



What's the Matter? Knowledge Acquisition by Unsupervised Multi-Topic Labeling for Spoken Utterances

Sebastian Weigelt, Jan Keim, Tobias Hey, and Walter F. Tichy

Karlsruhe Institute of Technology,
Institute for Program Structures and Data Organization,
Karlsruhe, Germany
{weigelt,jan.keim,hey,tichy}@kit.edu

Received (11/18/2019)

Revised (02/29/2020)

Accepted (03/12/2020)

Abstract. Systems such as Alexa, Cortana, and Siri appear rather smart. However, they only react to predefined wordings and do not actually grasp the user's intent. To overcome this limitation, a system must understand the topics the user is talking about. Therefore, we apply unsupervised multi-topic labeling to spoken utterances. Although topic labeling is a well-studied task on textual documents, its potential for spoken input is almost unexplored. Our approach for topic labeling is tailored to spoken utterances; it copes with short and ungrammatical input.

The approach is two-tiered. First, we disambiguate word senses. We utilize Wikipedia as pre-labeled corpus to train a naïve-bayes classifier. Second, we build topic graphs based on DBpedia relations. We use two strategies to determine central terms in the graphs, i.e. the shared topics. One focuses on the dominant senses in the utterance and the other covers as many distinct senses as possible. Our approach creates multiple distinct topics per utterance and ranks results.

The evaluation shows that the approach is feasible; the word sense disambiguation achieves a recall of 0.799. Concerning topic labeling, in a user study subjects assessed that in 90.9% of the cases at least one proposed topic label among the first four is a good fit. With regard to precision, the subjects judged that 77.2% of the top ranked labels are a good fit or good but somewhat too broad (Fleiss' kappa $\kappa = 0.27$).

We illustrate areas of application of topic labeling in the field of programming in spoken language. With topic labeling applied to the spoken input as well as ontologies that model the situational context we are able to select the most appropriate ontologies with an F_1 -score of 0.907.

Keywords: Topic Labeling; Topic Modeling; Unsupervised Machine Learning; Graph Centrality Measures; Word Sense Disambiguation; Ontology Selection; DBpedia; Wikipedia; Semantic Annotation; Spoken Language Interfaces; Spoken Language Understanding; Natural Language Processing;

1 Introduction

Conversational interfaces (CI) are a recent trend in human computer interaction. Today, millions of users communicate with virtual assistants such as Alexa, Cortana, or Siri. However, such systems often struggle to actually grasp the user’s intent. Although they appear rather smart, Alexa and the like merely react to predefined commands. Users will soon expect such systems to understand increasingly complex requests. Thus, techniques for (deep) spoken language understanding (SLU) are needed. We propose to apply topic labeling to spoken utterances as one building block of a comprehensive intent model. Topic modeling and labeling has already proved useful on textual documents; it has been applied to many tasks, such as text summarization, machine translation, and sentiment analysis [5]. However, topic labeling has rarely been adapted to spoken utterances scenarios [34]. Most likely this is the consequence of differing boundary conditions. Spoken language is typically ungrammatical. Thus, common techniques for natural language (pre-)processing (NLP) cannot be applied. Furthermore, utterances – be it dialog acts, virtual assistant interactions, or instructions for household robots – are short in comparison to text documents. This limits the usefulness of contextual information to a minimum. An exemplary input of that kind might be, “Hey robot, take – uhm – the apple – err – the orange from the fridge.” Even though the utterance is rather short, it encompasses three topics: *Domestic Robotics*, *Fruits*, and *Home Appliances*. Present approaches for topic labeling cannot cope with such conditions as they either rely on NLP or contextual models.

Our approach is influenced by a number of related approaches to topic labeling on documents. However, it is customized to the challenges of short, spoken utterances. Our approach is two-tiered. First, we perform word sense disambiguation (WSD). We have adapted the approach by Mihalcea [24] and Mihalcea and Csomai [25]. The approach uses Wikipedia as a pre-labeled corpus and applies a naïve-bayes classifier. Nouns are labeled with Wikipedia articles. Second, we use the word sense labels to determine topic labels. To this end, we build so-called sense graphs. Beginning with the Wikipedia articles attached to nouns in the utterance, we use relations in DBpedia to construct graphs. Afterwards, we determine the *most central terms*, which we take to be the topic labels for the utterance. We have implemented two different strategies for graph centrality. The first generates topic labels for dominant terms, i.e. the most frequent senses, in the utterance. The latter covers all terms. Both produce multiple labels for each utterance. The labels carry confidences, which we derive from the graph centrality value. The contribution of the paper is three-fold:

1. An adaptation of the WSD approach by Mihalcea and Csomai to short utterances, including an evaluation on a Wikipedia data set plus an additional evaluation on a corpus for programming in spoken language.
2. An implementation and evaluation of unsupervised multi-topic labeling tailored to short, spoken utterances.
3. An application of unsupervised multi-topic labeling to ontology selection for programming in spoken language, including an evaluation on several domains.

We also discuss further areas of application, such as context modeling, CyC micro-theory matching, and dialog interaction.

The remainder of the paper is structured as follows: First, we discuss related work in Section 2. In Section 3 we introduce our approach for unsupervised multi-topic labeling and evaluate it in Section 4. Finally, we discuss areas of application (Section 5) before we conclude the paper in Section 6.

2 Related Work

Topic labeling is typically preceded by a topic modeling step that determines sets of terms that are supposed to share the same topic. Afterwards, meaningful labels are assigned to these topics. Many approaches rely on the so-called *Latent Dirichlet Allocation (LDA)* introduced by Blei et al. [4] to create a topic model [3, 13, 15, 22, 23]. *LDA* is a generative probabilistic model for collections of discrete data such as text documents. It uses word distributions across a given set of documents to derive topics from word occurrences. Hence, an *LDA* topic model comprises a fixed number of topics that consist of words which often occur together.

To determine meaningful labels, some approaches derive labels directly from the given text [23, 26], assuming that a label can be found within the given text. However, this assumption may not hold. Often, a document does not contain appropriate labels; i.e. for certain topics no abstract term is ever mentioned. Additionally, text-based approaches usually suffer from challenges such as synonyms or spelling errors. Thus, advanced approaches incorporate additional information to gain a deeper understanding of a topic. Usually, these approaches map words that represent a topic to knowledge databases. Then, they create graph or tree structures based on relations in the knowledge database (e.g. Magatti et al. [22]). Another approach of that kind was introduced by Hulpus et al. [15]. The authors calculate a topic model with *LDA* and then determine central nodes in a so-called topic graph, which they build from DBpedia concepts and relations. Central concepts form the topic labels. All above-mentioned approaches use *LDA* to some extent, which is a statistical model. Therefore, its performance depends on the available amount of data. As spoken utterances are rather short, *LDA* does not produce reliable results. Hence, *LDA*-based approaches are infeasible in our context.

Some related approaches do not rely on *LDA*. Coursey et al. [8] create graphs based on Wikipedia articles (nodes) and the proximity of the containing words

(edges). They determine central nodes with the help of a biased PageRank algorithm and use the article names of these nodes as topic labels. Aker et al. [1] use the Markov Clustering Algorithm for topic modeling. Allahyari and Kochut [2] adapt *LDA*; they introduce a latent variable called *concept*. The concepts are DBpedia concepts and are used to build graphs. Recently, combined approaches are used; they either join different topic labeling approaches (e.g. Gourru et al. [11]), or incorporate concepts from other research areas, e.g. word embeddings [31]. However, they also require long documents to unfold their potential. Thus, they are inappropriate for short, spoken utterances.

In the field of SLU various approaches use phonetic information to model topics. Cerisara [6] creates a semantic lexicon from phonetic information and creates topic models by hierarchical clustering. Hazen et al. [12] and Siu et al. [32] propose similar approaches to model topics. However, none of the approaches actually label topics.

In summary, all above are inapplicable to determine topic labels for short utterances. *LDA*-based approaches (and others intended for texts) require long documents and present SLU approaches only model topics but do not determine labels.

