Using a Lexical Dictionary and a Folksonomy to Automatically Construct Domain Ontologies

Daniel Macías-Galindo, Wilson Wong,
Lawrence Cavedon, and John Thangarajah

School of Computer Science and I.T., RMIT University, Melbourne, Australia
{daniel.macias,wilson.wong,lawrence.cavedon,john.thangarajah}@rmit.edu.au

Abstract. We present and evaluate *MKBUILD*, a tool for creating domain-specific ontologies. These ontologies, which we call Modular Knowledge Bases (MKBs), contain concepts and associations imported from existing large-scale knowledge resources, in particular WordNet and Wikipedia. The combination of WordNet's human-crafted taxonomy and Wikipedia's semantic associations between articles produces a highly connected resource. Our MKBs are used by a *conversational agent* operating in a small computational environment. We constructed several domains with our technique, and then conducted an evaluation by asking human subjects to rate the domain-relevance of the concepts included in each MKB on a 3-point scale. The proposed methodology achieved precision values between 71% and 88% and recall between 37% and 95% in the evaluation, depending on how the middle-score judgements are interpreted. The results are encouraging considering the cross-domain nature of the construction process and the difficulty of representing concepts as opposed to terms.

1 Introduction

Conventional approaches to building domain ontologies typically rely on collections of domain text (i.e., ontology learning from text) or expert-crafted structured knowledge resources (e.g., WordNet [6], Cyc [11]). Such centralised approaches require enormous effort from domain experts and knowledge engineers; hence, these resources are slow to keep up with new knowledge and have considerably smaller coverage. The realisation of these drawbacks has resulted in the rise of an ontology construction approach using collaboratively-maintained resources: e.g., Freebase [3], YAGO [20] and DBPedia [1]. Despite the advantages of collaboratively maintained resources, issues of trustworthiness and subjectiveness related to social tagging can translate to poorer quality categorisations. For this reason, a backbone provided by expert-crafted resources is still desirable.

In this paper we present a methodology for construcing modular knowledge bases (MKBs) using WordNet and Wikipedia. As both resources provide their own strengths and shortcomings, their amalgamation increases the coverage and reliability of the resulting knowledge bases [8]. These MKBs combine the strengths of both resources as follows. The developer of the MKB first defines a domain using a Wikipedia article. A set of relevant concepts are extracted based

D. Wang and M. Reynolds (Eds.): AI 2011, LNAI 7106, pp. 638–647, 2011.

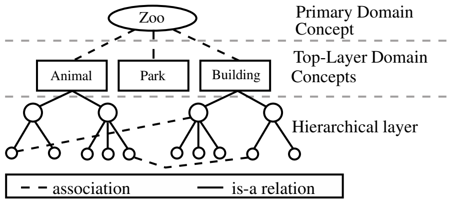

Fig. 1. Schema of an MKB

on being linked from the article. WordNet is then used to add parent and child concepts. Our methodology has been implemented as a tool called *MKBUILD* to construct MKBs for specific domains with minimal involvement from the developer. Our work on MKBs is motivated by the need to provide knowledge bases for a conversational agent designed to operate on a mobile platform with a small computational footprint. This agent is unable to accommodate large knowledge resources such as Cyc or DBPedia due to issues related to memory and storage size, and efficient access and processing. Our approach allows different MKBs to be loaded onto the platform as required depending on conversational flow.

The tasks of extracting domain-specific terms and of automatically constructing ontologies (typically using language processing techniques over Wikipedia or text corpora) have been widely studied: e.g., [12,15] for the former and [17,22] for the latter. To some degree, our approach combines these tasks. First, a set of "maximally general" concepts are extracted from Wikipedia and WordNet, which form the roots of the multiple sub-ontologies associated with a target domain. Second, the sub-ontologies rooted at each of these concepts are constructed, including *association* links between the concepts. These links form the basis of a generic *semantic relatedness* technique (not described here). The construction process of the MKBs is outlined in the following section, followed by the description of a user-based evaluation and a discussion.

