

Lexical Enrichment of a Human Anatomy Ontology using WordNet

Nils Reiter¹ and Paul Buitelaar²

¹ Department of Computational Linguistics, Heidelberg University,
Heidelberg, Germany**

`reiter@cl.uni-heidelberg.de`

² Language Technology Lab & Competence Center Semantic Web, DFKI,
Saarbrücken, Germany
`paulb@dfki.de`

Abstract. This paper is concerned with lexical enrichment of ontologies, i.e. how to enrich a given ontology with lexical entries derived from a semantic lexicon. We present an approach towards the integration of both types of resources, in particular for the human anatomy domain as represented by the Foundational Model of Anatomy (FMA). The paper describes our approach on combining the FMA with WordNet by use of a simple algorithm for domain-specific word sense disambiguation, which selects the most likely sense for an FMA term by computing statistical significance of synsets on a corpus of Wikipedia pages on human anatomy. The approach is evaluated on a benchmark of 50 ambiguous FMA terms with manually assigned WordNet synsets (i.e. senses).

1 Introduction

This paper is concerned with lexical enrichment of ontologies, i.e. how to enrich a given ontology with lexical entries derived from a semantic lexicon. The assumption here is that an ontology represents domain knowledge with less emphasis on the linguistic realizations (i.e. words) of knowledge objects, whereas a semantic lexicon such as WordNet defines lexical entries (words with their linguistic meaning and possibly morpho-syntactic features) with less emphasis on the domain knowledge associated with these.

1.1 Ontologies

An ontology is an explicit and formal description of the conceptualization of a domain of discourse (see e.g. Gruber Gruber (1993), Guarino Guarino (1998)). In its most basic form an ontology consists of a set of classes and a set of relations that describe the properties of each class. Ontologies formally define relevant knowledge in a domain of discourse and can be used to interpret data in this domain (e.g. medical data such as patient reports), to reason over knowledge that

** This work has been done while the first author was affiliated to the Department of Computational Linguistics at Saarland University.

can be extracted or inferred from this data and to integrate extracted knowledge with other data or knowledge extracted elsewhere. With recent developments towards knowledge-based applications such as intelligent Question Answering, Semantic Web applications and semantic-level multimedia indexing and retrieval, the interest in large-scale ontologies has increased. Here we use a standard ontology in human anatomy, the FMA: Foundational Model of Anatomy³ (Rosse and Mejino Jr (2003)). The FMA describes the domain of human anatomy in much detail by way of class descriptions for anatomical objects and their properties. Additionally, the FMA lists terms in several languages for many classes, which makes it a lexically enriched ontology already. However, our main concern here is to extend this lexical representation further by automatically deriving synonyms from WordNet.

1.2 Lexicons

A lexicon describes the linguistic meaning and morpho-syntactic features of words and possibly also of more complex linguistic units such as idioms, collocations and other fixed phrases. Semantically organized lexicons such as WordNet and FrameNet define word meaning through formalized associations between words, i.e. in the form of synsets in the case of WordNet and with frames in the case of FrameNet. Although such a representation defines some semantic aspects of a word relative to other words, it does not represent any knowledge about the objects that are referred to by these words. For instance, the English noun “ball” may be represented by two synsets in WordNet ({ball, globe}, {ball, dance}) each of which reflects another interpretation of this word. However, deeper knowledge about what a “ball” in the sense of a “dance” involves (a group of people, a room, music to which the group of people move according to a certain pattern, etc.) cannot be represented in this way. Frame-based definitions in FrameNet allow for such a deeper representation into some respect, but also in this case the description of word meaning is concerned with the relation between words and not so much with object classes and their properties that are referred to by these words. However, for instance in the case of BioFrameNet (Dolbey et al. (2006)), an extension of FrameNet for use in the biomedical domain, such an attempt has been made and is therefore much in line with our work described here.