3 Approach

Our approach for unsupervised multi-topic labeling is inspired by topic modeling and labeling approaches for text documents. However, it does not rely on a generative probabilistic model such as *LDA*. This is mainly because *LDA* is not applicable on short documents. Additionally, *LDA* can only distinguish a fixed number of topics. However, in our context the number of topics is uncertain in advance. Unlike *LDA*-based approaches, we build topic graphs for the entire input, i.e. each spoken utterance. We use data from DBpedia to create these graphs; articles are nodes and relations form edges. We use a biased PageRank algorithm to determine multiple central articles per utterance, which we use as topic labels. Therefore, we are relieved from the challenge of creating meaningful labels. Instead, we only have to determine which term is the most fitting for a topic. The approach requires word sense labels as starting point for the construction of sense graphs. For this, we adapt the approach by Mihalcea and Csomai [25] that uses Wikipedia as a pre-labeled corpus for WSD. It uses naïve-bayes classification to attach Wikipedia articles (as senses) to nouns.

In Subsection 3.1 we present our adapted re-implementation of their WSD method. Afterwards, we describe our unsupervised multi-topic labeling approach for spoken utterances in detail in Subsection 3.2.

3.1 Word Sense Disambiguation

Supervised classification tasks require manually attached labels for training, which is time-consuming and costly. Additionally, in the case of word sense disambiguation human annotators often disagree. Mihalcea and Csomai tackle this issue by using Wikipedia as a pre-labeled corpus for word senses. The basic idea is as follows. Relevant terms (mostly nouns) in a Wikipedia article each

have a link attached to the respective explanatory article. Thus, links can serve as manually annotated word senses.

Links are added by the article’s authors (most commonly), who are supposed to be domain experts. Therefore, Mihalcea and Csomai assume that the links are correct. Also, Wikipedia is growing steadily and the quality of articles improves over time through continuous inspection by the community. Even though the latter is arguable, the quality of Wikipedia articles surely has improved since Mihalcea and Csomai first implemented their approach in 2007. Further details about the original approach may be found in Mihalcea [24] and Mihalcea and Csomai [25].

We adopt the idea to use Wikipedia as a pre-labeled corpus for word senses. However, we altered the classification process slightly. We also use a naïve-bayes classifier and similar features: the ambiguous word, its part-of-speech (POS) tag, the three words to the left and right of the ambiguous word, and their POS tags, as well as the first nouns and verbs to the left and right. To increase the impact of the ambiguous word over its contextual features we weighted it tenfold (contrary to Mihalcea and Csomai, who did not alter weights). We filter out stop-words. Mihalcea and Csomai additionally use so-called context words, which are simply the most frequent words of the paragraph in which the ambiguous word appears. Context words are not feasible in our context, because short, spoken utterances do not consist of paragraphs. Even if we define a full utterance as paragraph, it is rather short with barely multiple occurrences of words that are not stop-words. We also skip the disambiguation of named entities. As they are usually unambiguous, there is no need to disambiguate these terms and their mere number impairs our classification model.

To train the classifier, we use a Wikipedia dump from August 2017. We prepare the data like Mihalcea and Csomai. We remove disambiguation pages, as they do not contain full sentences. For the same reason, we ignore info boxes and lists. Additionally, lists rarely contain links, which makes them useless. The same applies to quotes. We also remove links that lead to an article that considers a named entity; those are simply unusable topic labels.

We extracted 5,188,470 training instances. Among them are 283,173 different senses, of which 136,964 are unique. Unique senses are senses that are mentioned in one instance only. These unique senses account for 2.64% of the instances and 48.37% of the senses.

We can then use the trained model as a WSD classifier. Note that the classifier can only disambiguate nouns. However, contrary to Mihalcea and Csomai, our classifier attaches a label – i.e an Wikipedia article – to *all* nouns in the input.

3.2 Topic Labeling

Our approach for unsupervised multi-topic labeling is inspired by the sense graph idea proposed by Hulpus et al. [15]. However, they used *LDA* to determine topic models. As discussed before, we cannot use a statistical model such as *LDA* on short, spoken utterances. Instead, we directly determine topic labels and perform topic modeling implicitly. We assume that all nouns in the input are related at first (we discard that assumption later). We build sense graphs beginning with

Table 1: DBpedia relations used to build sense graphs.

relation	relates a concept to
dcterms:subject	its Wikipedia category
skos:broader	less specific concepts
skos:narrower	more specific concepts
purlg:hypernym	superordinate concepts
purlg:meronym	concepts that form parts
purlg:synonym	synonymous concepts
rdfs:type	its DBpedia ontology entity
rdfs:subClassOf	its subclasses in DBpedia
rdfs:seeAlso	related concepts

the word sense for each noun; we call these senses *initial senses*. We think of DBpedia as a graph with concepts (i.e. articles) as nodes and relations as edges and extract subgraphs. We traverse all chains of relations up to a distance of two to create the sense graphs¹. The relations we use to build the sense graphs are listed in Table 1.

Finally, we merge all sense graphs; the result is a topic graph. Hulpus et al. remove all disconnected subgraphs and proceed with the main graph only. Instead, we proceed with the entire graph, including all disconnected subgraphs. As each subgraph originates from different sense graphs, we assume that subgraphs represent different topic areas. Thus, subgraphs compensate the missing topic modeling step in our approach. Then, we can determine topics for each of the subgraphs, i.e. topic areas.

We continue with determining the central nodes of the graph. Hulpus et al. discuss different algorithms to determine graph centrality. However, none of them can cope with disconnected graphs. Instead, we apply a biased PageRank algorithm [7, 8]; it gives more weight to nodes that correspond to the initial senses. The biased PageRank, i.e. the score $S(V_i)$, is calculated as follows:

$$S(V_i) = (1 - d) * B(V_i) + d * \sum_{j \in I(V_i)} \frac{S(V_j)}{|O(V_j)|} \quad (1)$$

where $I(V_i)$ is the set of incoming edges of node i and $O(V_i)$ is the set of outgoing edges of node i . The constant d is the *damping factor*; in our implementation, we leave the default value unchanged ($d = 0.85$). The bias $B(V_i)$ is defined as:

$$B(V_i) = \frac{f(V_i)}{\sum_{j \in \text{InitNodes}} f(V_j)} \quad (2)$$

¹ Ideally, the distance is as short as possible to generate meaningful sense graphs. Longer distances introduce an increasing semantic drift. However, if we traverse only one relation, less connected graphs are constructed, i.e. we might be unable to discover shared senses. Therefore, we follow the choice of Hulpus et al. [15] and use two as the distance value.

where InitNodes is the set of nodes that correspond to the initial senses. Coursey et al. state, that $f(V_i)$ may vary regarding complexity and can be chosen freely. They also discuss different options for the choice of $f(V_i)$. For their approach they choose an $f(V_i)$ that was determined by a so-called *keyphraseness score*. However, this score requires a set of documents. In our context we consider only one utterance at once. Thus, we determine $B(V_i)$ differently; we simply set $f(V_i)$ to 1 if V_i is a member of the initial node set:

$$B(V_i) = \begin{cases} 0 & , V_i \notin \text{InitNodes} \\ \frac{1}{|\text{InitNodes}|} & , V_i \in \text{InitNodes} \end{cases} \quad (3)$$

Finally, we select the nodes from the topic graphs that will serve as topic labels. The selection of labels is contingent on the number of labels we create per utterance. On the one hand, with an increasing number of labels we observed that labels get too broad. Thus, the precision of our approach decreases. On the other hand, if we create too few labels, some senses from the utterance are not represented. As a consequence, the coverage decreases. We found that a good rule of thumb is to create twice as many labels as there are distinct senses in the input. This number of labels allows us to discover appropriate labels even for utterances with many distinct topics. At the same time, there is still a sufficient selection of labels for small inputs with only two or three senses.

