

Folktale similarity based on ontological abstraction

Marijn Schraagen

Digital Humanities Lab, Utrecht University, The Netherlands

<M.P.Schraagen@uu.nl>

Abstract

This paper presents a method to compute similarity of folktales based on conceptual overlap at various levels of abstraction as defined in Dutch WordNet. The method is applied on a corpus of Dutch folktales and evaluated using a comparison to traditional folktale similarity analysis based on the Aarne–Thompson–Uther (ATU) classification system. Document similarity computed by the presented method is in agreement with traditional analysis for a certain amount of folktale pairs, but differs for other pairs. However, it can be argued that the current approach computes an alternative, data-driven type of similarity. Using WordNet instead of a domain-specific ontology or classification system ensures applicability of the method outside of the folktale domain.

1 Introduction

A folktale is a specific type of narrative that is particularly suitable for analysis of semantic structure. Although folktales may differ in various aspects, such as the characteristics of the main actors or the sequence of events, often similarities can be identified on a more general or more abstract level. In this paper similarity between folktales is investigated using an explicit abstraction of text according to the WordNet concept hierarchy. A comparison is provided to conventional folktale motif analysis. An example of folktale similarity on various levels of abstraction is provided by the folktales *Sleeping Beauty* and *Snow White*, which both feature a princess as specific character and a variable number of enchanted objects at a more general level. Similarities regarding events occur at various levels as well, for example the princess in *Snow White* is asked by the seven dwarfs to perform household tasks, whereas the girl protagonist

from *Hansel & Gretel* is ordered by the witch to do housework. In this case both the actors and the actions are similar at various levels, depicted in Figure 1. This notion of abstraction-based semantic similarity can be computed automatically using a machine-readable concept hierarchy such as WordNet. The paper is structured as follows: Section 2 describes characteristics of folktales and provides an overview of the resources used in the analysis, Section 3 discusses related work, Section 4 provides details on similarity computation, Section 5 contains experimental results and a comparison to existing folktale analysis approaches, and Section 6 concludes.

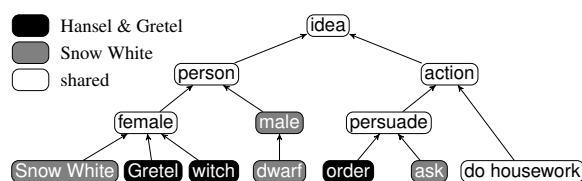


Figure 1: Example partial hierarchy showing concepts from *Snow White* and *Hansel & Gretel*.

2 Folktale similarity

The folktale texts used in the current research are extracted from the Dutch Folktale Database (Meder, 2010). This collection contains over 40,000 folktales (including jokes, urban legends, etc.) from written and oral sources, in Dutch, Frisian and several contemporary and historical Dutch dialects¹. The database is partially annotated using Aarne-Thompson-Uther (ATU) codes (see the remainder of this section for details). The database, maintained by the Meertens Institute, is available for research purposes upon request. For the current research a pilot set of 16 traditional fairy tales is used from a single original source (van Dongen and Grooten, 2009), with a total of

¹<http://www.verhalenbank.nl>, in Dutch

<i>id</i>	<i>description</i>	<i>category</i>	<i>ATU</i>	
F451.5.1.2	Dwarfs adopt girl as sister	Marvels → Spirits and demons → Underground Spirits	709	Snow White
Q2	Kind & Unkind	Rewards and punishments	440, 480, 513, 571, ...	Frog King, Kind & Unkind Girls, Wonderful Helpers, Golden Goose
F911.3	Animal swallows man (not fatally)	Marvels → Extraordinary occurrences → extraordinary swallowings	123, 333, 700	The Wolf & the Seven Kids, Little Red Riding Hood, Tom Thumb
F913	Victims rescued from swallower's belly	Marvels → Extraordinary occurrences → extraordinary swallowings	123, 333, ...	The Wolf & the Seven Kids, Little Red Riding Hood
J144	Well-trained kid does not open to wolf	The wise and the foolish → Wisdom acquisition → education	123	The Wolf & the Seven Kids
K1832	Disguise by changing voice	Deceptions → Deception by disguise	123	The Wolf & the Seven Kids

Table 1: TMI motif examples.

33,022 words. This set provides folktales in grammatically correct, modern Dutch, which increases applicability of natural language processing tools and methods. Several folktales in this set do not appear in the ATU catalog, illustrating the applicability of the current method on non-traditional folktale sources.