1.3 Related Work

Other related work is on word sense disambiguation (WSD) and specifically domain-specific WSD as this is a central aspect of our algorithm in selecting the most likely sense of words occurring in FMA terms. The work presented here is based directly on Buitelaar and Sacaleanu (2001) and similar approaches (McCarthy et al. (2004a); Koeling and McCarthy (2007)). Related to this work is the

³ See <http://sig.biostr.washington.edu/projects/fm/AboutFM.html> for more details on the FMA

assignment of domain tags to WordNet synsets (Magnini and Cavaglia (2000)), which would obviously help in the automatic assignment of the most likely sense in a given domain – as shown in Magnini et al. (2001). An alternative to this idea is to simply extract that part of WordNet that is directly relevant to the domain of discourse (Cucchiarelli and Velardi (1998); Navigli and Velardi (2002)). However, more directly in line with our work on enriching a given ontology with lexical information derived from WordNet is presented in Pazienza and Stellato (2006), but the main difference here is that we use a domain corpus as additional evidence for statistical significance of a selected word sense (i.e. synset). Finally, also some recent work on the definition of ontology-based lexicon models (Alexa et al. (2002); Gangemi et al. (2003); Buitelaar et al. (2006)) is of (indirect) relevance to the work presented here as the derived lexical information needs to be represented in such a way that it can be easily accessed and used by NLP components as well as ontology management and reasoning tools.

2 Approach

Our approach to lexical enrichment of ontologies consists of a number of steps, each of which will be addressed in the remainder of this section:

1. extract all terms from the term descriptions of all classes in the ontology, lookup of terms in WordNet
2. for ambiguous terms: apply domain-specific WSD by ranking senses (synsets) according to statistical relevance in the domain corpus
3. select most relevant synset and add the synonyms of this synset to the corresponding term representation

2.1 Term Extraction and WordNet Lookup

Ontologies, such as the FMA, describe objects and their relations to each other. Additionally, each such object (or rather the class descriptions for such objects) may carry terminological information in one or more languages. In the FMA, terms for classes are defined in several languages, i.e. 100,000 English terms, 8,000 Latin, 4,000 French, 500 Spanish and 300 German terms. Terms in the FMA can be simple, consisting of just one word, or complex multiword terms, e.g. “muscular branch of lateral branch of dorsal branch of right third posterior intercostal artery”. In our approach we considered simple as well as complex terms although only a small number of such domain-specific terms will actually occur in WordNet as will be reported below in section 3.

2.2 WSD Algorithm

The core of our approach is the word sense disambiguation algorithm as shown in figure 1. The algorithm iterates over every synonym of every synset of the term in question. It calculates the χ^2 values of each synonym and adds them up for each synset.

```

function getWeightForSynset(synset) {
    synonyms = all synonyms of synset
    weight = 0
    foreach synonym in synonyms
        c = chi-square(synonym)
        weight = weight + c
    end foreach
    return weight
}

s = synsets to which t belongs
highest_weight = 0
best_synsets = {}
foreach synset in s
    synonyms = all synonyms of synset

    weight = getWeightForSynset(synset)

    if (weight == highest_weight)
        best_synsets = best_synsets + { synset }
    else if (weight > highest_weight)
        best_synsets = { synset }
    end if
end foreach
return best_synsets

```

Fig. 1. Algorithm for the sense disambiguation of the term *t*

Using the χ^2 -test (see, for instance, (Manning and Schütze, 1999, p. 169)), one can compare the frequencies of terms in different corpora. In our case, we use a reference and a domain corpus and assume that the terms occurring (relatively) more often in the domain corpus than in the reference corpus are “domain terms”, i.e., are specific to this domain. If it is a domain term, it should be defined in the ontology.

$$\chi^2(t) = \frac{N * (O_{11}^t O_{22}^t - O_{12}^t O_{21}^t)^2}{(O_{11}^t + O_{12}^t)(O_{11}^t + O_{21}^t)(O_{12}^t + O_{22}^t)(O_{21}^t O_{22}^t)} \quad (1)$$

χ^2 , calculated according to formula 1, allows us to measure exactly this. O_{11}^t and O_{12}^t denote the frequencies of the term t in the domain and reference corpora while O_{21}^t and O_{22}^t denote the frequency of any term but t in the domain and reference corpora:

- O_{11}^t = frequency of t in the domain corpus
- O_{12}^t = frequency of t in the reference corpus
- O_{21}^t = frequency of $\neg t$ in the domain corpus
- O_{22}^t = frequency of $\neg t$ in the reference corpus
- N = Added size of the two corpora

The algorithm finally choses the synset with the highest weight as the appropriate one.