To select labels our approach is configured with one of two strategies. The *top strategy* selects nodes as labels that are strongly connected to the original sense nodes. Therefore, we count the number of sense graphs in which each particular node occurs. We call this value *connectivity*. As the connectivity might be equal for multiple nodes, we use PageRank as second criterion. Therefore, the *top strategy* first selects the node with the highest connectivity. If there is a draw, the node with the higher PageRank is chosen. We repeat this procedure until the maximum number of labels is reached. However, we found that some parts of the utterances are inadequately represented. If a topic is mentioned with a few words only, it gets dominated by other topics. Therefore, we implemented a second strategy: the *max strategy*. This strategy determines the first label in the same manner as the *top strategy*. However, it then examines the set of senses from the initial input. If not all senses are represented through a label yet, it selects the node that covers the highest number of previously unrepresented senses. As before, if there is a draw, the node with the higher PageRank is chosen. If all senses are covered, the strategy continues as the *top strategy*. Again, the procedure is repeated until the maximum number of topic labels is reached. We use the PageRank values as confidences to rank the labels.

3.3 Example

In order to illustrate our approach, we discuss the exemplary utterance, “take the orange from the fridge and close the dishwasher.” For the sake of simplicity we configure our approach to create a maximum of two labels. The first step is the disambiguation of nouns. For *orange* the WSD model picks the sense

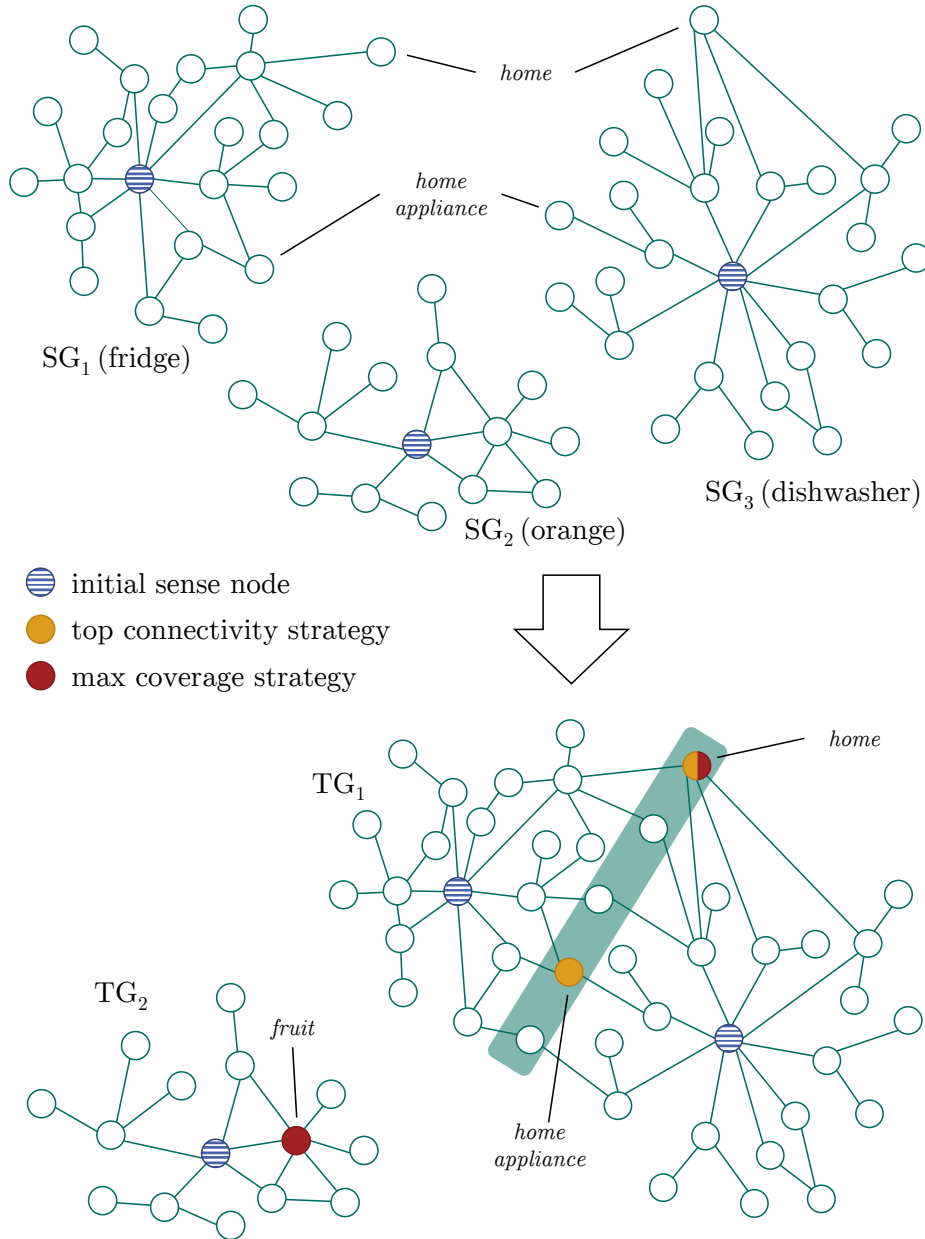


Fig. 1: The sense (SG_i) and topic graphs (TG_i) for the utterance, "take the orange from the fridge and close the dishwasher afterwards." The initial sense nodes are depicted blue-striped. The nodes selected as topic label by the *top strategy* are orange and the nodes selected by *max strategy* are red. The set of candidate nodes with equal connectivity in TG₁ are highlighted in green.

Orange(fruit) rather than other possible senses such as *Orange(color)* or *Orange(word)*. The approach determines *Refrigerator* and *Dishwasher* as senses for the two remaining nouns. Then, our approach creates a sense graph for each sense. The resulting sense graphs (SG₁, SG₂, and SG₃) are depicted in the upper half of Figure 1. Graphs may share sense nodes. In the example, the sense graphs SG₁ and SG₃ share – among others – the sense nodes *home* and *home appliance*. All sense graphs are merged at these shared nodes. The result is a topic graph that may consist of disconnected subgraphs. In the example, we have two topic subgraphs (TG₁ and TG₂) after merging. Connected sense graphs indicate that senses are topically similar. With the help of the topic graph we determine the labels. We obtain different results depending on the selected strategy. If we use the *top strategy*, graph connectivity is the key aspect. Thus, only sense nodes from the topic subgraph TG₁ are considered. From the set of sense nodes that connect the two sense graphs (highlighted in green) our approach selects the nodes with the highest PageRank (i.e. the yellow nodes). Thus, the *top strategy* creates the labels *home* and *home appliance*. However, the sense *Orange(fruit)* is not covered by these topic labels. Here, the result of the *max strategy* differs. First, it selects a node from those with the highest connectivity, too. Thus, the label *home* is also selected. However, the next label is drawn from the topic subgraph with senses that were not represented previously, here TG₂. Again, the node from the candidate set with the highest PageRank is chosen, here *fruit*. Thus, the *max strategy* selects the labels *home* and *fruit* (i.e. the red nodes).

4 Evaluation

To evaluate our approach we first assess the quality of the word sense disambiguation; its performance directly affects our unsupervised multi-topic labeling approach. As data sets we use Wikipedia and a speech corpus. The latter consists of 168 voice recordings from different user studies, gathered from 65 subjects. The subjects are between 18 and 50 years old, 21 are female and 44 male. Most of them are undergraduate and graduate students. All are non-native English speakers. However, their (self-assessed) English level is *advanced* on average. All recordings are instruction sequences for a household robot in eight different scenarios such as doing the laundry or preparing an instant meal. The recordings vary in length from 5 up to 80 seconds and in instructions from 2 up to 22. All recordings were manually transcribed according to the guideline by Kiesling et al. [17].

We also evaluate our topic labeling approach on this corpus. To broaden the range of topics, we added synthetic utterances from other domains. We conducted a user study, where subjects manually evaluated the quality of the topic labels.