In the current research the Dutch Cornetto database is used (Vossen et al., 2013) to obtain term abstractions. Cornetto is modeled after the Princeton WordNet, which is a widely used ontology for English concepts (Fellbaum, 1998) containing a comprehensive set of terms and (hierarchical) relations for an extensive variety of domains. Concepts are organized in sets of (approximate) synonyms, called *synsets*, which are connected by relations such as hypernymy and meronymy. Cornetto contains over 92,000 lemmas and is available under academic license².

Traditionally, folktales are analyzed using the Thompson Motif Index (TMI). This index is a set of over 12,000 story elements (motifs), classified in semantic categories and subcategories (Thompson, 1960). Some examples are provided in Table 1. The motifs in this index are often specific to a single folktale (or *folktale type*, i.e., the set of variants of a story that are considered the same folktale), however more general motifs are used as well. The Aarne–Thompson–Uther (ATU) classification system (Uther, 2004) describes a folktale type as a list of motifs (typically two or three to about 20) from a subset of nearly 1,900 elements from the TMI, divided into thematic categories and subcategories. The ATU classification is centered around the type of protagonists and the general theme of the folktale, while the TMI is centered around events and relations. This may in-

²An open source database using the structure of Cornetto and translations of the content of English WordNet is available as an alternative (Postma and Vossen, 2014a).

roduce semantic relatedness differences between the two systems, for example the classification of ATU 123 as *Animal tales–Wild animals and domestic animals* compared to ATU 333 which is classified as *Fairy tale–supernatural opponent*, while two out of the total of four motifs of ATU 123 are also found in ATU 333 (see Table 1).

3 Related work

Folktale similarity using WordNet-based term matching has been previously investigated (McIntyre and Lapata, 2010; Lestari and Manurung, 2015) using the hierarchical similarity measure of Wu and Palmer (1994). In this approach a folktale is considered sequential, with similarity computation based on alignment of the sequence of actions and actors. In contrast, the current approach considers the (non-sequential) presence of terms and term abstractions, similar to a bag-of-words approach, while preserving event or situation similarity by comparing folktales on a sentence level. Abstraction based on Dutch WordNet for folktale similarity has been used by Nguyen et al. (2013), using abstractions of verbs as one of several features involved in similarity computation. The abstraction feature did not improve the results significantly, which is attributed to limited coverage of the abstraction lexicon and inaccuracy of the grammatical analysis.

Characterizing semantic relations between folktales using TMI motifs is discussed by Karsdorp et al. (2012), presenting the conclusion that motif-based methods suffer from the limited amount of motif overlap between folktales. A search tool for TMI motifs using WordNet based semantic abstraction is presented in (Karsdorp et al., 2015). A mapping of nominal phrases to folktale actors using a domain-specific ontology for term abstraction is described by Declerck et al. (2012).

An unsupervised exploration and visualization method for concept clustering in folktales has been proposed by Honkela (1997), using self-organizing maps trained on word trigrams. Natural computing approaches using (phylo)genetic algorithms are used to study variation within folktale types and between closely related types, using TMI motifs and other story elements as features (Ross et al., 2013; Tehrani, 2013). A vector-based method for semantic folktale clustering using Latent Semantic Mapping is described in (Vaz Lobo and Martins de Matos, 2010).

Several semantic relatedness measures that use WordNet as knowledge base have been proposed, see, e.g., (Pedersen et al., 2004) for an overview (although the use of WordNet as an ontological resource can be criticized (Guarino, 1998)). Considering hierarchy traversal, well-known approaches include the Wu-Palmer measure mentioned above, which defines similarity between two nodes as the path length from the first shared parent node to the root node of the hierarchy, and the Leacock-Chodorow measure, which finds the shortest path between two concepts (scaled for specificity of the hierarchy). Further graph-topological information is incorporated using PageRank (Agirre et al., 2009). Evaluation of graph-based semantic relatedness measures has been performed using comparison to human word-pair similarity ratings, e.g., (Postma and Vossen, 2014b). Recent approaches of similarity computation include path length weighting strategies (Gao et al., 2015) and domain-specific data (McInnes and Pedersen, 2015).

