



Rules for Inducing Hierarchies from Social Tagging Data

Hang Dong^{1,2(✉)}, Wei Wang², and Frans Coenen¹

¹ Department of Computer Science, University of Liverpool, Liverpool, UK
{HangDong, Coenen}@liverpool.ac.uk

² Department of Computer Science and Software Engineering,
Xi'an Jiaotong-Liverpool University, Suzhou, China
Wei.Wang03@xjtlu.edu.cn

Abstract. Automatic generation of hierarchies from social tags is a challenging task. We identified three rules, set inclusion, graph centrality and information-theoretic condition from the literature and proposed two new rules, fuzzy set inclusion and probabilistic association to induce hierarchical relations. We proposed an hierarchy generation algorithm, which can incorporate each rule with different data representations, i.e., resource and Probabilistic Topic Model based representations. The learned hierarchies were compared to some of the widely used reference concept hierarchies. We found that probabilistic association and set inclusion based rules helped produce better quality hierarchies according to the evaluation metrics.

1 Introduction

Tagging is a popular functionality on many social media platforms. Users can add “free” words (tags) to describe shared resources to facilitate searching and recommendation. These user-generated tags form a taxonomy of users’ terminologies called folksonomy. There have been great interests and motivation to automatically induce knowledge structures from social media data. The most challenging issue is the intrinsic difficulties (e.g., sparse data, weak context, polysemy, concatenated tags and typos [10]) to learn large-scale concept hierarchies in various domains. Even for human being this requires considerable cognitive work [19]. Although there have been some general guidelines to associate semantics to tags [7], no theoretical consensus were reached on the rules and assumptions to determine semantic relations among tags.

There have been methods using different rules to derive hierarchical structures from social tags. However, there is no comprehensive research on analysis and comparison of these rules in a rigorous manner. The work in [16] compared four clustering and generality based techniques and found that generality-based methods in general outperform the clustering based ones. It primarily focused on evaluation techniques and did not analyse how the underlying rules and their assumptions affected the results. Furthermore, other important rules such as set inclusion and information-theoretic condition were not discussed.

A rule in the context of this work is defined as certain conditions that need to be satisfied to determine whether a pair of words (tags) has a directed hierarchical or subsumption relation. We summarise existing rules used in learning hierarchical relations from social tagging data and propose two new rules: probabilistic association and fuzzy set inclusion based rules. Then we evaluate the performance of these rules in learning relations and compare the results to three well-know reference hierarchies. We aim to conduct an in-depth study for the following two questions:

- Q1.** Which rule can effectively capture the hierarchical relations between social tags?
- Q2.** How do rules and data representation techniques affect the quality of the induced concept hierarchies?

Our contributions can be summarised as follows. First, we performed a thorough analysis on the existing rules and proposed two new rules. Second, we designed an algorithm that can accommodate different rules to iteratively learn a knowledge hierarchy. The algorithm ranks all pairwise tag pairs by similarity in descending order, and passes the tag pairs to a rule to determine their relations. Third, we used three reference knowledge hierarchies (DBPedia, Microsoft Concept Graph and ACM Computing Classification System) to evaluate the quality of the learned hierarchies based on standard metrics (taxonomic precision, recall and F-measure). To our best knowledge, this is the first comprehensive study on the rules for learning knowledge hierarchies from social tagging data.

2 Hierarchical Relation: Definitions and Rules

Creation of concept hierarchies representing the world knowledge is a difficult cognitive task that requires lot of effort [19, p. 139]. Hierarchical relation is a paradigmatic relation, which means that two words should fit into the same grammatical slot or be of the same semantic type [3]. It is different from syntagmatic relations which exist between concepts in specific documents or other contexts [15]. In this way, co-occurrence of tags is syntagmatic relation, but provides great source for paradigmatic relations, as stated in [14, p. 223]. Hierarchical relations can be divided into hyponymy and meronymy. In some ontology models, they are merged and referred to as broader/narrower relations (e.g., SKOS¹).