4.1 Word Sense Disambiguation

We first evaluate WSD on Wikipedia. We performed a customized ten-fold cross-validation. For each of the ten runs we drew 10,000 instances at random for testing; the remaining were used for training. Note that a full-blown ten-fold cross-

validation with over five million instances is infeasible in our context. We determine the correctly predicted (true positives) and incorrectly predicted senses. As the number of instances is known in advance and our classifier predicts labels for all instances, every incorrectly predicted sense accounts for a false positive *and* a false negative. Therefore, precision and recall are the same here². Mihalcea and Csomai distinguished false negatives and false positives in their evaluation. Their approach does not predict instances with a previously unseen surface form (during training phase). Thus, they removed all of these instances from the set of false positives. Consequently, we can compare recall only. Mihalcea and Csomai evaluated on a set of 85 Wikipedia articles drawn at random, which contained 7286 instances; we evaluated on ten times 10,000 random instances (see above). The results for our approach shown in Table 2 are encouraging.

Table 2: Results of the WSD evaluation.

Wikipedia avg. recall	Speech Corpus		
	precision	recall	F_1
.799	.894	.876	.885

We achieve a recall/precision of 0.799. Despite the adaptations of the original approach (see Subsection 3.1) and a test set differing largely in content and extent, this is comparable to the recall of 0.831 Mihalcea and Csomai reported.

In a second evaluation we used the speech corpus. Here, we prepared a gold standard for each noun. The manual transcriptions of the 168 recordings contained 1060 nouns in total. Note that in this evaluation we do not know instances in advance. To obtain the instances we have to identify nouns (except named entities) with a POS tagger. Again, false positives and false negatives encompass all incorrect labels. Additionally, all missed instances are false negatives. The results (shown in Table 2) are promising. We expected a drop in classification quality, as the task is more complex (additionally determine instances) and the domain is different from the training set. Instead, our approach achieves a recall of 0.876 and a precision of 0.894 (F_1 0.885). 21 instances were not disambiguated due to incorrect POS tags produced by our POS tagger. Some incorrectly classified senses are due to nouns that have no corresponding article on Wikipedia, e.g. the word “front”. There is no Wikipedia article describing the concept of *the side that is forward or prominent*. In such cases our approach retrieves incorrect senses.

Nevertheless, our results are encouraging. They show that the approach is feasible, even for domains where the content differs largely from Wikipedia articles (ungrammatical sequences mostly uttered in imperative mood vs. descriptive texts). Thus, the approach proves highly advisable in all contexts, where training

² It is more common to use *accuracy* in this case, which is calculated equally. We kept the notions *precision* and *recall* for comparability with Mihalcea and Csomai

of a custom WSD classifier is impossible because of data sparseness (as in our case) or too expensive.

4.2 Topic Labeling

Evaluating the quality of a topic labeling approach is demanding; one cannot easily provide a gold standard. Usually it is unclear what the correct label is and if it is the only one fitting. Therefore, we performed a user study; it is similar to the study conducted by Hulpus et al. [15]. Six subjects participated in this study; all were graduate students from different faculties, four male and two female, aged 22 to 27. We drew 16 recordings from the speech corpus at random and provided manual transcripts for each. Additionally, we created six synthetic utterance transcriptions. They are comparable to the corpus recordings in regard to linguistic complexity and length. However, these transcripts are from other domains: drone control, child’s playroom, and virtual assistants. Thus, we can evaluate our approach on a broader range of domains. We used our WSD classifier trained on the entire Wikipedia dump to label each noun. Based on the sense labels we created topic labels for each utterance with the *max strategy* and the *top strategy*. We presented the utterances to the subjects together with the topic labels. The labels are ordered according to their confidence values. The total number of topics per utterance varies from four to ten. The subjects were asked to rate each label either as *good fit*, *related but too broad*, *related but inconvenient* or *unrelated*. We divided the subjects into two groups that assessed the labels of eleven utterances each. Thus, all labels were assessed by three annotators. We use Fleiss’ Kappa (κ) to measure the inter-annotator agreement; the determined κ value is 0.27. According to Landis and Koch, this indicates a fair agreement [19]. Hulpus et al. reported a similar κ -value. This outcome illustrates that, although topic labels are quite subjective, shared preferences between annotators are present.

Table 3: Distribution of the assessed quality of the top-k ranked topics produced with *max* and *top strategy*.

k	good fit		too broad		inconvenient		unrelated	
	max	top	max	top	max	top	max	top
1	.530	.530	.242	.242	.045	.045	.182	.182
2	.447	.424	.167	.182	.106	.106	.280	.288
3	.449	.444	.141	.152	.157	.146	.253	.258
4	.432	.420	.144	.140	.155	.167	.269	.273
5	.381	.369	.145	.136	.176	.182	.298	.312
all	.368	.340	.138	.132	.179	.200	.315	.329

The assessment results are depicted in Table 3. It shows the distribution for the top-k topic labels and for all labels (the best results per rank and category are printed in bold). The *good fit*-labels can be interpreted as accurate. Thus,

for the *max strategy* the overall accuracy is 0.368 for all labels and 0.530 for the top-ranked. However, *related but too broad*-labels are also meaningful in most cases, depending on the application at hand (see Section 5). If we consider a combined accuracy of all *good fit*- and *too broad*-labels the value is 0.506; the combined accuracy of the labels at rank one is 0.772. The numbers for the category *related but inconvenient* are less informative. Apparently, a negligible share exists (0.045 of the top ranked labels). However, this category is particularly subjective. Therefore, one has to examine individual cases instead. The table shows the distribution of the annotators assessments. Thus, lower *unrelated*-values are better. Overall, the performance of the *max strategy* is slightly better than the *top strategy* on our test set. This result demonstrates the capability of the *max strategy* to discover small topic areas and label them correctly.

The overall distributions do not take the majority decision of the annotators into account. Therefore, we introduce two additional measures. First, *Precision@k* ($P@k$) as proposed by Hulpus et al.:

$$P@k = \frac{\#Hits \text{ with rank } \leq k}{k} \quad (4)$$

It determines how many labels of the first k labels ($k = [1,5]$) are a *Hit*. A *Hit* is a topic that was assessed a *good fit* (or *good fit or broader* respectively) by at least two of the three annotators. The second is an adaptation of *Coverage@k* ($C@k$) used by Hulpus et al. They measured the fraction of topics that have at least one fitting label. Hulpus et al. model topics explicitly. Hence, they were able to determine $C@k$ on a per-topic-level. Instead, we model topics implicitly, i.e. we determine labels for an entire utterance at once (see Subsection 3.2). Thus, we adapted $C@k$ to fit our set up:

$$C@k = \frac{\#\text{utterances w/ at least 1 Hit at rank } \leq k}{\#\text{utterances}} \quad (5)$$

Our adaptation determines the fraction of utterances for which our approach produces at least one *Hit* in the top- k . Since we want to determine the precision and coverage of accurate or reasonably accurate labels, we consider *good fit* and *good fit or broader* only. Plots for both are presented in Figure 2. Again, the *max strategy* outperforms the *top strategy* almost always. As expected, increasing values for k decreases precision but increases coverage. $C@k$ ranges from 0.801 up to 0.954 for the *good fit or broader*-case, which is encouraging. Even for the *good fit*-only case $C@k$ exceeds 0.909 at $k \geq 4$. Of course, higher precision values would be preferable. Nevertheless, our results are comparable to Hulpus et al. However, one must consider that although precision is calculated equally and values are similar, results have to be interpreted differently. As discussed before, our approach aims at labeling short, spoken utterances. In contrast to Hulpus et al. our approach saves an explicit topic modeling step and labels multiple topics at once. Therefore, precision is semantically slightly different and comparison needs to be considered with caution. A comparison to approaches for topic labeling on spoken language can not be drawn. As discussed in Section 2 these approaches model topics but do not attach labels.