Using WordNet for similarity of documents has been investigated by, e.g., (Hotho et al., 2003; Sedding and Kazakov, 2004) for the task of document clustering. These approaches represent a document as a bag-of-words, consisting of terms in the document as well as term synonyms and hypernyms from WordNet. However, it is concluded that the investigated approach of adding WordNet relations does not improve clustering results significantly. Similar methods for document clustering do show improved results, e.g., (Wang and Taylor, 2007), suggesting a large impact of preprocessing and sense selection procedures. Further applications include information retrieval, matching a WordNet-expanded query to a set of (non-expanded) documents (Varelas et al., 2005).

Note that many approaches using WordNet for

semantic similarity focus either on pairs of concepts (or synsets, words, lemmas, etc.), document clusters, or, in the folktale setting, variants of the same story. These tasks are generally motivated by the availability of evaluation resources, such as human concept similarity ratings, the Reuters categorized news corpus, or folktale corpora tagged by story type, respectively. In contrast, the current approach attempts to construct a network of documents based on semantic relatedness, by comparing document pairs on (non-sequential) sentence level. Evaluation of this approach is arguably less straightforward, however this task and the proposed WordNet-based method provide a shift in focus compared to traditional approaches.

4 Method

In the current approach a document collection is compared at sentence level. First, sentence boundaries, lemmas and part-of-speech tags are obtained using the Frog toolkit (van den Bosch et al., 2007). Lemmas tagged as noun (including proper names), adjective, or verb are selected (except for the common verbs *be*, *have*, *can* and *will*). The set of lemmas for a sentence is compared to the set of lemmas for all other sentences in all other folktales in the corpus. If a matching lemma is not found in the compared sentence then the WordNet hierarchy is consulted for a match at a higher level of abstraction (i.e., the *lowest common subsumer*), using the match level to adjust the similarity score. The similarity of two sentences is computed as the total of all match scores relative to the combined size of the lemma sets. Formally, the score $s \in [0, 1]$ equals $(\sum_i \frac{1}{level(a_i)} + \sum_j \frac{1}{level(b_j)}) / (|A| + |B|)$ for sentences A and B as sequences of WordNet lemmas and $level(\ell)$ defined as the minimum level of the lemma ℓ that matches a lemma (at any level) in the compared sentence. After computing similarity scores for all sentence pairs, for each (ordered) folktale combination (f_A, f_B) the relative number of sentences in f_A is counted for which the most similar sentence in the corpus originated from f_B . The procedure is described formally in Algorithm 1, an example is provided in Figure 2. For the mapping of sentence lemmas to WordNet the synset with the lowest WordNet sense number is selected, corresponding to some extent to the ‘default’ sense. Incorrect senses are assumed to be related or to have a minimal effect given the document size (cf. (Hotho et al.,

on the correlation values. The overall values are somewhat lower than the correlation for the hierarchy traversal measures reported by Postma and Vossen (2014b) on Dutch WordNet, which might be caused by the lack of hierarchy depth awareness of the current method. However, the current measure is not intended as a stand-alone word pair similarity computation, instead it is part of an asymmetrical sentence matching procedure intentionally designed for matching on any level of the hierarchy.

5 Results and evaluation

Application of the current method on the folktale test corpus results in a matrix of pair-wise directed similarity scores, shown in Table 3. The graph of scores above a threshold of 10.0 (i.e., the most similar sentence for at least 10% of the sentences in folktale *A* was found in folktale *B*) is provided in Figure 3. The graph contains a number of central nodes, most notably *Snow White*, *Hansel and Gretel*, and *The Wonderful Helpers*. These nodes can be interpreted as representing a prototypical folktale, more specifically the fairy tale subgenre.

Increasing the similarity threshold to 13.0 (solid lines in Figure 3) reveals two clusters in the graph. The left cluster contains folktales featuring civilian protagonists, who find themselves in potentially harmful circumstances. The right cluster contains royal protagonists dealing with issues of moral values. The exception is *Snow White*, which has a royal protagonist, who is however banned from the royal court, living as a civilian house guest annex maid, and subject of murder attempts.

For comparison, the same method is applied without using WordNet abstractions, i.e., counting overlap in (lemmatized) terms as found in the text. This comparison (see Figure 4) shows that plain term overlap is less structured or partitioned in general and pair-wise relations display less topic overlap as compared to the abstraction method.