We take a mixed view of hierarchical relations, with a focus on hyponymy. Although it seems intuitive to tell whether there is a hierarchical relation between two words, the idea of hyponymy is not straightforward. In the linguistic domain, Cruse [3] gave three definitions of hyponymy and proposed a prototype-theoretical characterisation to derive hyponymy. In the first definition, hyponymy was conceptualised in a **logical** way both extensionally and intensionally. Extensionally, X is a hyponym of Y iff the form $\forall x[X'(x) \rightarrow Y'(x)]$

¹ <https://www.w3.org/TR/skos-primer/>.

is satisfied, but none of the form $\forall x[Y'(x) \rightarrow X'(x)]$ holds, where X' and Y' are the logical constants corresponding to the concepts X and Y , and x can be understood as an instance object. This is equivalent to say that the extension of X' should be included in the extension of Y' . Intensionally, “ X is a hyponym of Y iff $F(X)$ entails, but is not entailed by $F(Y)$ ”, where $F(-)$ is a sentential function satisfied by X or Y . The second definition utilises the **collocational** property. The more restricted a word is through the collocational normality, then the more specific it is, or more formally, “ X is a hyponym of Y iff the normal context of X is a subset of the normal context of Y ”. The third idea for defining hyponymy is **componential**. “ X is a hyponym of Y iff the features defining Y are a proper subset of features defining X ”. Based on this idea, to analyse the degree of inclusion between X and Y , “prototypes” of characterisation of hyponymy were proposed [3]. Stock also pointed out a case that holds in most occasions: there is reciprocity between the extension and intension of concepts in a hierarchical chain [15], in other words, specific terms, with further restrictive properties, tend to have less number of objects than general terms.

Moving from linguistic definitions to machine computation, we summarised the main rules found in literature below. We adapted the fuzzy set theory with probabilistic topic representation and named it fuzzy set inclusion. We also proposed another new rule called probabilistic association, which has a strong probabilistic foundation. We use **R1** to **R5** to refer to these rules.

Data Representation. A tag can be represented as a vector over all resources. Let $v[R_i]$ refer to the number of times the tag used to annotate the resource R_i , and i be the index for resources. This **resource-based representation** can capture co-occurrence of tags and is usually high dimensional and sparse. Another representation is based on a Probabilistic Topic Model (**PTM**), representing a tag as a vector over latent (or hidden) topics. PTM is generally a dimension reduction technique and each topic can be represented as a probabilistic distribution of the vocabulary. We used the standard Latent Dirichlet Allocation (LDA) [8] for PTM representation. With LDA inference, we can obtain two matrices, the distribution of hidden topics for each document, $p(\text{Topic}|\text{Doc})$ and the distribution of tags for each hidden topic $p(\text{Tag}|\text{Topic})$. Using the Bayesian’s Theorem, we can estimate $p(\text{Topic}|\text{Tag})$.

R1: Set Inclusion. The logical extension of a tag can be quantified using its contexts, i.e., the set of resources it annotated or the set of users who used it for annotation. Through measuring the degree of inclusion between the contexts of two tags, Mika [13] generated a lightweight taxonomy from social tags. De Meo [12] defined this formally as Inclusion measure, $\text{set-inc}(A, B) = \frac{|R_A \cap R_B|}{|R_A|}$, where R is the set of resource annotated by the tag A or B . Formally, **tag A is a hyponym of tag B if $\text{set-inc}(A, B) \geq p \wedge \text{set-inc}(B, A) < p \wedge \text{sim}(A, B) > s$** , where p is experimentally set as 0.5. For all the rules R1–R5, the function $\text{sim}(A, B)$ measures the similarity (cosine similarity in this study) between A

and B; and s is a similarity threshold to be determined by the researcher (see **Data Collection** in Sect. 4).