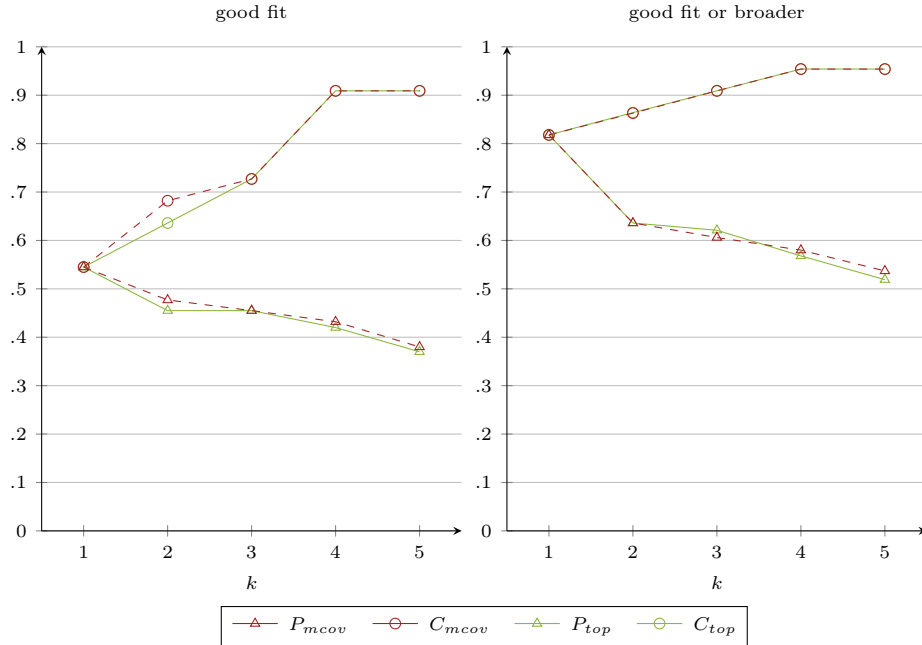



Fig. 2: *Precision@k* (P) and *Coverage@k* (C) for topics that are considered a *good fit* and “*or broader*,” respectively, as achieved with the *max strategy* (*mcov*) and the *top strategy* (*top*).

During evaluation, we discovered another interesting aspect: our results improve with the number of senses available. In other words, our approach has a bias towards long utterances with a broad vocabulary. This behavior is due to the graph centrality approach. It only determines meaningful labels if many connections between senses exist. In general, this is more likely the more senses contribute to the topic graph. Additionally, the homogeneity of senses has a direct influence on the performance of our two strategies. Since the *max strategy* considers all available senses, it performs better on homogeneous inputs, but is more easily diverted by discrete (irrelevant) senses. In such cases, the *top strategy* is more resilient but sometimes discards relevant senses too easily.

5 Areas of Application

In Section 1 we argue that unsupervised multi-topic labeling is a potential building block for a deeper understanding of spoken language. Subsequently, we will justify this point by discussing areas of application for our approach. With topic information at hand, a conversational interface is able to access knowledge more precisely; queries may be restricted to certain parts of a knowledge base that cover the topics mentioned in an utterance. This refines the understanding of the domain, the situational context, and even the common sense capabilities of a system.

Fig. 3: The architecture of *PARSE*.

Our work on unsupervised multi-topic labeling for spoken utterances is part of the project *PARSE* [39]. The goal of the project is to enable laypersons to program in plain, spoken English. Typical application areas of *PARSE* are robotics, home automation, and the like. To facilitate programming with spoken language the system must understand the user’s intents. Capturing the topics of an utterance or conversation is a crucial step towards universal spoken language understanding (SLU) [34]. *PARSE* is equipped with agents for deep SLU. All agents work in parallel and therefore may benefit from results of other agents. The strict separation of concerns between the agents enables us to either build them as knowledge-based or probabilistic agents depending on the task at hand and evaluate them intrinsically.

The architecture of *PARSE*, which is illustrated in Figure 3, is separated in three independent parts: a pipeline for pre-processing, an agent-based main execution, and a pipeline for post-processing. The agents operate on a graph-based shared data structure. Common natural language processing (NLP) tasks, e.g. automatic speech recognition (ASR), part-of-speech (POS) tagging, or shallow parsing, are carried out in the pre-processing pipeline. The user’s utterance is processed sequentially here. In the last pre-processing step, the initial graph is built and passed to the SLU module. There, agents work in parallel and transform the graph to publish their results; they create a semantic representation of the input incrementally. SLU tasks include detection of actions and control structures [35, 37], analysis of coreference and context [38], or – as proposed here – topic labeling. If the graph cannot be transformed into a proper intent model, the utterance is likely to be incomplete or ambiguous. In such situations the user is queried for clarification [36]. During the post-processing pipeline the user’s intents – modeled in the graph – are mapped to the functions of the target system. We first create an abstract syntax tree (AST) from *PARSE*’s graph, which is subsequently transformed to source code. *PARSE* is capable to synthesize code for most common programming languages, including Java, Python, and C.

While most of the process is independent of the domain, the target systems and environments are modeled in ontologies. However, if the target system offers a (high-level) API, the modeling process may be automated. In our previous project *NLCI* [18] we have shown that system ontologies can be extracted (semi-)automatically from most APIs with little effort. We model systems and environmental information in small, independent ontologies. Thus, ambiguities, either in terms of system target functionality or environments, are almost ruled-

Table 4: Domain ontology structure for systems.

class	description
Thing	Top concept of the ontology
└ System	Systems and sub-systems, i.e. API classes
└ Method	System functions, i.e. API methods
└ Parameter	Parameter names used by the system
└ DataType	Data types used by the system, e.g., <i>int</i> or <i>Graspable</i>
└ Object	External objects <empty here>
└ State	States of the external objects <empty here>

Table 5: Domain ontology structure for environments.

class	description
Thing	Top concept of the ontology
└ Object	Objects in environment
└ Graspable	Graspable objects, e.g., <i>cup</i> or <i>fork</i>
└ Openable	Openable objects, e.g., <i>fridge</i> or <i>cupboard</i>
└ Closeable	Closeable objects, e.g., <i>fridge</i> or <i>cupboard</i>
...	
└ State	States of the objects, e.g., <i>opened</i> and <i>closed</i>

out. On one hand, this enhances the understandability and reusability of the ontologies [10, 33]. On the other hand, it enhances the accuracy of *PARSE*'s language analyses that use the ontology. Table 4 and Table 5 show the base structure of the system and environment ontologies. For the time being, *PARSE* is manually configured with the appropriate system and environment ontologies for the use-case, e.g. a robot in a kitchen setting. We merge the system ontology with at least one environment ontology to fill in the objects and states available in the current environment.

In the upcoming subsections we discuss improvements to *PARSE*'s language understanding abilities with the help of topic labeling.

5.1 Ontology selection

The first idea tackles the issue that albeit *PARSE* is almost domain-agnostic it still must be configured with the fitting domain ontologies to work properly. With the help of topic labeling we can determine required ontologies at runtime.

The challenge of selecting ontologies for a given task is faced in research areas like ontology-based question answering or semantically enriched web browsing [28, 29]. A common approach for ontology selection is using lexical information. Park et al. [27] propose an approach that employs semantic similarity matching. The approach selects the ontology that represents the concepts and relations of the intended application the best. They include requirements and contexts for the application domain. ONSET [16] by Khan and Keet faces another ontology selection challenge: the selection of foundational ontologies as base for further ontologies. However, there is (to the best of our knowledge) no approach that employs a comparison of topics for ontology selection.

In our context, the selected ontologies either need to be mapped onto each other or merged. Related approaches to ontology merging employ divide-and-conquer [14, 30] or pairwise linguistic similarity-based methods [9, 21]. Most of them are semi-automatically, since merging ontologies correctly is a challenging task on unrestricted ontologies. Since *PARSE*'s ontology structure is designed to ease merging, we have no need for sophisticated merging strategies. We simply inject the concepts *Object* and *State* of the selected environmental into the system ontology in which these concepts are left empty. Thus, the primary challenge is to select the appropriate ontologies.

Approach Our approach to ontology selection for *PARSE* uses the topics from the spoken utterances to select the best fitting ontologies. First, we attach topics to all ontologies. Second, we compare the extracted topics from the spoken utterance with the extracted ontology topics. We select the ontologies that have the most similar topics with the user's utterance.