The term “gum”, for instance, has six noun senses with on average 2 synonyms. The χ^2 value of the synonym “gum” itself is 6.22. Since this synonym occurs obviously in every synset of the term, it makes no difference for the rating. But the synonym “gingiva”, which belongs to the second synset and is the medical term for the gums in the mouth, has a χ^2 value of 20.65. Adding up the domain relevance scores of the synonyms for each synsets, we find that the second synset gets the highest weight and is therefore selected as the appropriate one.

Relations The algorithm as shown in figure 1 uses the synonyms found in WordNet. However, other relations that are provided by WordNet can be used as well. Figure 2 shows the improved algorithm. The main difference is that we calculate and add the weights for each synonym of each synset to which a synset of the original term is related.

2.3 Lexical Representation

Finally, after the synsets for an ambiguous term t have been ranked according to relevance to the domain, we can select the top one or more to be included as (additional) lexical/terminological information in the ontology, i.e., the synonyms that are contained in this synset can be added as (further) terms for the ontology class c that corresponds to term t .

```

r = WordNet relations
s = synsets to which t belongs
highest_weight = 0
best_synsets = {}
foreach synset in s
  weight = getWeightForSynset(synset)

  related = with r related synsets

  foreach rsynset in related
    weight += getWeightForSynset(rsynset)
  end foreach

  if (weight == highest_weight)
    best_synsets = best_synsets + { synset }
  else if (weight > highest_weight)
    best_synsets = { synset }
  end if
end foreach
return best_synsets

```

Fig. 2. Improved algorithm – As in figure 1 but including WordNet relations

Here, we actually propose to extend the ontology with an ontology-based lexicon format, LingInfo, which has been developed for this purpose in the context of previous work (Buitelaar et al. (2006)). By use of the LingInfo model we will be able to represent each synonym for t as a linguistic object l that is connected to the corresponding class c . The object l is an instance of the LingInfo class of such linguistic objects that cover the representation of the orthographic form of terms as well as relevant morpho-syntactic information, e.g. stemming, head-modifier decomposition, part-of-speech. The implementation of a LingInfo-based linguistic knowledge base for the FMA is ongoing work, but a first version of a similar knowledge base for the football domain has been developed in the context of the SmartWeb project (Buitelaar et al. (2006); Oberle et al. (forthcoming)).

3 Experiment

In an empirical experiment, we enrich the FMA (“Foundational Model of Anatomy”) ontology with lexical information (synonyms) derived from WordNet using Wikipedia pages on human anatomy as domain corpus.

3.1 Data Sources

Ontology: Foundational Model of Anatomy “The Foundational Model of Anatomy (FMA) ontology was developed by the Structural Informatics Group⁴

⁴ <http://sig.biostr.washington.edu/index.html>

at the University of Washington. It contains approximately 75,000 classes and over 120,000 terms; over 2.1 million relationship instances from 168 relationship types link the FMA's classes into a coherent symbolic model. The FMA is one of the largest computer-based knowledge sources in the biomedical sciences. The most comprehensive component of the FMA is the Anatomy taxonomy" (FMA website), organized around the top class Anatomical Structure. "Anatomical structures include all material objects generated by the coordinated expression of groups of the organism's own structural genes. Thus, they include biological macromolecules, cells and their parts, tissues, organs and their parts, as well as organ systems and body parts (body regions)" (FMA website). For the purpose of the experiment reported on here we used the taxonomy component of the FMA, extracted all English terms and did a lookup for each of these terms in WordNet.

Semantic Lexicon: WordNet The most recent version of WordNet (3.0) was used in our experiment. As an interface to our own implementation, we use the Java WordNet interface⁵. The number of English terms (simple and complex) we were able to extract from the FMA was 120,417, of which 118,785 were not in WordNet. This left us with a set of 1,382 terms that were in WordNet but only 250 of these were actually ambiguous and therefore of interest to our experiment. Interestingly, 10 of these were in fact multiword terms. The experiment as reported below is therefore concerned with the disambiguation of these 250 FMA terms, given their sense assignments in WordNet.

Medical Corpus: Wikipedia Pages on Human Anatomy Our approach requires the use of a domain corpus. As corpus from the anatomy domain, we use the Wikipedia pages from the category "Human Anatomy" and all its subordinated categories⁶. These are 7,251 single pages, containing over 4.4 million words.