2 Building Domain-Specific MKBs

In this section, we briefly describe the proposed methodology for building Modular Knowledge Bases[1]. An MKB is an ontology built around a main concept representative of a domain, and features a set of sub-taxonomies linked by the associations amongst its nodes (concepts). The target architecture of MKBs is shown in Figure 1. To build MKBs, we use two knowledge resources, namely, WordNet [6], a lexical dictionary that contains multiple word senses grouped by their meaning, and Wikipedia[2], an online encyclopaedia that operates like a collaborative wiki. WordNet features a taxonomy of concepts, but lacks relationships that are not lexical (e.g., ***Lion*** `lives in` ***Savannah***). On the other

[1] We omit some details, such as related work here for reasons of space: full details can be found in [13].

[2] `http://en.wikipedia.org/`

hand, Wikipedia does not have a taxonomical organisation; rather, we focus on the *wikilinks* featured in every article. A *wikilink* represents a concept that helps in the understanding of definitions[3]. Although *wikilinks* do not always describe a positive association, we are interested only in existence of such associations rather than their nature. The combination of Wikipedia *wikilinks* ("flexible" in the sense that humans themselves choose what to link in Wikipedia articles) and the WordNet hierarchy ("rigid" because property inheritance cannot be changed by humans) helps us to produce richer MKBs. *Wikilinks* have been previously analysed as a reliable set, though not absolute, of associations between articles [9,16]. For our approach, we use unidirectional *wikilinks* instead of mutual (from article a to b and vice versa) since we are interested in using such associations for conversational topic transitions. Thus, we are prepared to tolerate a more liberal notion of "relatedness".

Our process for constructing domain-specific MKBs consists of the following three stages. An overview of the process is shown in Figure 2:

1. **Define the domain**, i.e., select the *primary domain concept* by choosing a Wikipedia article that unambiguously reflects the main concept of the target domain;
2. **Build the top-layer** by extracting concepts to represent the most general and representative concepts associated with the *primary domain concept*;
3. **Extend the MKB** by adding sub-concepts to each top layer concept and analysing, for each concept's articles, the corresponding *wikilinks*.

The first stage of this process is performed manually, where the module designer chooses a Wikipedia entry that best matches the domain of the MKB. In this work, we refer to the selected entry's identifier (which may be qualified by a specific "sense" for ambiguous terms) as the *primary domain concept*.

The next two stages of the process are executed by the *MKBUILD* tool. *MKBUILD* performs all tasks necessary for those stages and produces an MKB *automatically*. *MKBUILD* has been developed in Java and uses the OWL-API Library[4] for handling the ontology. All these stages may be performed separately using *MKBUILD*, thus allowing intermediate manual modifications to the MKB in order to improve the coverage of the module. The rest of this section contains a brief description of the process. For full details, see [13].

Stage 2. Building the MKB Top Layer

The *primary domain concept* identified in the first stage is used as the input to *MKBUILD*, which performs Stages 2 and 3 automatically. In Stage 2, the concepts that form the top layer of the MKB are discovered. The tasks that comprise this stage are briefly described below.

2.1. Page link extraction. *MKBUILD* retrieves all terms that appear as *wikilinks* in the article referenced by the *primary domain concept*. This extraction