To select topics for the ontologies we reuse our approach for unsupervised multi-topic labeling (see Subsection 3.2). We add Wikipedia articles (i.e. word senses) to all the individuals in the ontology. Since there might be individuals that cannot be represented with Wikipedia articles, the annotation process is not applicable for all ontology individuals³. The annotations serve as disambiguation for the individuals in the ontology similar to the word sense disambiguation in the spoken utterances. While this step might be automated in the future, we currently annotate the senses manually. With the annotated senses at hand, we can apply our unsupervised multi-topic labeling approach to the ontologies just as we did to spoken utterances. Since the ontologies represent closed areas of topics, we use the *top strategy* instead of *max strategy* to determine the topics of an ontology.

Then, we calculate the similarities of the topics extracted from the spoken utterance and the topics of the ontologies. Therefore, we again create topic graphs that cover the topics of both sources. We use the distance of topics in the utterance to all topics of an ontology to determine the similarity. The similarity $Sim_T(t, T_o)$ of an utterance topic t and the set T_o of topics of an ontology is calculated as follows:

$$Sim_T(t, T_o) = \frac{1}{\min_{t_o \in T_o} (dist(t, t_o)) + 1}, \quad t \in T_u \quad (6)$$

The distance $dist(t, t_o)$ is defined as the shortest path between topic t and ontology topic t_o in the ontology topic graph. If the topic t is contained in the ontology topic graph, we calculate the minimal distance of t to any of the topics selected for the ontology. The minimal distance is 0, if topic t is one of the selected topics. Therefore, the similarity is 1 if the topic extracted from the utterance is one of the topics of the ontology. Otherwise, the similarity value ranges from 0 to 1. If the utterance topic is not contained in the topic graph, we set the similarity to

³ For instance, there is no Wikipedia article for the term *front* in the sense of *the side that is forward or prominent*.

0. We accumulate the similarities of all topics T_u of an utterance to one ontology to calculate the similarity of each utterance to this ontology (see Equation 7).

$$Sim_{Acc}(T_u, T_o) = \frac{\sum_{t \in T_u} Sim_T(t, T_o)}{|T_u|} \quad (7)$$

After the similarity is calculated for all ontologies, we can select appropriate ontologies. Multiple options on how to select fitting ontologies are on the cards. One option is to select the n most similar ontologies. However, defining a fixed number entails the risk of selecting too few (or too many) ontologies. Another option is to define a (fixed) similarity threshold. However, such a threshold is inflexible and may cause different issues. If the threshold is set too high, there is the possibility that no ontology is selected. If the threshold is set too low, then a lot of inappropriate ontologies are selected. The latter would counteract *PARSE*'s design principle of utilizing small, precise, and independent, ontologies.

We decided on a third, more flexible, option that is basically a combination of the previous two. We derive a flexible threshold based on the most similar ontology and then select ontologies with an accumulated similarity value above that threshold (see Equation 8).

$$Select(T_u, O) = \{o \in O | Sim_{Acc}(T_u, T_o) > \theta * \arg \max_{i \in O} Sim_{Acc}(T_u, T_i)\} \quad (8)$$

The threshold is calculated by reducing the accumulated similarity value of the best ontology based on a fixed threshold factor θ . For example, for a threshold factor of 0.9 and a maximum accumulated similarity value of 0.7, we set the threshold to $0.9 * 0.7 = 0.63$.

However, there is a problem with this approach when selecting system ontologies. We observed that in utterances the system seldomly is a major factor; i.e. system related terms are mentioned in rare cases. Therefore, system concepts are hardly ever among the topics for an utterance. Thus, this approach might have a hard time selecting a fitting system ontology. To cope with that, we adapt the approach slightly for the selection of system ontologies. After the environment ontologies are selected, we compute how each system can cope with the environment.

$$Sim_{Sys}(o_s, T_u, O_{sel}) = \frac{|DT_{o_s} \cap \bigcup_{i \in O_{sel}} Objects_i|}{3 * |DT_{o_s}|} + \frac{2 * Sim_{Acc}(T_u, T_{o_s})}{3} \quad (9)$$

For this, we calculate the so-called compatibility. We define compatibility as the share of data types of a system that are also present as objects in the selected environment ontologies. The corresponding similarity Sim_{Sys} is composed of two-thirds the similarity value Sim_{Acc} (see Equation 7) and one third the just explained compatibility value. Equation 9 shows how Sim_{Sys} is calculated. To compute the compatibility, we calculate the intersection between all data types of the system ontology (DT_{o_s}) and the union of all objects of the selected environment ontologies ($\bigcup_{i \in O_{sel}} Objects_i$). Then, we normalize the value in terms of dividing it by the total amount of data types of the system ontology.

We therefore also apply the ontology selection twice; first to select environment ontologies and afterwards to select a system ontology. This ensures that there is at least one selected ontology for each, system and environment. Finally, we merge all selected ontologies using the simple merging strategy discussed above that exploits the fixed structure of *PARSE*'s ontologies.

Table 6: Domain ontologies used in the evaluation. *Type* shows whether the ontology is a system ontology (S) or an environment ontology (E).

type	name	description
S	Household Robot	Household robots, like ARMAR ⁴
S	Virtual Assistant	Virtual assistants, like Amazons Echo (Alexa)
S	Drone	Drones, like a quadcopter
S	Lego Mindstorm	Lego Mindstorm Roboter
E	Kitchen	Kitchen with, e.g., utensils, furniture, groceries
E	Bar	Bar, including cocktails, furniture etc.
E	Garden	Garden with garden furniture, plants, tools etc.
E	Bedroom	Bedroom with furniture etc.
E	Playroom	Children's room with furniture, toys, etc.
E	Music	Music concepts and genres, instruments and alike
E	Heating	Heating domain (heaters, air conditioning, ...)
E	Laundry	Things like washing machine, dryer, detergent etc.

Evaluation We used the 24 utterances from Subsection 4.2 to evaluate our ontology selection approach. Additionally, we added nine synthetic utterances to enrich the evaluation corpus with two further domains and utterances that need multiple environment ontologies. The evaluation corpus thus contains 33 utterances with eight ontologies modeling environment and four ontologies modeling systems. Five utterances include objects from more than one environment ontology. The assembled domains (system ontology plus environment ontologies) for the three scenarios are heating and music, music and bar, and kitchen and garden. We use the ontologies listed in Table 6 to model the systems and environments.

We prepared a gold standard for each utterance containing the required ontologies for execution. After the application of the approach for ontology selection on the utterances, we compare the selected ontologies with the gold standard. We first evaluate the selection of environment ontologies. An ontology that was correctly selected counts as a *true positive*, an incorrectly selected one is a *false positive*. Similarly, an ontology that was correctly not selected is a *true negative*, while an ontology that was incorrectly not selected counts as a *false negative*. We determine precision, recall, and F_1 -score. Additionally, we use the false positive rate (fpr) to capture the chance of selecting one of the inappropriate ontologies.

Table 7 shows the results of our evaluation corpus for different configurations. The first two columns display different options for the amount of topics labeled

Table 7: Different configurations for topic labeling (TL) and ontology selection with the corresponding results for precision, recall, F_1 -score, and false positive rate (fpr).

TL		ontology selection		precision	recall	F_1	fpr
#topics	#topics	θ					
2*n	5	0.90		0.919	0.895	0.907	0.013
2*n	5	0.85		0.809	0.895	0.850	0.035
2*n	5	0.80		0.708	0.895	0.791	0.062
5	5	0.90		0.850	0.895	0.872	0.027
5	5	0.85		0.791	0.895	0.840	0.040
5	5	0.80		0.761	0.921	0.833	0.049
2*n	10	0.90		0.778	0.921	0.843	0.044
5	10	0.90		0.790	0.895	0.840	0.040

to the spoken utterances and the ontologies, respectively. The third column varies the threshold factor θ . All configurations yield good results with F_1 -scores between 0.79 and 0.91. The false positive rate is relatively low and in line with the other metrics. The best results are obtained by a configuration with 2*n topics for the topic extraction, five topics for each ontology and a threshold factor of 0.90. While recall is stable throughout the configurations, precision decreases for lower thresholds.