R2: Graph Centrality. Graph centrality measures the influence or popularity of an object in a social network. The work in [9] assumed that there is a latent taxonomy underlying the similarity graph of social data, and centrality can be used to mine the taxonomy. A *tag similarity graph* is generated through linking each pair of tags having similarity greater than a threshold. It is assumed that in such a graph, nodes with higher centrality are more popular and thus are more general. Both degree and betweenness centrality can be used [1,9]. The rule states that **tag A is a hyponym of tag B if $\text{graph-cent}(A) < \text{graph-cent}(B) \wedge \text{sim}(A, B) > s$** , where *graph-cent()* computes the centrality of a node, we used degree centrality in this study.

R3: Fuzzy Set Inclusion. We propose this rule adapting the set inclusion rule to the PTM representation. We define a fuzzy set of a tag A as a pair (U, m) , where U is the set of topics for the tag and $m:U \rightarrow [0, 1]$ is a membership function defined by $p(\text{Tag}|\text{Topic})$: for each topic $z \in U$, $m(z) = p(A|z)$. Therefore, the inclusion degree [17] of a fuzzy set S_A of S_B for a pair of tags A and B is defined as $\text{fuzzy-set-inc}(S_A, S_B) = \frac{\sum_i \min(S_{A_i}, S_{B_i})}{\sum_i S_{B_i}}$. Instead of measuring the resource set inclusion between a pair of tags, the fuzzy set inclusion rule measures the topic set inclusion in a low dimensional space. We assume that **tag A is a hyponym of tag B if $\text{fuzzy-set-inc}(S_A, S_B) \geq p \wedge \text{fuzzy-set-inc}(S_B, S_A) < p \wedge \text{sim}(A, B) > s$** , where p is set as 0.5.

R4: Information-Theoretic Condition. The hierarchical relations between tags were captured using the information-theoretic principle which measures the difference of relative entropy, i.e., Kullback-Leibler (KL) Divergence, from one tag topic distribution to another [18]. KL Divergence is defined as $D_{KL}(P_A||P_B) = \sum_i P_{A_i} \log \frac{P_{A_i}}{P_{B_i}}$, where P_A and P_B are probabilistic distributions of hidden topics for a pair of tags A and B, based on the PTM representation. It represents the information wasted by encoding events with one distribution P_A with a code based on another distribution P_B . It is elaborated as the average “surprise” received from a tag A, when it is expected to receive B. The information-theoretic condition [18] is constructed as follows, **tag A is a hyponym of tag B if $D_{KL}(P_B||P_A) - D_{KL}(P_A||P_B) < f \wedge \text{sim}(A, B) > s$** , where f is a noise factor with a very small value (0.05 in this study).

R5: Probabilistic Association. We propose another rule called probabilistic association, which quantifies the (relative) componential features (in this case, topics) of a tag. Based on the PTM-representation, the probabilistic association [8], $p(A|B)$, can be computed by marginalising over topics, $p(A|B) = \sum_z p(A|z)p(z|B)$, where $p(z|B) \propto p(B|z)p(z)$. Based on the fact that tag A and B are conditionally independent given a topic z , $p(z)$ can be estimated using

either uniform distribution or the fraction of tag annotations assigned to the topic z to all tag annotations. The rule is defined as: **tag A is a hyponym of tag B** if $p(A|B) < p(B|A) \wedge \text{sim}(A, B) > s$. This measures the relative topic components of a tag given another tag. If a tag A is highly associated by the topics contained in another tag B, while B is less associated from the topics contained in tag A, then tag B tends to be more general.

3 Methodology: Algorithm for Hierarchy Generation

To facilitate comparison of the rules R1–R5, we designed an algorithm for hierarchy generation. A hierarchy in this work is defined as a Directed Acyclic Graph G , $G = \{V, E\}$, where V is a set of vertices (tags) and $E \subseteq V \times V$ is a set of edges. Algorithm 1 incorporates a rule to select candidate hierarchical tag pairs whose similarity exceeds the threshold. Starting from tag pair with highest similarity, an edge is iteratively added to the graph G if the rule is satisfied. The algorithm assumes the learned hierarchies have a single parent tag for each tag (i.e. mono-hierarchies) and all established relations are direct and transitive.

Algorithm 1. Hierarchy generation algorithm using *any* rules.