[3] See `http://en.wikipedia.org/wiki/Wikipedia:Manual_of_Style`

[4] `http://owlapi.sourceforge.net/`

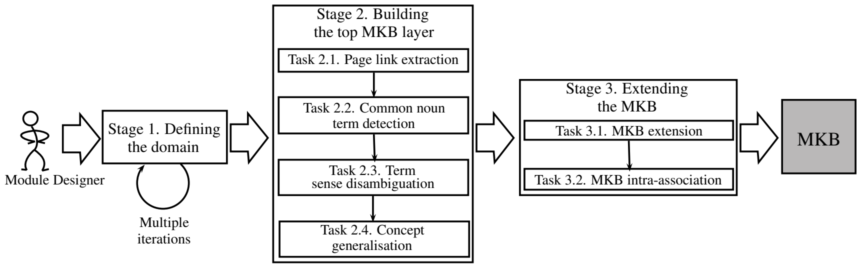

Fig. 2. An overview of the process to build MKBs

process is performed using DBPedia [1] (version 3.5.1), which contains Wikipedia links stored as triplets. MKB also extracts any *redirect links* that accompany each term, as these contain the original name of the Wikipedia article (i.e., *wikilinks* are proposed by authors; *redirect links* reconcile other concepts to point to the same article). In contrast to previous work that has considered the category structure provided by Wikipedia [7,10,18], we propose the use of *wikilinks* as the initial source of concepts directly related to a domain. We do, however, propose to leverage Wikipedia's category-based hierarchical *folksonomy* in future improvements, as discussed in the Evaluation section. The *wikilink* terms extracted are validated using a *named entity recognition (NER)* tool[5] and a "Wikipedia-to-WordNet" conversion table provided by DBPedia. At the end of this task, *MKBUILD* obtains a set of *preliminary concept terms*.

2.2. Common noun term detection. The preliminary terms may refer to either concepts or *instances* of concepts (e.g., specific people or places) as Wikipedia itself does not distinguish between the two [9]. This task performs a second detection and removal of terms that correspond to instances. These terms are detected using two tools: a Part-of-Speech (POS) tagger implemented in the Language Technology tool MorphAdorner[6], and (ii) WordNet word forms. Terms are retained for the next step as long as MorphAdorner determines that they contain at least one common noun and no proper nouns, proper adjectives nor non-English words. Additionally, WordNet helps with removing terms that start with a capital letter, as this has proven to be a sufficient heuristic to determine instances [14]. After this task is performed, a list of terms is obtained.

2.3. Term sense disambiguation. Terms retained in the above step may be ambiguous, in that they have multiple senses in WordNet. Consequently, to obtain concepts, a disambiguation process is required. This process finds the concepts that are related to the *primary domain concept* using semantic similarity measure of Lesk, adapted to WordNet glosses[7][2].

[5] The Stanford NER tool, that can be obtained from `http://nlp.stanford.edu/ner/`

[6] `http://morphadorner.northwestern.edu/`

[7] This value is obtained from the Java WordNet:Similarity Library, available in: `http://www.cogs.susx.ac.uk/users/drh21/`.

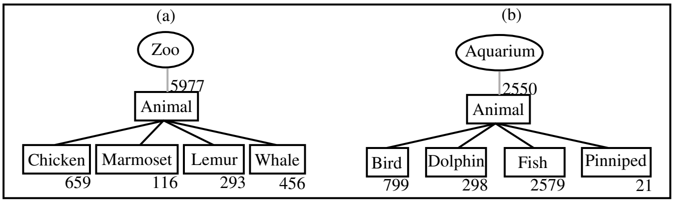

Fig. 3. Detection of more general classes via WordNet: (a) an accepted generalisation; and (b) a rejected generalisation. The top oval corresponds to the *primary domain concept*, and each number represents the co-occurrence between the *pdc* and a concept.

2.4. Concept generalisation. The concepts obtained in the previous task may not represent the level of generality required for the domain (i.e., the domain covers more general concepts than those identified). Concept generalisation requires extracting all WordNet *hypernyms* (super-classes) of the concepts obtained in step 2.3. This task is executed by the following two steps:
(i) *Generalisation using available concepts:* In this step, *MKBUILD* removes a concept if another concept in the list is its parent, as they will be later added as sub-concepts of the corresponding top-layer concept at a later stage.
(ii) *Generalisation using WordNet hierarchy: MKBUILD* detects if two or more concepts $sc_i, \ldots sc_j$ can be generalised using a common super-class h. If a super-class is detected, *MKBUILD* compares the co-occurrence of the *primary domain concept* (*pdc*) and h against the co-occurrence of *pdc* and each concept $sc_i, \ldots sc_j$ using Wikipedia articles as a corpus. If concept h is more commonly associated with the *pdc* than the sum of all *sc*, then the sub-concepts are replaced by h in the list of related concepts. An example of this is shown in Figure 3.