We removed the meta information (categories, tables of contents, weblinks, ...) using heuristic methods. By use of part of speech tagging with the Tree-Tagger (Schmid (1994)), we automatically extracted all nouns from this corpus, resulting in 1.3 million noun tokens and 92,927 noun types.

Reference Corpus: British National Corpus Our ranking of the domain relevance of a synset is based on comparing the frequencies of its synonyms in a domain corpus and a reference corpus. The reference corpus we use is the British National Corpus (BNC). Since we were only interested in the frequencies, we used the frequency lists provided by Leech et al. (2001).

⁵ <http://www.mit.edu/~markaf/projects/wordnet/>

⁶ http://en.wikipedia.org/wiki/Category:Human_anatomy

3.2 Benchmark

Our benchmark (gold standard) consists of randomly selected 50 ambiguous terms from the ontology. Four terms have been removed from the test set because none of their senses belong to the domain of human anatomy.

Two annotators manually disambiguated them according to the domain of human anatomy. Each term is associated with one (or more) WordNet synsets. More than one synset is used due to the high granularity of WordNet. The agreement between the two annotators is, generally speaking, high. In every single case, there is an overlap in the associated synsets, i.e., for every term, there is at least one synset chosen by both annotators. If we count only a perfect match, i.e., both annotators chose exactly the same set of senses, the kappa value κ according to Cohen (1960) is still $\kappa = 0.71$.

Baseline The synsets in WordNet are sorted according to frequency. The word “jaw”, for instance, occurs more often with its first synset than with its third. It is therefore a reasonable assumption for any kind of word sense disambiguation to pick always the first sense (see, for instance, McCarthy et al. (2004b)). We use this simple approach as baseline for our evaluation.

3.3 Evaluation

The system was evaluated with respect to precision, recall and f-score. Precision is the proportion of the meanings predicted by the system which are correct. Recall is the proportion of correct meanings which are predicted by the system. Finally, the f-score is the harmonic mean of precision and recall, and is the final measure to compare the performance of systems.

We calculated precision, recall and f-score separately for WordNet relations and test items. The test item “jaw”, for instance, was manually disambiguated to the first or second noun synset of the word. Our program, using the hyponymy relation, returned only the first noun synset. For this item, we calculate a precision of 100% and a recall of 50%.

Table 1 shows the results averaged over all 46 test items for the different relations. The relations not shown did not give other results than the algorithm without using any of the relations (results for “all other” are exactly the same as “only synonyms”).

The two lines at the bottom of the table are combinations of relations. In the first line, we use the three relations with the highest precision (and f-score, but that is a coincidence) together (hyponym, holonym (part) and meronym (part)). The last line shows the three relations with the highest recall taken together (topic, holonym (substance) and meronym (part)). Note, that the meronym (part) relation is the only relation that is among the top three in both cases.


3.4 Discussion

Our results show – in almost any configuration – a clear improvement compared to the baseline.

Table 1. Evaluation Results for the different WordNet relations

Relation	Precision	Recall	F-Measure
Baseline (first sense)	58.69	46.56	51.93
Only Synonyms	54.78	65.58	59.70
Hypernym	56.52	47.10	51.38
Hypernym (instance)	53.70	63.41	63.22
Hyponym	64.93	63.95	64.44
Topic	56.96	67.75	61.89
Holonym (part)	63.04	63.41	63.22
Holonym (substance)	55.87	65.58	60.34
Meronym (member)	52.61	63.41	57.51
Meronym (part)	58.05	68.12	62.68
Meronym (substance)	55.51	62.32	58.72
All other	54.78	65.58	59.70
Hyponym, Holonym (part), Meronym (part)	77.53	70.11	73.63
Topic, Holonym (substance), Meronym (part)	61.30	70.29	65.49

Using just the synonyms of WordNet and no additional relation(s), we observe an increase in recall (around 20%) and a relatively small decrease in precision (less than 5%). The increase in recall can easily be explained by the fact that our baseline takes only the first (and therefore: only one) synset – every term that is disambiguated to more than one synset already gets a recall of 50% or less. The decrease in precision can be explained by looking at the test samples. For some of the synsets, a synonym – especially when it comes to multi-word expressions – can not be found in the corpus. This leads to the same weight for a number of synsets and thus to more selected synsets, even if the evidence does not increase. The precision decreases because among the selected synsets, there are more inappropriate ones. Or, the other way around: if an appropriate synset has no synonyms (or only synonyms that do not appear in the corpus), the precision decreases.