Input: $L_{PairSim}$ is a pre-computed list of tag pairs $\langle t_i, t_j \rangle$ ranked in descending order by similarity. s is a similarity threshold for a tag pair. $\text{sim}(t_i, t_j)$ computes the cosine similarity between t_i and t_j ; $\text{hasParent}(t, G)$ returns a boolean indicating whether tag t has a parent node in G ; $\text{isHypo}(t_i, t_j)$ determines whether t_i is a hyponym of t_j .

Output: G , an induced hierarchy as a directed graph.

```

1 Initialise  $G$ ;
2 for  $i \leftarrow 1$  to  $|L_{PairSim}|$  do
3    $\langle t_i, t_j \rangle \leftarrow L_{PairSim}[i]$ ;
4   if  $\text{sim}(t_i, t_j) < s$  then
5     continue to the next  $i$ ;
6   end
7   if NOT  $\text{hasParent}(t_i, G)$  then
8     if  $\text{isHypo}(t_i, t_j)$  then
9        $G \leftarrow G \cup \langle t_i, t_j \rangle$ ;
10    end
11  end
12 end

```

To answer Q1, we replaced the $\text{isHypo}(t_i, t_j)$ with different rules to generate different hierarchies. To answer Q2, we represented tags using resource-based and PTM representations respectively to generate two different sets of $L_{PairSim}$. Different s for the two representations were used to ensure the same number of tag pairs passed to a rule. The set inclusion rule is based on resource-based representation; fuzzy set inclusion, information-theoretic condition and probabilistic

association rules are based on the PTM representation; and graph centrality can be based on any of the two representations.

4 Experiments

In this section, we present the experiments and evaluation results and discuss our findings in relation to the two research questions.

Data Collection. The academic social bookmarking data Bibsonomy was used. The whole dataset² comprises 3,794,882 annotations, 868,015 distinct resources, 283,858 distinct tags annotated by 11,103 users. We followed the steps in [6] to clean the dataset and further removed resources annotated with less than 3 tags. The cleaned dataset contains 7,846 tags and 128,782 resources. Each resource is a bag-of-tags, to represent tags using PTM, we set the number of topic as 1000 minimising perplexity of LDA model. The similarity threshold s were set as 0.2 and 0.038 for PTM and resource-based representation resp. to control same number (2.3 million) of pairwise tag pairs passed to a rule. For reference-based evaluation, we chose three gold-standard hierarchies of different nature, the crowd-sourced DBpedia, the web-mined Microsoft Concept Graph (MCG) and the expert-crafted ACM Computing Classification System (CCS). DBpedia³ is a knowledge base containing structured information from Wikipedia. MCG⁴ is a knowledge base mined from billions of web pages. ACM computing classification system (CCS)⁵ is a classification system to organise ACM publications by subjects and support retrieval. In Bibsonomy, there are 6,616 common concepts overlapped with DBpedia, 6,029 with MCG and 691 with CCS.

Evaluation Metrics. With the cleaned dataset, we used Algorithm 1 to generate several hierarchies based on rules R1–R5 and different data representations. A number of proposals on evaluation metrics for this purpose are available, see [4, 5, 11]. We chose the *taxonomic precision (TP)*, and *taxonomic recall (TR)* [4]. The basic idea for these metrics is to find common concepts presented in both hierarchies, and then for each common concept to extract a set of concepts that characterises it, termed as *characteristic excerpt*. The similarity of hierarchies is then computed based on the similarity of these characteristic excerpts.

The common semantic cotopy was proposed to define the characteristic excerpt in [5]. However, it cannot judge the direction of hierarchical relations. Thus, we adopted the direct common subconcepts, by traversing from the common concept to all subconcepts until a subconcept is found, which is contained in both hierarchies. We restricted this search to one level only in the hierarchy. This is because that the study requires evaluating the rules' ability to capture

² <http://www.kde.cs.uni-kassel.de/bibsonomy/dumps/>.

³ <http://downloads.dbpedia.org/2015-10/core/>.

⁴ <https://concept.research.microsoft.com/Home/Download>.