Table 8: Results of the utterances that required multiple environment ontologies.

utterance	gold standard	selected	recall
21	Heating \cup Music	Heating	0.500
22	Heating \cup Music	Heating	0.500
31	Bar \cup Music	Bar	0.500
32	Bar \cup Music	Bar \cup Music	1.000
33	Garden \cup Kitchen	Garden	0.500

A thorough analysis of the false positives revealed that in many cases the (wrongly) selected ontology actually had a lot in common with the intended ontologies of the gold standard; i.e. they covered overlapping concepts and individuals. For example, the ontologies about *bar* and *kitchen* shared individuals like *orange juice* or *fridge*.

The utterances requiring multiple environment ontologies show a lower recall (see Table 8). This is due to overrepresented topics that can be found in one ontology only. This influences the accumulated similarity and therefore the ontology selection. For these utterances on average 61.90% of the entities that are present in the utterance can be mapped to the corresponding individuals in the ontology. In comparison, the average for all scenarios is at 95.79%.

Regarding the selection of system ontologies, we must admit that the selection of a system does not work well. Without the adaptations (that uses the data types), our approach is unable to correctly select a system ontology in almost all cases. The approach can not distinguish between the different systems based on the topic labeling alone. The adaptation improves the selection, but new problems arise. In 94% of the cases, the system ontology for the household robot is selected. The household robot can operate in almost all environments; thus, its system ontology dominates the others. Therefore, the approach has to be improved further to assure a proper selection of system ontologies. Alternatively, instead of selecting the system ontology at runtime, the system ontology can be selected manually beforehand.

Planned Improvements Since the results obtained so far are promising, we plan to continue the development of this application. The major drawback of the current implementation is the manual word sense labeling step for the ontology individuals. Thus, we plan to automate this step as soon as possible. There are at least three possible approaches that we want to explore. The first possible approach incorporates the relations of an individual in the ontology. If there is more than one possible sense for an individual, the description of the senses – in the case of Wikipedia articles the article content – can be used to narrow down the possibilities until only one candidate is left. The inclusion of the description can be done for example by counting the number (or share) of terms within the descriptions that correspond to related individuals in the ontology. A second possible approach is even more about the structure of the ontology. There are usually some individuals that are unambiguous. These can be used as *root concepts* for the approach. We use the relations of the unambiguous and disambiguated individuals in resources like DBpedia to compare them to related individuals in our ontology. In this way, we can steadily disambiguate individuals. A combination of the two approaches might also pose a valid method. As a third approach, we want to explore the usefulness of natural language statements generated from the information within the ontologies. We utilize the relations defined in the ontologies to produce natural language statements like “A fork is a graspable object” and “A fridge can be opened and closed”. We are then able to reuse our WSD approach (see Subsection 3.1) on these natural language statements to disambiguate the individuals of the ontology.

5.2 Further Areas of Application

Besides the selection of domain ontologies, topic information might prove useful for other language understanding components of *PARSE*. In the following we will describe our plans to improve *PARSE*’s context modeling, knowledge acquisition, and dialog interaction.

Context modeling In *PARSE* we build a comprehensive context model. Among other information the model includes concept relations between entities [38]. The precision of the conceptualization might improve if we incorporate information about the current topics; e.g. we might be able to distinguish the concepts *cup(dishware)* and *cup(trophy)* in more contexts.

CyC (micro theories) Another idea concerns world knowledge bases such as CyC [20]. We have experimented largely with CyC to enrich different language inputs with world knowledge, e.g. to prove the feasibility of a described course of action. However, the knowledge CyC stores is vast and hard to handle if one does not use precise queries. Luckily, all information is stored in so-called *micro theories* that cover knowledge about a certain topic. Thus, if we match topics extracted from spoken utterances with micro theories, we might be able to reduce the search space and improve querying.

Dialog interaction *PARSE* also employs an extensible dialog component to resolve ambiguous situations [36]. With topic information at hand, we are able to pose more precise queries; e.g. if the system has understood that the topic is *kitchen* but missed some parts of the utterance, we might ask the user, “Do you mean ‘go to the fridge’?”, instead of replying that the system has not understood the last word.

6 Conclusion and Future Work

We have presented an approach for unsupervised multi-topic labeling that is tailored to spoken language. State-of-art approaches either depend on large textual corpora or model topics but do not attach labels. We see topic labeling as a fundamental building block to gain a deeper understanding of spoken utterances.

The contribution we presented here is three-fold. First, we have adapted the approach for word sense disambiguation by Mihalcea and Csomai [25] to short, spoken utterances. We can confirm their results; on the Wikipedia data set we achieve a similar recall (recall 0.799 vs. 0.831 in the original paper). An additional evaluation on a speech corpus with instructions for a robot shows that the method works properly on previously unseen input (F_1 : 0.887).

The second contribution is the approach for unsupervised multi-topic labeling for spoken utterances. Based on the word senses – i.e. the attached Wikipedia articles – we construct so-called topic graphs from DBpedia relations. We use graph centrality to determine the topics. We implemented two strategies to find the most central terms, called *top strategy* and *max strategy*. The first strategy creates topics that describe the dominant part of the utterance. The latter covers as many distinct senses as possible. Our approach creates multiple topics per utterance. Our evaluation shows that the *max strategy* slightly outperforms the *top strategy* in almost all cases.

However, the *max strategy* is more sensitive to single unrelated parts of the utterance. The overall results are promising. In a user study subjects assessed 53% of the top-ranked topic labels as *good fit*. Furthermore, for 90.9% of the evaluated utterances at least one of the top four topic labels was considered a *good fit*. If we also take labels into account that were assessed as *related but too broad* the results are even more encouraging. Subjects judged that 77.2% of the top ranked labels fit into this category; 95.4% of the utterances receive at least one topic label of this quality.

Fortunately, in most potential application areas *good or too broad*-labels are equally useful. As a third contribution, we presented an approach to automatically select the most appropriate ontologies that model the situational context

in our research project *PARSE*. Therefore, we attach topics to both, utterances and ontologies, simultaneously. Then, we determine the ontologies that share the most topics with the utterance. Consequently, *too broad* labels are valuable, as long as ontology labels are similar or related to the utterance labels. The approach obtained an F_1 -score of 90.7% in an evaluation with eight different domain ontologies.

Too broad labels are valuable as well for the selection of CyC micro theories. Refining the conceptualization of our context model also works with broader topics.

Beyond that, we will utilize the topics for more precise dialog management and explore other application areas. Furthermore, we plan to implement and evaluate additional strategies to determine central terms and experiment with differently weighted edges in sense graphs.

References

1. Aker, A., Kurtic, E., Balamurali, A.R., Paramita, M., Barker, E., Hepple, M., Gaizauskas, R.: A Graph-Based Approach to Topic Clustering for Online Comments to News. In: *Advances in Information Retrieval*. pp. 15–29. *Lecture Notes in Computer Science*, Springer International Publishing (2016)
2. Allahyari, M., Kochut, K.: Automatic Topic Labeling Using Ontology-Based Topic Models. In: *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*. pp. 259–264 (Dec 2015). <https://doi.org/10.1109/ICMLA.2015.88>
3. Blei, D.M.: Probabilistic topic models. *Commun. ACM* **55**(4), 77–84 (Apr 2012). <https://doi.org/10.1145/2133806.2133826>
4. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent Dirichlet Allocation. *Journal of Machine Learning Research* **3**, 993–1022 (Mar 2003)
5. Boyd-Graber, J., Hu, Y., Mimno, D.: Applications of topic models. *Foundations and Trends® in Information Retrieval* **11**(2-3), 143–296 (2017). <https://doi.org/10.1561/15000000030>
6. Cerisara, C.: Automatic discovery of topics and acoustic morphemes from speech. *Computer Speech & Language* **23**(2), 220–239 (2009)
7. Coursey, K., Mihalcea, R.: Topic Identification Using Wikipedia Graph Centrality. In: *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*. pp. 117–120. *NAACL-Short '09*, Association for Computational Linguistics, Stroudsburg, PA, USA (2009)
8. Coursey, K., Mihalcea, R., Moen, W.: Using Encyclopedic Knowledge for Automatic Topic Identification. In: *Proceedings of the Thirteenth Conference on Computational Natural Language Learning*. pp. 210–218. *CoNLL '09*, Association for Computational Linguistics, Stroudsburg, PA, USA (2009)
9. Cruz, I.F., Antonelli, F.P., Stroe, C.: Agreementmaker: efficient matching for large real-world schemas and ontologies. *Proceedings of the VLDB Endowment* **2**(2), 1586–1589 (2009)
10. d’Aquin, M., Schlicht, A., Stuckenschmidt, H., Sabou, M.: Ontology modularization for knowledge selection: Experiments and evaluations. In: Wagner, R., Revell, N., Pernul, G. (eds.) *Database and Expert Systems Applications*. pp. 874–883. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)