Stage 3. Building the Hierarchical Layer

With a top-layer of concepts obtained from Stage 2, two more tasks are performed before an initial version of the MKB is produced. In the first task, *MKBUILD* adds sub-classes from WordNet below each top-layer concept, which now become the root nodes of sub-ontologies. As in Stage 2, only WordNet senses that are common nouns are included. Finally, in the second task, *MKBUILD* adds association links between concepts that are not lexically based. These association links support a notion of *semantic relatedness* featuring more general links between concepts. These links are used by our conversational agent for concept-based topic transitions. *MKBUILD* inserts an association between two concepts if a *wikilink* between the articles corresponding to those concepts exists in DBPedia (as long as there is not already a lexical link from WordNet).

3 Evaluation

In this section, we describe an evaluation of the Stage 2 of *MKBUILD*, i.e., identifying the *top-layer domain concepts* of the sub-ontologies related to the specified domain.[8] We conducted a user study by asking subjects to judge whether the *top-layer domain concepts* extracted by *MKBUILD* were appropriate to the domain. We focus on evaluating *precision* and *recall* of the extraction process for *top-layer domain concepts* and not the hierarchy below, since concepts in the hierarchy below a *top-layer domain concept* are assumed to be related to it.

Setup. We used *MKBUILD* to construct MKBs for 14 domains, which are shown in the first column of Table 1. The total number of *top-layer domain concepts* (*tldc*) across all domains is 490 (set T). We extracted a subset of T, namely T', with the highest *idf* in each domain. T' was distributed across 6 different survey files. The breakdown of these concepts according to the different domains is summarised in columns 2 and 3 of Table 1. Each survey contains 3 domains, each domain comprises up to 10 concepts.

We asked 55 anonymous users to score how "related" each *tldc* is to a proposed domain D. Surveys were randomly assigned, following an even distribution across users. Users scored each domain-concept pair with an integer number of either 2, 1 or 0, where 2 indicates that the concept is highly related to D, 1 indicates it is related, and 0 for unrelated concepts. Users could also separately select *Unsure*. Users were also requested to add, for each domain, a set of up to five concepts that were not in the survey but what they considered to be highly related to the domain.

We obtained assessment scores from between 8 and 10 participants for each survey. We calculated the average Pearson correlation between subjects for each survey, obtaining values ranging from 0.28 (indicating medium low correlation) to 0.54 (strong correlation)[5]. Although these values indicate some agreement, these also show the difficulty of finding similarly scored participations.

Results. To determine users' agreement with the system for each *top-layer domain concept*, we calculated an aggregated value in three different ways, each representing a different assessment of relatedness. First, p_a (i.e., precision) was calculated by adding the number of participants scoring either 1 (i.e., "related") or 2 ("highly related") and subtracting the number of 0's ("unrelated") scored for each *tldc*. Second, the scores of 1's were changed to 0.5 to calculate p_b. Third, p_c took into consideration only the number of 0's and 2's, with the number of 1's used only to break any ties (i.e., the numbers of 0's and 2's were the same). The first criterion is standard according to the definition of our experiment, which is that both scores of 1 and 2 represent a certain degree of relatedness. The latter two criteria represent a less generous interpretation of the middle score (i.e., 1). These criteria bias against our system, hence we include them for comparison.

Using the total number of concepts together with the aggregated values obtained as per the three criteria, we calculated the *Precision* and *Recall* for all domains, as defined by [19]. We employ these measures as they reflect the coverage of the concepts with respect to the target domain. Our evaluation of the 36% of all the available 490 *top-layer domain concepts* resulted in the following

[8] Evaluating other stages would be effectively evaluating WordNet and DBpedia.

precision values, namely, $p_a = 0.88$, $p_b = 0.80$ and $p_c = 0.71$. These values reflect a high number of human participants agreeing with the *top-layer domain concepts* extracted by *MKBUILD*, particularly on the standard interpretation of the middle score.

Next, we estimated recall using the *top-layer domain concepts* deemed as related, plus the extra concepts provided by participants. Only 38 out of the 55 participants provided any extra concepts; a total of 366 extra concepts were provided, ranging from 12 to 46 per domain. Due to the lack of a gold standard, we artificially created one with these extra concepts and the scores obtained from the provided *top-layer domain concepts*. We analysed these extra concepts in two ways: first, assuming that it was due to a lack of coverage of WordNet or Wikipedia that such concepts were not added to the MKB (method d); and second, assuming that all the suggested extra concepts should be in the MKB (method e). These concepts are proposed as our *false negatives*, while the concepts with a positive score are the *true positives*. These criteria affected the results for recall, which are shown in Table 1 as $r_{m|i}$, where m is method d or e and i refers to the method for calculating precision, as described above.