⁵ <http://www.acm.org/about/class/class/2012>.

direct relations and the reference ontologies are generally sparse (especially for DBpedia and MCG). We named this new characteristic except as common direct subsumption (cdsub). The taxonomic precision is defined in Eqs. 1 and 2, where L is the learned hierarchy and G the referenced one. The taxonomic recall is computed as $TR(L, G) = TP(G, L)$. Similar to the standard F -measure, the taxonomic F -measure is the harmonic mean of taxonomic precision and recall $TF(L, G) = \frac{2 * TP(L, G) * TR(L, G)}{TP(L, G) + TR(L, G)}$. The metrics are defined for non-leaf nodes ensuring $|cdsub(c, L, G)|$ above zero; leaf concepts (have no direct subconcepts) in L or G have tp_{cdsub} as zero.

Another two metrics were also used: *taxonomic overlap* (TO) [11] measuring the quality of hierarchies and *taxonomic F' -measure* (TF' , note that it is different from TF) [2, 5] measuring the overall quality of the lexical and relation levels. TO is a symmetric measure and is calculated with Eqs. 3 and 4. As cdsub only concerns common concepts, TF and TO are not affected by non-common concepts in both hierarchies. Thus, to include lexical coverage into evaluation, *taxonomic F' -measure* (TF') is introduced, computed as the harmonic mean of Lexical Recall (LR) and TF , $TF'(L, G) = \frac{2 * LR(L, G) * TF(L, G)}{LR(L, G) + TF(L, G)}$. LR is defined as $LR(L, G) = \frac{|V_L \cap V_G|}{|V_G|}$, where $|V|$ is the size of vocabulary in a hierarchy.

$$tp_{cdsub}(c, L, G) = \frac{|cdsub(c, L, G) \cap cdsub(c, G, L)|}{|cdsub(c, L, G)|} \quad (1)$$

$$TP(L, G) = \frac{1}{|L \cap G|} \sum_{c \in L \cap G} tp_{cdsub}(c, L, G) \quad (2)$$

$$to_{cdsub}(c, L, G) = \frac{|cdsub(c, L, G) \cap cdsub(c, G, L)|}{|cdsub(c, L, G) \cup cdsub(c, G, L)|} \quad (3)$$

$$TO(L, G) = \frac{1}{|L \cap G|} \sum_{c \in L \cap G} to_{cdsub}(c, L, G) \quad (4)$$

Results and Discussion. From Fig. 1 it can be seen that in terms of the quality of hierarchies (measured by TF and TO , left column), with PTM representation, the set inclusion rule (set-inc) outperformed the others. With resource-based representation, the results were rather inconsistent. Probabilistic association rule (prob-asso) had comparable performance to set inclusion rule with DBpedia and MCG. The results obtained on CCS were more inconsistent compared to the other two due to the low overlapping with the learned hierarchy. For CCS, information-theoretical condition (info-theo) and fuzzy set inclusion (fuzzy-set-inc) with resource-based representation performed the best.

In terms of overall quality of the learned hierarchies (measured by TF'), due to the huge size of DBpedia, there is little difference among rules. Notable difference was found with MCG: using the PTM representation, the set inclusion rule showed the best TF' ; with the resource-based representation, prob-asso has the best TF' ; fuzzy-set-inc performed well with both representation techniques.




Fig. 1. Results of the reference-based evaluation. The figures illustrate the quality of each learned hierarchical structure based on its similarity to three gold-standard taxonomies (rows: DBpedia, Microsoft Concept Graph (MCG) and the ACM Computing Classification System ToC (CCS)), with different rules and data representation techniques (PTM and Resource-based). Three metrics were used for evaluation: in left column, *taxonomic F1-measure* (TF), *taxonomic overlap* (TO) that measure the quality of concept hierarchies; and in right column *taxonomic F'-measure* (TF') that measures the overall lexical and relational quality of the learned hierarchy.