To provide an evaluation of the proposed similarity measure, a comparison is performed to the traditional ATU classification and associated TMI motif sets. To address the problem of limited motif overlap between folktales, the hierarchical TMI numbering system can be used for partial or abstract motif overlap using an approach similar to the term similarity computation described in Section 4 (see Table 4 for examples of motif matching). Motif overlap for a pair of folktales can

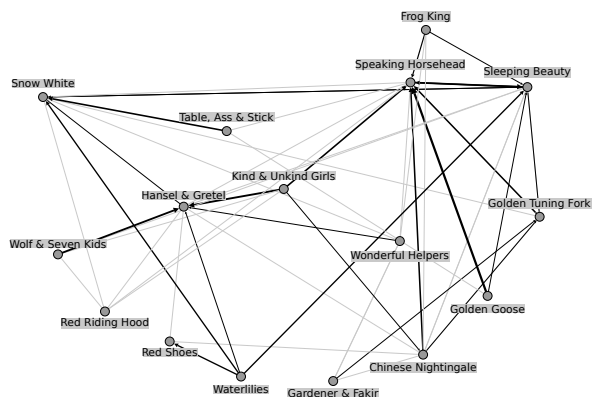


Figure 4: Plain term overlap scores for threshold 10.0 (grey edges) and 13.0 (black edges).

be aggregated in different ways, e.g., the number of overlapping motifs, the sum of the highest match levels for overlapping motifs, the number of matches at any level (in this case the motifs K1832 and K1839.1 from Table 4, for example, generate the four matches K, K1, K18, K183), or the sum of all match levels, each of which can optionally be weighted by the number of motifs assigned to the two folktales involved in the match. The directed score for a folktale pair (*A*, *B*) can be normalized using the rank of the similarity value of *A* and *B* as compared to other similarity values for *A*. Using this normalization on the values for the number of unique motif matches results in a clear separation of directionality, shown in Figure 5. Upon visual inspection, the TMI-based graph appears to confirm the central position of *Snow White* as folktale prototype. In contrast, for several other nodes the properties are inconsistent with similarity computed using WordNet. Considering the individual relations, 7 out of the 27 directed edges in Figure 5 (marked with *) are present in Figure 3 as well, three of which connected to *Snow White*.

5.1 Graph comparison

In order to provide graph-theoretical support for the visual correspondence claim, the differences in degree distribution for corresponding nodes can be quantified. For this analysis the *assortativity coefficient* (Newman, 2002; Piraveenan et al., 2008) is used, which considers the number of edges of a node and the direct neighbors of this node and compares these numbers to the overall degree distribution of the network. Formally, the assortativ-

ATU	title	123	327	333	410	440	480	513	533	563	571	709	GTF	NG	RS	GF	WL
123	Wolf & Seven Kids	—	8.45	17.61	4.93	7.04	4.23	3.52	7.75	5.63	9.15	<i>11.27</i>	2.11	2.82	6.34	2.11	2.11
327	Hansel & Gretel	3.04	—	9.57	3.91	6.09	5.65	<i>11.30</i>	8.70	4.35	6.52	14.78	2.17	6.96	9.13	2.61	2.17
333	Red Riding Hood	14.29	17.01	—	3.40	3.40	6.12	7.48	9.52	2.72	4.76	3.40	2.72	2.04	8.16	4.08	2.04
410	Sleeping Beauty	1.42	5.67	2.84	—	7.09	4.26	13.48	9.93	5.67	4.26	<i>12.06</i>	7.80	8.51	9.93	4.96	0
440	Frog King	4.03	8.05	3.36	<i>11.41</i>	—	5.37	14.09	14.09	4.03	4.03	<i>10.74</i>	4.70	8.05	1.34	3.36	0
480	Kind & Unkind Girls	5.15	22.06	5.88	2.94	4.41	—	6.62	<i>10.29</i>	8.09	6.62	5.88	1.47	5.88	<i>10.29</i>	2.21	0
513	Wonderful Helpers	1.45	<i>11.64</i>	6.18	9.45	8.00	3.64	—	<i>11.27</i>	4.36	5.09	8.36	3.64	9.45	5.45	2.91	1.09
533	Speaking Horsehead	1.83	<i>10.98</i>	1.83	7.93	13.41	4.27	16.46	—	4.27	6.71	8.54	4.88	6.10	6.71	3.05	1.22
563	Table, Ass & Stick	3.29	<i>11.84</i>	3.95	1.32	9.21	4.61	9.87	4.61	—	9.21	<i>10.53</i>	5.92	3.29	7.24	7.89	1.32
571	Golden Goose	3.68	<i>11.76</i>	5.15	5.15	8.09	2.21	14.71	8.09	5.88	—	8.82	5.88	5.15	5.88	5.15	3.68
709	Snow White	4.21	<i>12.30</i>	6.47	8.09	7.12	5.83	9.06	5.83	4.21	4.85	—	6.80	6.47	6.15	4.85	2.27
GTF	Golden Tuning Fork	0.80	7.20	4.80	<i>11.20</i>	2.40	4.00	<i>12.80</i>	16.00	1.60	6.40	8.80	—	8.00	6.40	7.20	0
NG	Nightingale	4.10	14.55	4.10	7.46	5.60	3.36	<i>11.19</i>	7.84	2.24	6.34	8.58	6.72	—	4.48	8.96	0.75
RS	Red Shoes	2.68	8.72	13.42	8.72	3.36	6.04	7.38	6.04	4.03	6.04	<i>12.08</i>	4.70	8.72	—	2.01	3.36
GF	Gardener & Fakir	2.79	<i>11.73</i>	3.91	6.15	4.47	5.59	<i>11.73</i>	7.82	5.03	2.79	8.38	6.15	14.53	4.47	—	2.79
WL	Waterlilies	2.08	6.25	8.33	<i>10.42</i>	2.08	6.25	<i>10.42</i>	6.25	0	6.25	14.58	0	8.33	<i>12.50</i>	4.17	—

