*, where the phrase translations are the following:

- *creation of* → *Schaffung von*
- *master* → *Herrn*
- *record for* → *für*
- *tangible assets* → *Sachanlagen*

The system managed to translate the phrase *tangible assets* as a whole, yielding a good translation of this phrase. But *master* was translated in isolation. *Herrn* (Engl. *Mister*) might be a good translation of this word in general, but in this context it is not. The system did not have another choice because the phrase *master record* does not occur in the News Commentary corpus. This example also demonstrates that the phrase alignment illustrated in Fig. 1 is not perfect. It sometimes makes questionable alignments like *record for* → *für*. With this alignment, only *for* is translated, ignoring the word *record*.

Unfortunately, the problem of data sparseness is not restricted to uncommon vocabulary. Even the phrase *check availability* neither occurs in the News Commentary nor in Europarl, another freely available language corpus. The only phrase in the Europarl corpus that is similar to *check availability* is *checking the availability*. Consequently, the words have to be translated individually, yielding *Schach Verfügbarkeit* (Engl. *chess availability*) as the *most probable* translation to German. In this case, the ambiguous word *check* was translated to *Schach* (Engl. chess) based on a wrong meaning.

3.1 Moses Translation Candidates

To solve this problem the BPMT makes use of the fact that Moses tries to find the most probable translation from a list of *translation candidates*. Table 1 shows an excerpt of the possible translations of *check availability* sorted by their probability (highest first) according to Moses. This table demonstrates the fact

Table 1: Translation candidates of *check availability* according to Moses

#	**Translation Candidate**
1	Schach Verfügbarkeit
2	Schach halten Verfügbarkeit
...	...
51	Verfügbarkeit überprüfen
...	...
553	Kreditfähigkeit Einkommenssegmenten
...	...

that although phrase translation was not possible there is a correct word translation including word re-ordering (#51) in which *check* was correctly translated to *überprüfen* instead of *Schach* (Engl. *chess*). How can we select #51 instead of #1? Our approach is based on the idea of *word sense disambiguation*, finding the meaning of a word in a given context. For example, we would have to find out that *check* is used as in *check the brakes* (translation *überprüfen*) and not as in *the chess player's king is placed in check* (translation *Schach*). This involves two things: First, the part-of-speech (POS) of *check* (whether it is a noun, verb, adjective etc.) needs to be identified. Second, after having determined that *check* is used as a verb, the actual disambiguation must be done. POS tagging of the words of a label is done on the basis of a label refactoring tool described in [2]. The refactorer assigns the categories *action*, *business object* and *addition* to the words of the activity labels in business process models (see Table 2). Moreover, it converts the labels to the recommended verb-object style [9]. This information

Table 2: Assignment of categories to words of labels in business process models

Label	**Action**	**Business Object**	**Addition**
Creation of master record for tangible assets	create	master record	for tangible assets
program analysis	analyse	program	

is used to infer the words' part-of-speeches: actions are declared as verbs and business objects as nouns. Additions can contain various part-of-speeches, thus it would be unwise to assign a single one to them. Instead, the *Stanford Part-of-Speech Tagger*[3] is used to assign the (potentially) correct POS. The tagger is not used for all words as taggers require *syntactically correct* context to give accurate results, which labels of process models typically do not provide.

Knowing the meaning of a word in a given context, provides valuable information about its correct translation, which is obvious when looking at the exemplary usage examples of *check* above.

[3] `http://nlp.stanford.edu/software/tagger.shtml`.

3.2 Finding the Best Translation

BPMT's algorithm for finding the best translation will be explained using the example label *check availability*. First, *check* and *availability* are disambiguated. This requires possible definitions of these words, which can be obtained from the English lexical database *WordNet*[4]. Basically, WordNet groups words together based on their meanings. Thus, it is possible to ask WordNet for (i) a definition of a word and (ii) other words that are related to it. Those relations include synonymy, hypernymy[5], hyponymy[6] and many more.

To find the correct definition in a given context, the Lesk algorithm [10] is used. For the BPMT it was implemented similar to [11]. The implementation determines the overlap of the WordNet definition from (i) and (ii) with its context. The BPMT defines the context of a label of a business process model as the concatenation of all the labels of the model. This is acceptable in this situation as we are mainly interested in the words themselves and not their syntactic combinations as it is the case for POS taggers. BPMT then scores the overlaps according to $\text{score} = \sum_{\text{overlap}} \text{length}(\text{overlap})^2$. This gives a higher score to longer *consecutive* overlaps. Thus, a definition that has one one-word sequence overlap and one two-word sequence overlap with the context gets a score of 5.