11. Gourru, A., Velcin, J., Roche, M., Gravier, C., Poncelet, P.: United We Stand: Using Multiple Strategies for Topic Labeling. In: Natural Language Processing and Information Systems. pp. 352–363. Lecture Notes in Computer Science, Springer International Publishing (2018)
12. Hazen, T.J., Siu, M., Gish, H., Lowe, S., Chan, A.: Topic modeling for spoken documents using only phonetic information. In: 2011 IEEE Workshop on Automatic Speech Recognition Understanding. pp. 395–400 (Dec 2011). <https://doi.org/10.1109/ASRU.2011.6163964>
13. Hingmire, S., Chougule, S., Palshikar, G.K., Chakraborti, S.: Document Classification by Topic Labeling. In: Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 877–880. SIGIR '13, ACM, New York, NY, USA (2013). <https://doi.org/10.1145/2484028.2484140>
14. Hu, W., Qu, Y.: Falcon-ao: A practical ontology matching system. *Web Semantics: Science, Services and Agents on the World Wide Web* **6**(3), 237–239 (2008)
15. Hulpus, I., Hayes, C., Karnstedt, M., Greene, D.: Unsupervised graph-based topic labelling using dbpedia. In: Proceedings of the Sixth ACM International Conference on Web Search and Data Mining. pp. 465–474. WSDM '13, ACM, New York, NY, USA (2013). <https://doi.org/10.1145/2433396.2433454>, <http://doi.acm.org/10.1145/2433396.2433454>
16. Khan, Z., Keet, C.M.: Onset: Automated foundational ontology selection and explanation. In: International Conference on Knowledge Engineering and Knowledge Management. pp. 237–251. Springer (2012)
17. Kiesling, S., Dilley, L., Raymond, W.D.: The variation in conversation (ViC) project: Creation of the Buckeye Corpus of Conversational Speech. Ohio State University, Columbus, OH (2006)
18. Landhäußer, M., Weigelt, S., Tichy, W.F.: NLCI: A Natural Language Command Interpreter. *Automated Software Engineering* (2016). <https://doi.org/10.1007/s10515-016-0202-1>
19. Landis, J.R., Koch, G.G.: The measurement of observer agreement for categorical data. *Biometrics* **33**(1), 159–174 (Jan 1977). <https://doi.org/10.2307/2529310>
20. Lenat, D.B.: CYC: A Large-scale Investment in Knowledge Infrastructure. *Commun. ACM* **38**(11), 33–38 (Nov 1995). <https://doi.org/10.1145/219717.219745>
21. Li, J., Tang, J., Li, Y., Luo, Q.: Rimom: A dynamic multistrategy ontology alignment framework. *IEEE Transactions on Knowledge and Data Engineering* **21**(8), 1218–1232 (2009)
22. Magatti, D., Calegari, S., Ciucci, D., Stella, F.: Automatic Labeling of Topics. In: 2009 Ninth International Conference on Intelligent Systems Design and Applications. pp. 1227–1232. ISDA '09, IEEE Computer Society, Washington, DC, USA (Nov 2009). <https://doi.org/10.1109/ISDA.2009.165>
23. Mei, Q., Shen, X., Zhai, C.: Automatic Labeling of Multinomial Topic Models. In: Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 490–499. KDD '07, ACM, New York, NY, USA (2007). <https://doi.org/10.1145/1281192.1281246>
24. Mihalcea, R.: Using Wikipedia for Automatic Word Sense Disambiguation. In: Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference. pp. 196–203. Association for Computational Linguistics, Rochester, New York (Apr 2007)
25. Mihalcea, R., Csomai, A.: Wikify!: Linking Documents to Encyclopedic Knowledge. In: Proceedings of the Sixteenth ACM Conference on Conference on Information

- and Knowledge Management. pp. 233–242. CIKM '07, ACM, New York, NY, USA (2007). <https://doi.org/10.1145/1321440.1321475>
26. Muhr, M., Kern, R., Granitzer, M.: Analysis of structural relationships for hierarchical cluster labeling. In: Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval. pp. 178–185. ACM (2010)
 27. Park, J., Oh, S., Ahn, J.: Ontology selection ranking model for knowledge reuse. *Expert Systems with Applications* **38**(5), 5133–5144 (2011)
 28. Sabou, M., Lopez, V., Motta, E., Uren, V.: Ontology selection: ontology evaluation on the real semantic web. In: 15th International World Wide Web Conference (WWW 2006) (2006), <http://oro.open.ac.uk/23648/>
 29. Sabou, M., Lopez, V., Motta, E.: Ontology selection for the real semantic web: How to cover the queen’s birthday dinner? In: Staab, S., Svátek, V. (eds.) *Managing Knowledge in a World of Networks*. pp. 96–111. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
 30. Seddiqui, M.H., Aono, M.: An efficient and scalable algorithm for segmented alignment of ontologies of arbitrary size. *Web Semantics: Science, Services and Agents on the World Wide Web* **7**(4), 344–356 (2009)
 31. Shi, B., Lam, W., Jameel, S., Schockaert, S., Lai, K.P.: Jointly Learning Word Embeddings and Latent Topics. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 375–384. SIGIR '17, ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3077136.3080806>
 32. Siu, M.h., Gish, H., Chan, A., Belfield, W., Lowe, S.: Unsupervised training of an hmm-based self-organizing unit recognizer with applications to topic classification and keyword discovery. *Computer Speech & Language* **28**(1), 210–223 (2014)
 33. Stuckenschmidt, H., Parent, C., Spaccapietra, S.: *Modular ontologies: concepts, theories and techniques for knowledge modularization*, vol. 5445. Springer (2009)
 34. Tur, G., De Mori, R.: *Spoken Language Understanding: Systems for Extracting Semantic Information from Speech*. John Wiley & Sons Ltd (Mar 2011)
 35. Weigelt, S., Hey, T., Steurer, V.: Detection of Conditionals in Spoken Utterances. In: 2018 IEEE 12th International Conference on Semantic Computing (ICSC). pp. 85–92 (Jan 2018). <https://doi.org/10.1109/ICSC.2018.00021>
 36. Weigelt, S., Hey, T., Landhäußer, M.: Integrating a Dialog Component into a Framework for Spoken Language Understanding. In: Proceedings of the 6th International Workshop on Realizing Artificial Intelligence Synergies in Software Engineering. pp. 1–7. RAISE '18, ACM, Gothenburg, Sweden (2018). <https://doi.org/10.1145/3194104.3194105>
 37. Weigelt, S., Hey, T., Steurer, V.: Detection of Control Structures in Spoken Utterances. *International Journal of Semantic Computing* **12**(03), 335–360 (Sep 2018). <https://doi.org/10.1142/S1793351X18400159>
 38. Weigelt, S., Hey, T., Tichy, W.F.: Context Model Acquisition from Spoken Utterances. In: The 29th International Conference on Software Engineering & Knowledge Engineering,. pp. 201–206. Pittsburgh, PA (Jul 2017). <https://doi.org/10.18293/SEKE2017-083>
 39. Weigelt, S., Tichy, W.F.: Poster: ProNat: An Agent-Based System Design for Programming in Spoken Natural Language. In: 2015 IEEE/ACM 37th IEEE International Conference on Software Engineering (ICSE). vol. 2, pp. 819–820 (May 2015). <https://doi.org/10.1109/ICSE.2015.264>