Table 3: Pair-wise WordNet-based similarity scores for the test corpus. Thresholds indicated in italics (10.0) and bold (13.0).

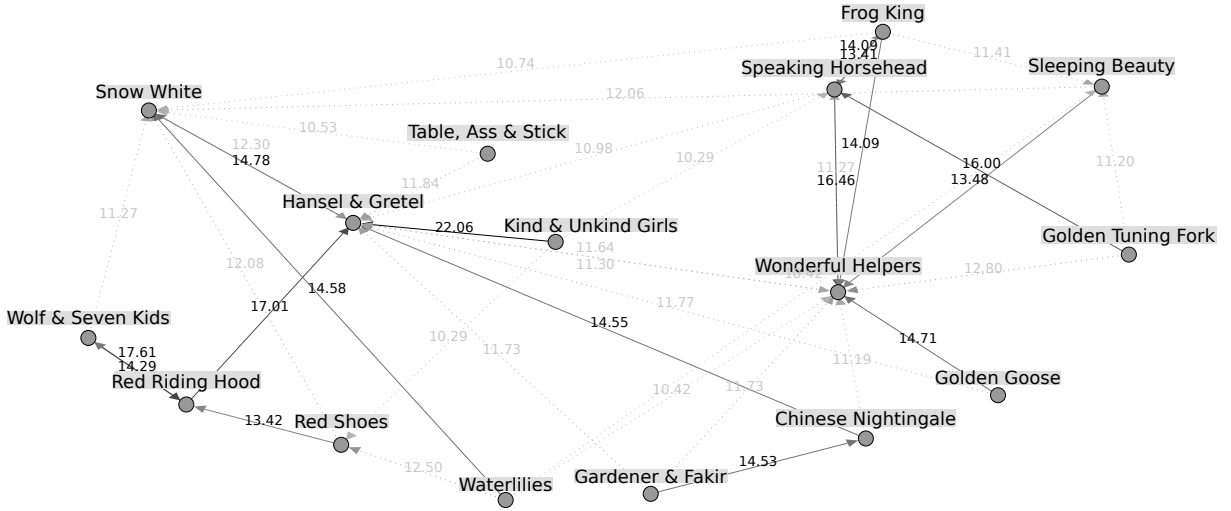


Figure 3: Graph of directed pairwise folktale similarity scores for threshold 10.0 (dotted edges) and 13.0 (solid edges).

ity coefficient ρ for a single node is defined as

$$\rho = \frac{j(j+1)(\bar{k} - \mu_q)}{2M\sigma_q^2} \quad (1)$$

In this equation j is the excess (or remaining) degree of the node, defined as the degree minus one, \bar{k} is the average excess degree of the neighbor nodes, μ_q is the mean of the excess degree distribution, M is the number of edges and σ_q^2 is the standard deviation of the excess degree distribu-

tion. The distribution parameters μ_q and σ_q^2 can be derived from the normalized excess degree distribution q which is based on the traditional degree distribution p as observed on the graph (i.e., $p(k)$ equals the number of nodes with degree k divided by the total number of nodes). The normalization for q is given by the following equation:

$$q(k) = \frac{(k+1)p(k+1)}{\sum_j jp(j)} \quad (2)$$

ATU	Title	Motif description	Motif code	match level
123	The Wolf & the Seven Kids	Disguise by changing voice	K1832	
333	Little Red Riding Hood	Wolf puts flour on his paw to disguise himself	K1839.1	4
533	The Speaking Horsehead	Disguise as goose-girl (turkey-girl)	K1816.5	3
533	The Speaking Horsehead	Imposter forces oath of secrecy	K1933	2
709	Snow White	Compassionate executioner: substituted heart	K0512.2	1

Table 4: Example motif matches for *The Wolf & the Seven Kids*.