Having obtained the correct definitions of *check* and *availability*, Moses is asked to provide a list of possible translations of each definition. This helps in extracting the correct translation of *check availability* from the list in Table 1. Again, the algorithm for disambiguating words just described is employed. But now its task is to disambiguate *translations*.

Drawing analogies to the different steps of word sense disambiguation makes this point clear: The input to a word sense disambiguation algorithm is a word and its context. WordNet is asked for possible definitions of the word. The output is the correct definition given the context. Analogously, the input to the BPMT translation disambiguation is a label and all the translations of the definitions of the words of this label. Moses is asked for possible translations of the label. The output is the correct translation given the translations of the definitions.

3.3 BPMT Architecture

The BPMT system consists of two core modules: Moses and the word/translation disambiguator. These modules take care of the translation and the selection of the best result. To feed the BPMT with business process model labels it was integrated into PromniCAT, a platform for research on process model collections [12][7]. Among others, the platform provides utility units to extract business process models and their activity labels from business process model repositories like the BPMAI[8]. The architecture of the BPMT is depicted in Fig. 2 as a UML

[4] `http://wordnet.princeton.edu/`.

[5] A more general word: *vehicle* is a hypernym of *car*.

[6] A more specific word: *car* is a hyponym of *vehicle*.

[7] `http://code.google.com/p/promnicat/`.

[8] `http://bpt.hpi.uni-potsdam.de/BPMAcademicInitiative/WebHome`.

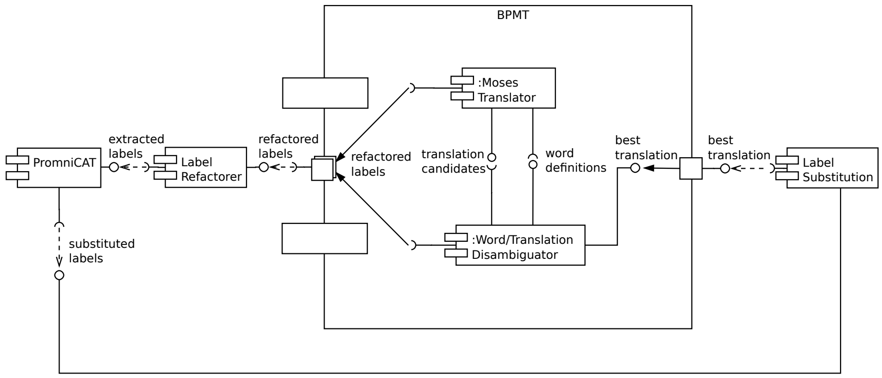

Fig. 2: The component architecture of the BPMT.

component diagram. Business processes of the BPMAI modeled in BPMN 2.0 are selected from the research platform's process repository. The activity labels are extracted and passed through the label refactoring tool described in Section 3. The refactored labels are then handed over to the BPMT where Moses interacts with the translation and word sense disambiguation module in order to find the best translation. Finally, the original label is replaced by its translation in the process model and the translated model is stored as a new revision in the research platform's repository.

4 Evaluation

An interesting question concerning the translations of the BPMT is how probable it is that they are different from the translations favored by Moses. In those cases the actual translation is not the first of the list of translation candidates but some other translation in that list. In order to answer this question, we let the two systems translate a subset of the activity label data set from the BPMAI, totaling 2084 labels. The evaluation revealed that nearly 63% of the Moses translations are discarded. Table 3 shows some of those differences. Basically, there are four classes of results: Rows (1), (2), (3,4) and (5).

The first row shows that there are cases in which the BPMT performs better than Moses. In this case, the BPMT disambiguated *receive* correctly and then chose a better translation for it based on the translations of the definition of *receive* and their overlap with the translation candidates of *receive* provided by Moses. The translation *eine der* for *receive* by Moses (which is incomprehensible) is due to a misalignment during training.