Table 1. Sample domains with their evaluated precision and recall values

Domain(D)	T	T'	p_a	p_b	p_c	$r_{d\|a}$	$r_{d\|b}$	$r_{d\|c}$	$r_{e\|a}$	$r_{e\|b}$	$r_{e\|c}$
Amusement park	26	10	0.9	0.7	0.6	1	1	1	0.6	0.54	0.5
Association football	25	10	0.9	0.8	0.8	1	1	1	0.56	0.53	0.53
Automobile	41	20	0.85	0.4	0.4	0.89	0.8	0.8	0.35	0.2	0.2
Beach	28	10	1	1	1	1	1	1	0.38	0.38	0.38
Computer	73	20	0.95	0.9	0.8	0.86	0.85	0.84	0.39	0.38	0.35
Economy	56	20	0.85	0.85	0.75	1	1	1	0.47	0.47	0.44
Food	88	20	0.9	0.8	0.55	1	1	1	0.49	0.46	0.37
Museum	32	10	1	0.8	0.5	1	1	1	0.56	0.5	0.38
Music	37	10	0.9	0.9	0.9	0.82	0.82	0.82	0.45	0.45	0.45
Public aquarium	11	10	0.5	0.5	0.4	0.83	0.83	0.8	0.25	0.25	0.21
School	25	10	0.8	0.7	0.7	1	1	1	0.32	0.29	0.29
Sport	24	10	0.9	0.8	0.7	1	1	1	0.53	0.5	0.47
Theatre	18	10	0.9	0.9	0.9	1	1	1	0.36	0.36	0.36
Zoo	8	8	1	0.75	0.5	1	1	1	0.42	0.35	0.27
Total	490	178	0.88	0.8	0.71	0.95	0.95	0.94	0.42	0.4	0.37

We do not have a comparable task for direct comparison, but can compare to performance in domain term extraction; e.g., for this task, [12] reported values of precision and recall of 0.354 and 0.183 respectively. [15] obtained an F1 quality score of 0.25 in term extraction using the Web. On the other hand, our lowest F1 score reported is 0.486 for p_c and $r_{e|c}$. Some care has to be taken when interpreting these figures because there are clear differences between our approach and domain term extraction which makes them not comparable. First, we focus on extracting concepts, not just terms, so we have to resolve against concepts (which includes performing word sense disambiguation). Second, term extraction is commonly applied in closed environments using well-defined domain corpora, whereas we extract from a resource as broad as Wikipedia. Hence, in comparison to this (related) baseline task, we consider our results as encouraging.

Error Analysis. We can analyse the set of concepts suggested by participants to obtain insights into *MKBUILD*'s inability to extract certain concepts. The suggested concepts can be classified in four ways, namely, (A) they ambiguously refer to proper instead of common nouns (e.g., Shakespeare) or to other parts of speech besides noun (e.g., play), (B) they are in an MKB but were not shown to the user, and (C) they do not appear anywhere in the MKB. From the set of 366 suggested concepts, 8 concepts fall under category (A), 151 under (B), and 207 under (C).

From category (C), we can create three subgroups. Group (C1) contains those concepts that do not appear in WordNet. Analogously, group (C2) contains concepts that do not appear in the Wikipedia article of the domain as *wikilinks*. Group (C3) contains those suggested concepts appearing in both WordNet and Wikipedia which did not appear in the resulting MKBs. Concepts in group (C2) represent the largest limitation of our approach, showing that using only the *primary domain concept* article is not enough to find concepts associated with the domain. Earlier, we mentioned that *MKBUILD* does not currently use Wikipedia's *folksonomy*. Therefore, a broader, more systematic exploration of related articles, considering the article categorisation in Wikipedia, should be performed in future work.