Now, the distribution parameters are conventionally given as $\mu_q = \sum_j jq(j)$ and $\sigma_q^2 = \sum_j q(j) \cdot (j - \mu_q)^2$. From Equation (1) the behavior of the coefficient is apparent: the magnitude is increased for higher values of j (i.e., nodes with high degree), and the direction is dependent on the neighboring nodes, resulting in a positive value when the nodes have high degree (compared to the average) and a negative value when the nodes have a low degree. The denominator term scales the coefficient between -1 and 1.

In Figure 6 the assortativity values for both types of similarity computation are shown as a correlation graph. The figure shows correspondences for peripheral nodes and the central *Snow White* node, as well as differences for nodes which are central in one of the two graphs only. Note that assortativity is not a measure of centrality as such. The definition takes into account the difference in degree between neighboring nodes, i.e., a larger part of the network is measured as compared to single degree count.

6 Discussion and future work

The current WordNet-based similarity measure divides the example folktale set into two clusters corresponding to civilian protagonists in threatening circumstances and royal protagonists presented with moral choices, respectively. This result shows that the method is able to differentiate general topics in folktales based on overlap in terms and term abstractions. Using term abstraction increases the level of clustering. The

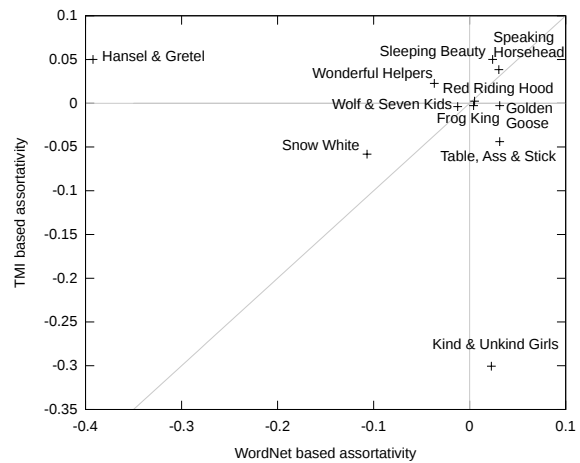


Figure 6: Correlation of assortativity of similarity graphs.

comparison with traditional folktale motif analysis shows corresponding similarity relations and centrality for a number of folktales, but deviating results for others. However, even though both analysis methods measure folktale similarity, the WordNet similarity measure considers the full text of a document, involving both syntax and semantics, while motif analysis is based on a small set of key events or themes, resulting in a highly specific semantic comparison on a considerably reduced and condensed representation of the document. The difference in approach leads to different results of similarity computation as well.

Rather than an alternative approach of computing TMI similarity, the WordNet method should be considered an alternative text-oriented measure of