Yet, in some cases, as in row 2, translation quality decreases. This is due to the fact that the BPMT wrongly disambiguated *produce* as *bring out for display*— instead of *create or manufacture a man-made product*. Now, the translation lists of both *bring out for display* and *produce dog food* contain translations with the word ***bringen***. For example, in the first case *für die Tag* ***bringen*** and in the

Table 3: Comparison of Moses and BPMT translations

Label	**Moses**	**BPMT**
receive Mail	einer der Mail	Mail erhalten
produce dog food	Nahrungsmittel zu produzieren Hund	Hund Nahrungsmittel bringen
buy book	kaufen Buch	Buch erwerben
send decision	schicken Entscheidung	Entscheidung schicken
create leased asset master record	verleaste schaffen , die Meister Bilanz	einem gepachteten Stück Vermögenswerte Meister Bilanz schaffen

second case *Hund Nahrungsmittel* ***bringen***. Thus, the BPMT concluded that *Hund Nahrungsmittel bringen* is the correct translation.

BPMT corrects syntactic structures of phrases. This is represented in another class that consists of rows 3 and 4, where all translations are fairly accurate. In this class the translation of an individual word changes slightly (in row 3, *kaufen* is a synonym of *erwerben*) or the syntactic structure of a phrase was changed. E.g., in German it is wrong if a verb in its infinitive form precedes the noun. The BPMT uses the Stanford Tagger to identify if a verb is the first word of a label. In those cases the verb is moved to the end of the label. Consequently, *erwerben Buch* will be changed to *Buch erwerben*. Also, the BPMT removes unnecessary elements for process model labels like the honorific form and commas, among others. For example, the BPMT removes the honorific form *Sie* from the translation of *send decision* (row 4) to *schicken Sie Entscheidung* and later switches words as the first word of the translated label is a verb. The resulting translation *Entscheidung schicken* is very good.

The last class demonstrates a substantial problem: If the language used in the label is so specific and compact as in row 5, the list of possible translations will not contain any correct result that could be identified by the BPMT. The phrase *asset master record* must be encountered during training, so that it can be translated as a whole. As neither the News Commentary nor the Europarl corpus contain this phrase, both translations in row 5 are fairly bad.

4.1 BLEU Score

The analysis of the different translations of the 2084 labels suggested that the BPMT performs better than Moses. To validate the results we quantitatively evaluated the BPMT and Moses translations using the BLEU Score [13]. It is computed using an algorithm that scores the translation system output by the number of N-gram overlaps with a reference translation that has to be created manually. Three reference translations were created by us for 207 of the 2084 labels as no reference translations of models from process repositories exist.

When using the BLEU algorithm one has to decide how long the sequence of overlapping words (i.e. the size of the N-gram) should be. Smaller N-grams yield higher scores, but may also be less meaningful. For example, when using 1-grams,

Table 4: BLEU Scores for BPMT and Moses translations

	1-grams	**2-grams**	**3-grams**
BPMT	38.96	22.95	14.94
Moses	38.63	21.72	13.82

the words that the two translations have in common can occur randomly in the text. In the case of 2-grams they have to occur in common two-word sequences. Since 2-word activity labels are very common in process models, the N-gram scores were computed for $N = 1, 2, 3$. The results are shown in Table 4. As one can see, the two translation systems perform nearly equally well in the choice of their words (1-gram score). Note that this does not imply that they choose the same words. However, larger N-grams lead to a more significant difference between the scores. This means that the BPMT translations are more *fluent* and thus better understandable than the Moses translations.

4.2 Expert Survey

Despite the clear difference between the BLEU scores, this does not represent the actual experienced improvement for two reasons: First, BLEU is rather suited to conduct an evaluation of translations of entire corpora, averaging out errors on individual sentences [13, 14]. In contrast, our evaluation data set is comparatively small and only consists of activity labels, not sentences. Second, BLEU does not consider synonymy and other forms of semantic relations—only identical words are scored.

For those reasons we conducted an expert evaluation, in which 20 subjects with a strong BPM background were asked to assess the quality of the translations of both Moses and BPMT. The survey included translations of six randomly selected process models from the BPMAI and was structured as follows: For each process model, the subjects were first asked to rate the translations of the entire models. Afterwards, they rated the translations of the model's labels that differed between Moses and BPMT. The six process models taken together contained 77 labels, 54 (70%) of which were translated differently by Moses and BPMT. The translations were rated using a Likert scale, ranging from 1 (*not acceptable*) to 7 (*excellent*). The ratings are visualized as box plots in Fig. 3a that shows the plots of the ratings of the entire models and Fig. 3b that shows the ratings of the differing labels. The plots illustrate that in both cases the median of the BPMT is one value higher than that of Moses. In addition to the higher BLEU score, this supports the impression that translation quality increases when the BPMT is used to translate process models instead of Moses.