**Fig. 3.** Synonyms for the synsets of “alveolus”

For the term “alveolus”, for instance, both noun synsets are annotated as appropriate in the gold standard. The baseline algorithm selects only the first synset and gets therefore a precision of 100% and a recall of 50%. Figure 3 shows

the synonyms for the term “alveolus” graphically. In the configuration, where we just use the WordNet synonyms, both synsets get the same weight, because the synonym alveolus appears 137 times in the domain corpus, and all other synonyms do not appear at all (not a single occurrence of “air sac”, “air cell” and “tooth socket”).

This problem diminishes, if WordNet relations are taken into account. By using WordNet relations, we increase the number of synonyms that we search in the corpus and thus increase the number of actually appearing synonyms.

The relation that leads to the lowest recall is the hypernymy relation (47.1%). In general, one can speculate that this is due to the fact that a hypernym of a term does not necessarily lay in the same domain – and therefore receives a lower relevance ranking. Nevertheless, it may be a very general term that occurs very often, such that the low relevance score is compensated or even overruled by the high frequency.

The term “plasma”, for instance, has three synsets, from which the first one is the most appropriate in our domain. Based on the synonyms only, our program returns all three synsets. But if we add the hypernymy relation, the third synset gets selected by our program. This mistake is due to the fact that this synset has a synset of “state” {state, state of matter} as one of its synonyms, which has not a high domain relevance but occurs extremely often. The first synset of “plasma” has “extracellular fluid” as hypernym, which does not occur at all.

The relations hyponymy, holonymy and meronymy clearly stay in the same domain. A term like “lip” is partially disambiguated by looking at its holonyms: “vessel” or “mouth”. Since “mouth” lays in the domain of human anatomy, its relevance score is higher than vessel.

It is no surprise either that the topic relation, that assigns a category to synsets, is among the relations leading to high recall values. However, as many synsets do have a related topic, it does not contribute to precision.

There is a clear benefit of using several relations together. This combination increases the number of included synonyms further than by using a single relation.

4 Conclusions and Future Work

We presented a domain-specific corpus-based approach to the lexical enrichment of ontologies, i.e. enriching a given ontology with lexical entries derived from a semantic lexicon such as WordNet. Our approach was empirically tested by an experiment on combining the FMA with WordNet synsets that were disambiguated by use of a corpus of Wikipedia pages on human anatomy. The approach was evaluated on a benchmark of 50 ambiguous FMA terms with manually assigned WordNet synsets. Results show that the approach performs better than a most-frequent sense baseline. Further refinements of the algorithm that include the use of WordNet relations such as hyponym, hypernym, meronym, etc. showed a much improved performance, which was again improved upon drastically by combining the best of these relations. In summary, we achieved good performance

on the defined task with relatively cheap methods. This will allow us to use our approach in large-scale automatic enrichment of ontologies with WordNet derived lexical information, i.e. in the context of the OntoSelect ontology library and search engine⁷ (Buitelaar et al. (2004)). In this context, lexically enriched ontologies will be represented by use of the LingInfo model for ontology-based lexicon representation (Buitelaar et al. (2006)).

5 Acknowledgements

We would like to thank Hans Hjelm of Stockholm University (Computational Linguistics Dept.) for making available the FMA term set and the Wikipedia anatomy corpus.

This research has been supported in part by the THESEUS Program in the MEDICO Project, which is funded by the German Federal Ministry of Economics and Technology under the grant number 01MQ07016. The responsibility for this publication lies with the authors.