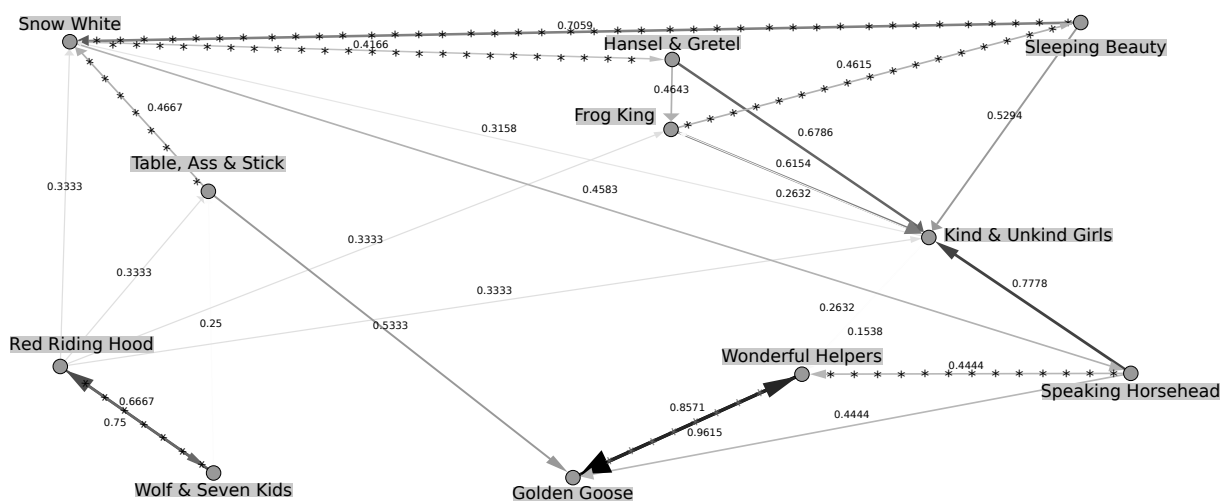


Figure 5: Graph of directed pairwise folktale similarity scores using the Thompson Motif Index. The number of motif pairs with the highest overlap is shown for pairs of documents (relative to the number of source document motifs), restricted to the top two highest ranked target nodes for each document.

similarity of folktales. The current approach has the advantage that a domain-specific classification system is no longer required. Within the folktale domain this addresses the issue of selective motif attribution and differences in motif granularity for folktales featured in the existing catalogs, as well as the possibility to include folktales outside of the catalog, as demonstrated in the current test corpus. This advantage extends to potential use outside of the folktale domain, e.g., using general literary works or non-fictional narratives.

The current method is ranking-based, therefore a strong match between two documents (e.g., two variants of the same narrative) may cause less pronounced similarities to remain undetected. This behaviour can be exploited for incremental clustering, by leaving out the comparison of highly similar document pairs in subsequent iterations.

In future work, the granularity of the WordNet hierarchy and the relative position in the concept tree can be used to adjust term matching weights. Word sense disambiguation can be taken into account. The method can be applied on larger or more heterogeneous corpora, e.g., folktale documents lacking standardized spelling or grammatical sentences could be used to test the robustness of knowledge-based approaches. The approach could be extended towards discourse analysis to accommodate story element matching across sentence boundaries. Scalability issues resulting from the current method of comparing every pair of sentences for every pair of documents could be addressed using various precomputing, pruning or selection mechanisms. Finally, the development of an informative baseline (e.g., using existing clustering toolkits) and an automatic evaluation procedure tailored towards the current notion of narrative similarity (e.g., using story variants as in (Nguyen et al., 2013)) is desired to increase understanding and interpretation of current results.

References

- Eneko Agirre, Enrique Alfonseca, Keith Hall, Jana Kravalova, Marius Paşca, and Aitor Soroa. 2009. A study on similarity and relatedness using distributional and WordNet-based approaches. In *Proceedings of the 2009 Annual Conference of the North American Chapter of the ACL (NAACL HLT)*, pages 19–27. ACL.
- Antal van den Bosch, Bertjan Busser, Sander Canisius, and Walter Daelemans. 2007. An efficient memory-based morphosyntactic tagger and parser for Dutch. In *Selected Papers of the 17th Computational Linguistics in the Netherlands Meeting*, pages 99–114. LOT Utrecht.
- Thierry Declerck, Nikolina Koleva, and Hans-Ulrich Krieger. 2012. Ontology-based incremental annotation of characters in folktales. In *Proceedings of the 6th EACL Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH 2012)*, pages 30–34. ACL.
- Gerrie van Dongen and Ad Grooten. 2009. *Sprookjesboek van De Efteling*. Ploegsma. (in Dutch).
- Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.
- Jian-Bo Gao, Bao-Wen Zhang, and Xiao-Hua Chen. 2015. A WordNet-based semantic similarity measurement combining edge-counting and information content theory. *Engineering Applications of Artificial Intelligence*, 29:80–88.
- Nicola Guarino. 1998. Some ontological principles for designing upper level lexical resources. In *Proceedings of LREC 1998*, pages 527–534. ELRA.
- Timo Honkela. 1997. Self-organizing maps of words for natural language processing applications. In *Proceedings of the International ICSC Symposium of Sof Computing*, pages 401–407. ICSC.
- Andreas Hotho, Steffen Staab, and Gerd Stumme. 2003. Wordnet improves text document clustering. In *Proceedings of the Semantic Web Workshop at SIGIR-2003*. ACM.
- Folger Karsdorp, Peter van Kranenburg, Theo Meder, Dolf Trieschnigg, and Antal van den Bosch. 2012. In search of an appropriate abstraction level for motif annotations. In *Proceedings of the 2012 Computational Models of Narrative Workshop*, pages 22–26.
- Folger Karsdorp, Marten van der Meulen, Theo Meder, and Antal van den Bosch. 2015. MOMFER: A search engine of Thompson’s motif-index of folk literature. *Folklore*, 126(1):37–52.
- Victoria Lestari and Ruli Manurung. 2015. Measuring the structural and conceptual similarity of folktales using plot graphs. In *Proceedings of the 9th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH 2015)*. ACL.
- Bridget McInnes and Ted Pedersen. 2015. Evaluating semantic similarity and relatedness over the semantic grouping of clinical term pairs. *Journal of Biomedical Informatics*, 54:329–336.
- Neil McIntyre and Mirella Lapata. 2010. Plot induction and evolutionary search for story generation. In *Proceedings of the 48th annual meeting of the Association for Computational Linguistics (ACL 2010)*, pages 1562–1572. ACL.