Only 8 suggested concepts fall within group (C3). These suggested concepts missing from the MKBs are classified into four types. The first type of missing suggested concept, namely (C3-a) features those concepts with an ambiguous WordNet taxonomy. For example, the concept ***Dolphin*** has two different senses, where one corresponds to its meat, and the other defines a type of ***Mammal***. If two concepts have similar names and no other synonyms available, *MKBUILD* is unable to create a new concept, thus the concept referring to the second sense and its children concepts are not included. This issue can be resolved if by analysing the definitions of concepts according to WordNet. In cases where some definitions for different senses of a word are complimentary (e.g., Dolphin is an edible fish AND a mammal) we must merge both senses in our produced MKB.

The second type (C3-b) occurs with the NER tool (Step 2.1), which performs suboptimally due to the lack of context for terms. Therefore, some concepts that correspond to common nouns are treated as referring to instances, and are removed from the process. For example, the term ***Algorithm*** is recognised as an entity expressing a location.

Finally, the third type, (C3-c) occurs due to our heuristic to identify instances using WordNet. Our approach automatically eliminates a term if it contains a word form (a synonym) starting with a capitalised letter. This applies to concepts such as ***hydrogen***, which can be also represented with the letter "H".

Error type (C3-c) is the most frequent, occurring with four suggested concepts. Error (C3-b) was detected on three occasions and (C3-a) only once. This means that in order to improve entity recognition, we have to use longer texts rather than only terms. One possibility is to feed the NER tool with a sentence from the short abstract extracted from Wikipedia containing the analysed term.

4 Conclusions and Future Work

We have described a process for constructing domain-specific ontologies, called Modular Knowledge Bases, to be used by a conversational agent with a modular infrastructure. The process has been programmed as *MKBUILD*, a tool that allows the *automatic* extraction of concepts and relations specific to a given domain using large resources such as WordNet and Wikipedia/DBPedia. The ontology construction process we described saves developers a significant amount of effort in constructing an ontology specific to a conversational domain, while at the same time allowing the developers to easily intervene at any point in time to correct any egregious errors.

We have conducted an experiment involving human assessors to determine the precision and recall of a critical stage of the construction process, namely identifying the top-level concepts for the domain-specific ontologies. We obtained encouraging results considering the difficulty of cross-domain concept extraction. This experiment has also allowed us to determine that the exploration of only the Wikipedia article associated with the *primary domain concept* of the MKB is insufficient. Other related Wikipedia articles have to be considered in order to extract a broader range of domain-specific concepts. We also discussed limitations in the current extraction process and proposed solutions.

Our main application of the domain ontologies constructed using *MKBUILD* is to generate a *Topic Network* that can be used to link conversational fragments together into more coherent longer-running threads, using ontology-based semantic similarity measures. We are also conducting an evaluation to measure ontology-based semantic relatedness involving sets of relations that go beyond previously considered (e.g., [4,18]), and evaluate its efficacy in topic transitioning in conversational dialogue. Other future work includes extending the coverage of the concepts and relations in the MKBs through the use of other large knowledge bases constructed using information extraction techniques (e.g., [21]).

Acknowledgements. The first author acknowledges scholarship 201228 from the Consejo Nacional de Ciencia y Tecnologia (CONACYT), Mexico. This work was partially supported by Australian Research Council Linkage Projects LP0882013 and LP110100050. We acknowledge the support and collaboration of our partner, Realthing Entertainment Pty Ltd. Aidan Martin implemented the framework for data collection.

References

1. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.G.: DBpedia: A Nucleus for a Web of Open Data. In: Aberer, K., Choi, K.-S., Noy, N., Allemang, D., Lee, K.-I., Nixon, L.J.B., Golbeck, J., Mika, P., Maynard, D., Mizoguchi, R., Schreiber, G., Cudré-Mauroux, P. (eds.) ASWC 2007 and ISWC 2007. LNCS, vol. 4825, pp. 722–735. Springer, Heidelberg (2007)