To prove that there is a significant difference between the medians of the translation quality ratings, we conducted the nonparametric sign test. This test makes no assumptions about the distribution of the population from which the samples are drawn. The results show that for both the ratings of the entire models as well as the individual labels we can reject the hypothesis that there is no difference between the ratings with a confidence of 99.9%.

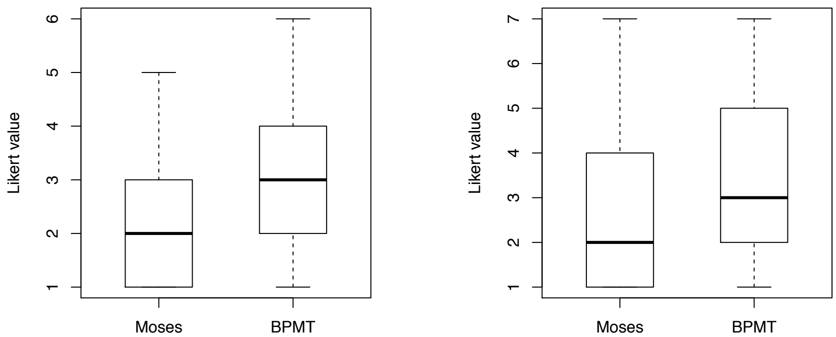

(a) Entire Models (b) Individual Labels

Fig. 3: Box-plots of translation quality ratings for Moses and BPMT.

In summary, an improvement to the overall translation quality has been achieved, but more work has to be devoted to making it more stable so that *every* translation is understandable.

5 Related Work

The work presented in this paper can be related to two major streams of research: statistical machine translation and natural language processing in process models.

Applying natural language processing in process models that pays attention to the structure of process model labels is a quite recent endeavor. In this context, techniques for the enforcement of naming conventions [15] or the refactoring of activity labels have been proposed [2]. Particularly, the latter technique is an important foundation for the technique proposed in this paper. Only by explicitly building on the label components such as action and business object, the translation of the state-of-the-art machine translation system Moses could be improved. Prior works which did not consider the structure of process model labels have been mainly concerned with improving the language in terms of the consistent usage of words [16, 17, 18]. These works represent complementary techniques to the translation approach as they can help to assure a consistent use of natural language before the model is translated.

Statistical machine translation is nowadays dominated by a phrase-based approach relying on huge bilingual corpora on which the translation engine is trained without using any linguistic information [4]. However, there are several works that try to improve the quality of translations on the basis of semantic information. For example, [19, 20] report that they achieve a statistically significantly higher BLEU score by disambiguating the phrases of the source text and thus selecting a translation that matches the obtained definitions. In [21] a system is described that reduces the amount of training data necessary to build a well-performing translator. This is achieved by creating models of the interdependencies of related inflected word forms, which can also help in finding the right word form during translation if not enough context information is available. Also, methods for language-specific sentence-level restructuring transformations are applied after the translation has been obtained. Finally, in [22] a very large amount of context of the target language is used to find and rank N-grams containing the phrase translations of the source text. However, to the best of our knowledge, there are no works that systematically evaluate their techniques using very short phrases only, like those in business process models.

6 Conclusion and Future Work

This paper presented the BPMT, a tool for automatic translation of activity labels of business process models to facilitate their re-use in an international context. Our use case, the translation from English to German, can be easily adapted to different languages by training the BPMT on corpora of the languages in question. The presented concept of word and translation disambiguation does not change. The functionality of the correction of syntactic structure described in Section 4 can be adjusted to the rules of the target language as well. BPMT extends Moses to the specific BPM domain to improve translation quality. Our evaluation results showed that BPMT performs better than Moses by itself.

In future work, two aspects need to be considered: First, language specific tuning of the algorithm could further improve translation quality. For example, moving a verb from the beginning of a label to the end, can lead to problems if the label contains a conjunction of actions. Second, more generally, training of a *factored* translation system [6] should be taken into consideration.