⁷ <http://olp.dfki.de/ontoselect/>

Bibliography

- M. Alexa, B. Kreissig, M. Liepert, K. Reichenberger, L. Rostek, K. Rautmann, W. Scholze-Stubenrecht, and S. Stoye. The Duden Ontology: An Integrated Representation of Lexical and Ontological Information. *Proceedings of the OntoLex Workshop at LREC, Spain, May, 2002*.
- P. Buitelaar and B. Sacaleanu. Ranking and selecting synsets by domain relevance. *Proceedings of WordNet and Other Lexical Resources: Applications, Extensions and Customizations, NAACL 2001 Workshop, 2001*.
- P. Buitelaar, T. Eigner, and T. Declerck. OntoSelect: A Dynamic Ontology Library with Support for Ontology Selection. *Proceedings of the Demo Session at the International Semantic Web Conference. Hiroshima, Japan, 2004*.
- P. Buitelaar, T. Declerck, A. Frank, S. Racioppa, M. Kiesel, M. Sintek, R. Engel, M. Romanelli, D. Sonntag, B. Loos, et al. LingInfo: Design and Applications of a Model for the Integration of Linguistic Information in Ontologies. *Proceedings of OntoLex 2006, 2006*.
- J. Cohen. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement, 20(1):37, 1960*.
- A. Cucchiarelli and P. Velardi. Finding a domain-appropriate sense inventory for semantically tagging a corpus. *Natural Language Engineering, 4(04):325–344, 1998*.
- A. Dolbey, M. Ellsworth, and J. Scheffczyk. BioFrameNet: A Domain-specific FrameNet Extension with Links to Biomedical Ontologies. In *Proceedings of the "Biomedical Ontology in Action" Workshop at KR-MED, Baltimore, MD, USA, pages 87–94, 2006*.
- A. Gangemi, R. Navigli, and P. Velardi. The OntoWordNet Project: extension and axiomatization of conceptual relations in WordNet. In *Proceedings of ODBASE03 Conference. Springer, 2003*.
- T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition, 5(2):199–220, 1993*.
- N. Guarino. Formal ontology and information systems. In Nicola Guarino, editor, *Formal ontology in information systems*, pages 3–15. IOS Press, June 1998.
- R. Koeling and D. McCarthy. Sussx: WSD using Automatically Acquired Predominant Senses. In *Proceedings of the Fourth International Workshop on Semantic Evaluations*, pages 314–317. Association for Computational Linguistics, 2007.
- G. Leech, P. Rayson, and A. Wilson. *Word Frequencies in Written and Spoken English: Based on the British National Corpus*. Longman, 2001.
- B. Magnini and G. Cavaglia. Integrating subject field codes into WordNet. *Proceedings of LREC-2000, Second International Conference on Language Resources and Evaluation*, pages 1413–1418, 2000.
- B. Magnini, C. Strapparava, G. Pezzulo, and A. Gliozzo. Using domain information for word sense disambiguation. *Proceeding of SENSEVAL-2: Second*

- International Workshop on Evaluating Word Sense Disambiguation Systems*, pages 111–114, 2001.
- C. D. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. The MIT Press, Cambridge, Massachusetts, 1999. URL <http://nlp.stanford.edu/fsnlp/>.
- D. McCarthy, R. Koeling, J. Weeds, and J. Carroll. Finding predominant senses in untagged text. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*, pages 280–287, 2004a.
- D. McCarthy, R. Koeling, J. Weeds, and J. Carroll. Using automatically acquired predominant senses for word sense disambiguation. In *Proceedings of the ACL SENSEVAL-3 workshop*, pages 151–154, 2004b.
- R. Navigli and P. Velardi. Automatic Adaptation of WordNet to Domains. In *Proceedings of 3rd International Conference on Language Resources and Evaluation Conference (LREC) and OntoLex2002 workshop*, pages 1023–1027, 2002.
- D. Oberle, A. Ankolekar, P. Hitzler, P. Cimiano, C. Schmidt, M. Weiten, B. Loos, R. Porzel, H.-P. Zorn, V. Micelli, M. Sintek, M. Kiesel, B. Mougouie, S. Vembu, S. Baumann, M. Romanelli, P. Buitelaar, R. Engel, D. Sonntag, N. Reithinger, F. Burkhardt, and J. Zhou. Dolce ergo sumo: On foundational and domain models in swinto. *Journal of Web Semantics (accepted for publication)*, forthcoming.
- M. T. Paziienza and A. Stellato. An environment for semi-automatic annotation of ontological knowledge with linguistic content. In *Proceedings of the 3rd European Semantic Web Conference*, 2006.
- C. Rosse and J.L.V. Mejino Jr. A reference ontology for biomedical informatics: the foundational model of anatomy. *Journal of Biomedical Informatics*, 36(6): 478–500, 2003.
- H. Schmid. Probabilistic part-of-speech tagging using decision trees. *Proceedings of the conference on New Methods in Language Processing*, 12, 1994.