2. Banerjee, S., Pedersen, T.: An Adapted Lesk Algorithm for Word Sense Disambiguation Using Wordnet. In: Gelbukh, A. (ed.) CICLing 2002. LNCS, vol. 2276, pp. 136–145. Springer, Heidelberg (2002)
3. Bollacker, K., Evans, C., Paritosh, P., Sturge, T., Taylor, J.: Freebase: a Collaboratively Created Graph Database for Structuring Human Knowledge. In: SIGMOD, pp. 1247–1250 (2008)
4. Budanitsky, A., Hirst, G.: Evaluating WordNet-based Measures of Lexical Semantic Relatedness. Computational Linguistics 32(1), 13–47 (2006)
5. Cohen, J.: Statistical power analysis for the behavioral sciences, 2nd rev. edn. Academic Press, London (1977)
6. Fellbaum, C.: WordNet: an electronic lexical database. The MIT Press (1998)
7. Grieser, K., Baldwin, T., Bohnert, F., Sonenberg, L.: Using Ontological and Document Similarity to Estimate Museum Exhibit Relatedness. J. Computing and Cultural Heritage 3(3), 10:1–10:20 (2011)
8. Gruber, T.: Ontology of Folksonomy: A Mash-up of Apples and Oranges. Semantic Web and Information Systems 3(2), 1–11 (2007)
9. Hepp, M., Siorpaes, K., Bachlechner, D.: Harvesting Wiki Consensus: Using Wikipedia Entries as Vocabulary for Knowledge Management. Internet Computing 11, 54–65 (2007)
10. Herbelot, A., Copestake, A.: Acquiring Ontological Relationships from Wikipedia Using RMRS. In: Web Content Mining with Human Language Technologies, pp. 1–10 (2006)
11. Lenat, D., Guha, R.V.: Building Large Knowledge-Based Systems; Representation and Inference in the Cyc Project. Addison-Wesley (1990)
12. Liu, Z., Huang, W., Zheng, Y., Sun, M.: Automatic Keyphrase Extraction via Topic Decomposition. In: EMNLP, pp. 366–376 (2010)
13. Macias-Galindo, D., Cavedon, L., Thangarajah, J.: Building Modular Knowledge Bases for Conversational Agents. In: Knowledge Representation and Reasoning for Practical Dialogue Systems, pp. 16–23 (2011)
14. Martin, P.: Correction and Extension of WordNet 1.7. In: Ganter, B., de Moor, A., Lex, W. (eds.) ICCS 2003. LNCS, vol. 2746, pp. 160–173. Springer, Heidelberg (2003)
15. Massey, L., Wong, W.: A Cognitive-Based Approach to Identify Topics in Text Using the Web as a Knowledge Source. In: Wong, W., Liu, W., Bennamoun, M. (eds.) Ontology Learning and Knowledge Discovery Using the Web. IGI Global (2011)
16. Milne, D., Medelyan, O., Witten, I.H.: Mining Domain-specific Thesauri from Wikipedia: A case study. In: Web Intelligence, pp. 442–448 (2006)
17. Ponzetto, S.P., Strube, M.: Deriving a Large Scale Taxonomy from Wikipedia. In: The National Conference on Artificial Intelligence, vol. 22, pp. 1440–1446 (2007)
18. Ponzetto, S.P., Strube, M.: Knowledge Derived From Wikipedia For Computing Semantic Relatedness. J. AI Research 30, 181–212 (2007)
19. Sabou, M., Wroe, C., Goble, C., Mishne, G.: Learning domain ontologies for web service descriptions: an experiment in bioinformatics. In: World Wide Web, pp. 190–198 (2005)
20. Suchanek, F.M., Kasneci, G., Weikum, G.: YAGO: A Large Ontology from Wikipedia and WordNet. Web Semantics 6(3), 203–217 (2008)
21. Yates, A., Cafarella, M., Banko, M., Etzioni, O., Broadhead, M., Soderland, S.: TextRunner: open information extraction on the web. In: NAACL-Demos, pp. 25–26 (2007)
22. Zirn, C., Nastase, V., Strube, M.: Distinguishing between Instances and Classes in the Wikipedia Taxonomy. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) ESWC 2008. LNCS, vol. 5021, pp. 376–387. Springer, Heidelberg (2008)