References

1. Indulska, M., Recker, J., Rosemann, M., Green, P.: Business process modeling: Current issues and future challenges. In: Advanced Information Systems Engineering, Springer (2009) 501–514
2. Leopold, H., Smirnov, S., Mendling, J.: On the refactoring of activity labels in business process models. Inf. Syst. **37**(5) (July 2012) 443–459
3. Becker, J., Bergener, P., Breuker, D., Räckers, M.: An empirical assessment of the usefulness of weakness patterns in business process redesign. In: 20th European Conference on Information Systems (ECIS 2012). (2012) Barcelona, Spain.
4. Jurafsky, D., Martin, J.H.: Speech and Language Processing (2nd Edition) (Prentice Hall Series in Artificial Intelligence). 2 edn. Prentice Hall (2008)
5. Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach. 2 edn. Prentice-Hall, Englewood Cliffs, NJ (2003)
6. Koehn, P.: MOSES, Statistical Machine Translation System, User Manual and Code Guide (2010)
7. Koehn, P., Och, F.J., Marcu, D.: Statistical phrase-based translation. In: Proceedings of the 2003 Conference of the NAACL on Human Language Technology - Volume 1. NAACL '03, Stroudsburg, PA, USA, ACL (2003) 48–54

8. Axelrod, A., Mayne, R.B., Callison-burch, C., Osborne, M., Talbot, D.: Edinburgh system description for the 2005 iwslt speech translation evaluation. In: In Proc. International Workshop on Spoken Language Translation (IWSLT. (2005)
9. Mendling, J., Reijers, H.A., Recker, J.: Activity labeling in process modeling: Empirical insights and recommendations. Inf. Syst. **35**(4) (June 2010) 467–482
10. Lesk, M.: Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. In: Proceedings of the 5th annual international conference on Systems documentation. SIGDOC '86, New York, NY, USA, ACM (1986) 24–26
11. Banerjee, S., Pedersen, T.: An adapted lesk algorithm for word sense disambiguation using wordnet. In: Proceedings of the Third International CICLing. CICLing '02, London, UK, UK, Springer-Verlag (2002) 136–145
12. Eid-Sabbagh, R.H., Kunze, M., Meyer, A., Weske, M.: A Platform for Research on Process Model Collections. In: BPMN2012 Workshop proceedings. (2012)
13. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting on ACL. ACL '02, Stroudsburg, PA, USA, ACL (2002) 311–318
14. Agarwal, A., Lavie, A.: Meteor, m-bleu and m-ter: evaluation metrics for high-correlation with human rankings of machine translation output. In: Proceedings of the Third Workshop on StatMT. StatMT '08, Stroudsburg, PA, USA, ACL (2008) 115–118
15. Becker, J., Delfmann, P., Herwig, S., Lis, L., Stein, A.: Towards Increased Comparability of Conceptual Models - Enforcing Naming Conventions through Domain Thesauri and Linguistic Grammars. In: ECIS 2009. (June 2009)
16. Koschmider, A., Blanchard, E.: User assistance for business process model decomposition. In: In First IEEE International Conference on Research Challenges in Information Science. (2007) 445–454
17. van der Vos, B., Gulla, J.A., van de Riet, R.: Verification of conceptual models based on linguistic knowledge. Data and Knowledge Engineering **21**(2) (1997) 147 – 163
18. Dominic Breuker, Daniel Pfeiffer, J.B.: Reducing the variations in intra- and interorganizational business process modeling - an empirical evaluation. In: Proceedings of the Internationale Tagung Wirtschaftsinformatik. (2009)
19. Carpuat, M., Wu, D.: Improving statistical machine translation using word sense disambiguation. In: In EMNLP-CoNLL 2007. (2007) 61–72
20. Chan, Y.S., Ng, H.T.: Word sense disambiguation improves statistical machine translation. In: In 45th Annual Meeting of the Association for Computational Linguistics (ACL-07. (2007) 33–40
21. Nießen, S., Ney, H.: Statistical machine translation with scarce resources using morpho-syntactic information. Comput. Linguist. **30**(2) (June 2004) 181–204
22. Carbonell, J., Klein, S., Miller, D., Steinbaum, M., Grassiany, T., Frey, J.: Context-based machine translation. In: In Proceedings of the 7th Conference of the AMTA. (2006) 8–12