- Theo Meder. 2010. From a Dutch folktale database towards an international folktale database. *Fabula*, 51:6–22.
- Mark Newman. 2002. Assortative mixing in networks. *Physical Review Letters*, 89(20).
- Dong Nguyen, Dolf Trieschnigg, and Mariët Theune. 2013. Folktale classification using learning to rank. In *Proceedings of the 35th European Conference on Information Retrieval (ECIR 2013)*, pages 195–206. Springer.
- Ted Pedersen, Siddharth Patwardhan, and Jason Michellizzi. 2004. WordNet::similarity – measuring the relatedness of concepts. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI-04)*, pages 1024–1025. ACM.
- Mahendra Piraveenan, Mikhail Prokopenko, and Albert Zomaya. 2008. Local assortativeness in scale-free networks. *Europhysics Letters*, 84.
- Marten Postma and Piek Vossen. 2014a. Open source Dutch WordNet. Technical report, Free University of Amsterdam.
- Marten Postma and Piek Vossen. 2014b. What implementation and translation teach us: the case of semantic similarity measures in wordnets. In *Proceedings of the 7th Global WordNet Conference (GWC 2014)*, pages 133–142.
- Robert Ross, Simon Greenhill, and Quentin Atkinson. 2013. Population structure and cultural geography of a folktale in Europe. *Proceedings of the Royal Society B*, 280(20123065).
- Julian Sedding and Dimitar Kazakov. 2004. WordNet-based text document clustering. In *Proceedings of the 3rd Workshop on Robust Methods in Analysis of Natural Language Data*, pages 104–113. ACL.
- Jamshid Tehrani. 2013. The phylogeny of Little Red Riding Hood. *PLoS One*, 8(11).
- Stith Thompson. 1960. *Motif-index of folk-literature: a classification of narrative elements in folktales, ballads, myths, fables, mediaeval romances, exempla, fabliaux, jest-books and local legends*. Indiana University Press.
- Hans-Jörg Uther. 2004. *The Types of International Folktales: A Classification and Bibliography Based on the System of Antti Aarne and Stith Thompson*. Finnish Academy of Science and Letters.
- Giannis Varelas, Epimenidis Voutsakis, Paraskevi Raftopoulou, Euripides Petrakis, and Evangelos Milios. 2005. Semantic similarity methods in WordNet and their application to information retrieval on the web. In *Proceedings of the 7th ACM International Workshop on Web Information and Data Management (WIDM 2005)*, pages 10–16. ACM.
- Paula Vaz Lobo and David Martins de Matos. 2010. Fairy tale corpus organization using latent semantic mapping and an item-to-item top-n recommendation algorithm. In *Proceedings of LREC 2010*, pages 1472–1475. ELRA.
- Piek Vossen, Isa Maks, Roxane Segers, Hennie van der Vliet, Marie-Francine Moens, Katja Hofman, Erik Tjong Kim Sang, and Maarten de Rijke. 2013. Cornetto: A combinatorial lexical semantic database for Dutch. In Peter Spyns and Jan Odijk, editors, *Essential Speech and Language Technology for Dutch: Results by the Stevin programme*, pages 165–184. Springer.
- James Wang and William Taylor. 2007. Concept forest: A new ontology-assisted text document similarity measurement method. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*, pages 395–401. IEEE.
- Zhibiao Wu and Martha Palmer. 1994. Verb semantics and lexical selection. In *Proceedings of the 32nd annual meeting of the Association for Computational Linguistics (ACL '94)*, pages 133–138. ACL.