Again, the results using CCS demonstrated considerable inconsistency, with the PTM representation, the set inclusion performed quite well, while graph centrality rule has the lowest TF' . With the resource-based representations, information theoretic condition and fuzzy set inclusion rule performed almost equally well, both slightly better than the set inclusion rule. The hierarchy learned using set inclusion rule showed high overall quality with both representations.

With regard to **Q1** and **Q2**, the set inclusion rule generated the best result with PTM representations, while the probabilistic association and fuzzy set inclusion rules with the resource-based representation generated competitive results. This finding is interesting as the simplest rule produced the best result. It is probably because the set theory provides the most precise simulation to the extension of a concept. The probabilistic association rule simulated the semantic association of concepts and achieved best results when the resource-based representation is used. In terms of the overall quality of hierarchy taking into consideration lexical coverage, results were considerably inconsistent. This can

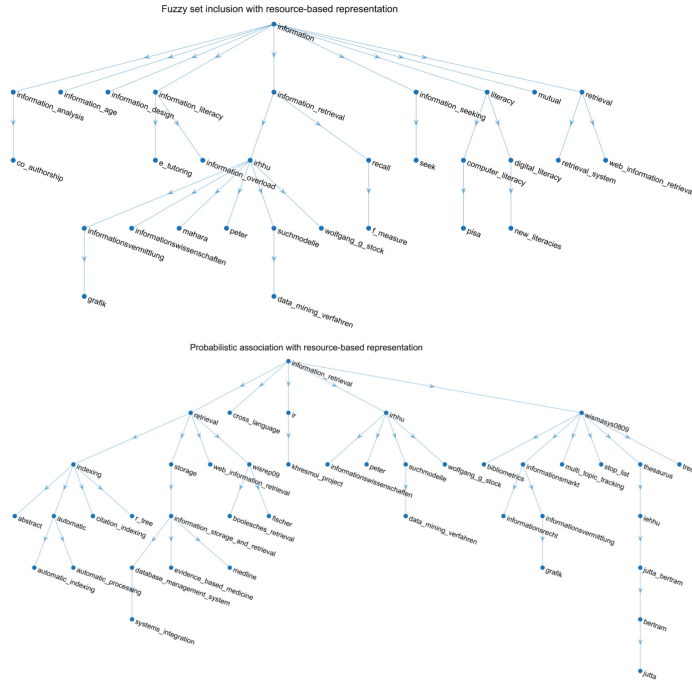


Fig. 2. Hierarchies learned with the two proposed rules. (We suggest to view the figures by zooming in on the digital version of this paper.)

be attributed to different nature of reference hierarchies and also demonstrated the usefulness of different rules for this task. There might be issues with the reference-based evaluation, as it only measures the global similarity between the learned hierarchy to the reference hierarchy. However, it is possible that some branches of the learned hierarchy were very similar to the reference hierarchy [16]. Two excerpts (due to the limited space) of the learned hierarchies are illustrated in Fig. 2 with fuzzy set inclusion and probabilistic association rule.

We also realised that the current evaluation strategies might not be perfect due to the significantly different nature of the datasets. In general, quality of the learned hierarchies from social data cannot be compared to those created by human beings or learned from sentential data using lexico-syntactic patterns. Nevertheless, useful and unseen concepts and relations can still be discovered.

5 Conclusion and Future Study

Hierarchical relations are essential to Knowledge Organisation Systems and research has shown these relations can be learned from folksonomies. This work summarised the representative rules used to induce such relations and conducted

an in-depth comparison on these rules. We believe that with a deeper understanding of the rules, it is possible to design methods to extract valuable information even from low-quality social data to enrich existing knowledge bases. The reference-based evaluation showed that with the proposed hierarchy generation algorithm, the set inclusion rule and probabilistic association rule can derive useful hierarchical knowledge with resource-based and PTM representations. For future studies, we plan to include more fine-grained rules and evaluate them with other folksonomies. We are also combining these rules with machine learning approaches to generate hierarchies with higher quality. Other representations, such as neural word embeddings, are also to be explored.

References

1. Benz, D., Hotho, A., Stumme, G., Stützer, S.: Semantics made by you and me: Self-emerging ontologies can capture the diversity of shared knowledge. In: Proceedings of the 2nd Web Science Conference (WebSci10) (2010)
2. Cimiano, P., Hotho, A., Staab, S.: Learning concept hierarchies from text corpora using formal concept analysis. *J. Artif. Intell. Res. (JAIR)* **24**(1), 305–339 (2005)
3. Cruse, D.A.: Hyponymy and its varieties. In: Green, R., Bean, C.A., Myaeng, S.H. (eds.) *The Semantics of Relationships: An Interdisciplinary Perspective*, pp. 3–21. Springer, Dordrecht (2002). https://doi.org/10.1007/978-94-017-0073-3_1
4. Dellschaft, K., Staab, S.: Measuring the similarity of concept hierarchies and its influence on the evaluation of learning procedures. Master's thesis (Diplomarbeit), University of Koblenz-Landau (2005)
5. Dellschaft, K., Staab, S.: On how to perform a gold standard based evaluation of ontology learning. In: Cruz, I., Decker, S., Allemang, D., Preist, C., Schwabe, D., Mika, P., Uschold, M., Aroyo, L.M. (eds.) *ISWC 2006. LNCS*, vol. 4273, pp. 228–241. Springer, Heidelberg (2006). https://doi.org/10.1007/11926078_17
6. Dong, H., Wang, W., Frans, C.: Deriving dynamic knowledge from academic social tagging data: a novel research direction. In: *iConference 2017 Proceedings* (2017)
7. García-Silva, A., Corcho, O., Alani, H., Gómez-Pérez, A.: Review of the state of the art: discovering and associating semantics to tags in folksonomies. *Knowl. Eng. Rev.* **27**(1), 57–85 (2012)
8. Griffiths, T.L., Steyvers, M.: Prediction and semantic association. In: *Proceedings of the 15th International Conference on Neural Information Processing Systems*, pp. 11–18. MIT Press (2002)
9. Heymann, P., Garcia-Molina, H.: Collaborative creation of communal hierarchical taxonomies in social tagging systems. Technical report, Stanford University (2006)
10. Jabeen, F., Khusro, S.: Quality-protected folksonomy maintenance approaches: a brief survey. *Knowl. Eng. Rev.* **30**(5), 521–544 (2015)
11. Maedche, A., Staab, S.: Measuring similarity between ontologies. In: Gómez-Pérez, A., Benjamins, V.R. (eds.) *EKAW 2002. LNCS (LNAI)*, vol. 2473, pp. 251–263. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45810-7_24
12. Meo, P.D., Quattrone, G., Ursino, D.: Exploitation of semantic relationships and hierarchical data structures to support a user in his annotation and browsing activities in folksonomies. *Inf. Syst.* **34**(6), 511–535 (2009)
13. Mika, P.: Ontologies are us: a unified model of social networks and semantics. *Web Semant.: Sci. Serv. Agents World Wide Web* **5**(1), 5–15 (2007)

14. Peters, I., Becker, P.: *Folksonomies: Indexing and Retrieval in Web 2.0*. De Gruyter/Saur, Berlin (2009)
15. Stock, W.G.: Concepts and semantic relations in information science. *J. Am. Soc. Inf. Sci. Technol.* **61**(10), 1951–1969 (2010)
16. Strohmaier, M., Helic, D., Benz, D., Körner, C., Kern, R.: Evaluation of folksonomy induction algorithms. *ACM Trans. Intell. Syst. Technol.* **3**(4), 1–22 (2012)
17. Tho, Q.T., Hui, S.C., Fong, A.C.M., Cao, T.H.: Automatic fuzzy ontology generation for semantic web. *IEEE Trans. Knowl. Data Eng.* **18**(6), 842–856 (2006)
18. Wang, W., Barnaghi, P.M., Bargiela, A.: Probabilistic topic models for learning terminological ontologies. *IEEE Trans. Knowl. Data Eng.* **22**(7), 1028–1040 (2010)
19. Weller, K.: *Knowledge Representation in the Social Semantic Web*. De Gruyter Saur, Berlin/New York (2010)