Using a distributional neighbourhood graph to enrich semantic frames in the field of the environment

Gabriel Bernier-Colborne **Marie-Claude L'Homme**
Observatoire de linguistique Sens-Texte (OLST)
Université de Montréal
C.P. 6128, succ. Centre-Ville
Montréal (QC) Canada, H3C 3J7
{gabriel.bernier-colborne|mc.lhomme}@umontreal.ca

Abstract

This paper presents a semi-automatic method for identifying terms that evoke semantic frames (Fillmore, 1982). The method is tested as a means of identifying lexical units that can be added to existing frames or to new, related frames, using a large corpus on the environment. It is hypothesized that a method based on distributional semantics, which exploits the assumption that words that appear in similar contexts have similar meanings, can help unveil lexical units that evoke the same frame or related frames. The method employs a distributional neighbourhood graph, in which each word is connected to its nearest neighbours according to a distributional semantic model. Results show that most lexical units identified using this method can in fact be assigned to frames related to the field of the environment.

1 Introduction[1]

Recent work has shown that Frame Semantics (Fillmore, 1982; Fillmore and Baker, 2010) is an extremely useful framework to account for the lexical structure of specialized fields of knowledge (Dolbey et al., 2006; Faber et al., 2006; Schmidt, 2009; L'Homme et al., 2014). It is especially attractive in terminology since it provides an apparatus to connect linguistic properties of terms to a more abstract conceptual representation level.

[1]The work reported in this paper is carried out within a larger project entitled "Understanding the environment linguistically and textually", whose objective is to develop methods for characterizing the contents of texts on two different levels: 1. textual (using methods and techniques derived from corpus linguistics and text mining); and 2. linguistic (based on lexical semantic models).

Frame Semantics has proved especially useful to represent predicative units (verbs such as *deforest*, *recycle*, *warm*; predicative nouns such as *impact*, *pollution*, *salinization*; adjectives such as *clean*, *green*, *sustainable*), units that are often ignored in terminological resources. L'Homme et al. (2014) showed that the framework and more specifically the methodology devised within the FrameNet Project (Ruppenhofer et al., 2010) could be used to represent various lexico-semantic properties of predicative terms (in English and in French). L'Homme and Robichaud (2014) showed that frames could be connected via a series of relations and contribute to help us understand how terms are used to express environmental knowledge. However, as will be seen below, the work that led to the definition of frames and relations between frames mentioned above was done manually and turns out to be quite time-consuming. In this paper, we explore the potential of a semi-automatic, graph-based method to discover frame-relevant lexical units based on corpus evidence.

This paper is structured as follows. Section 2 explains how semantic frames help reveal part of the lexical structure of a specialized field of knowledge. Section 3 describes the graph-based method used to identify frame-relevant lexical units. Section 4 discusses how the model used in the manual evaluation of this method was selected. Section 5 presents the evaluation methodology and the results of the evaluation.

2 Frame Semantics applied to the field of the environment

In a specialized field such as the environment, many concepts correspond to processes, events and properties which are typically expressed linguistically by predicative terms (verbs, predicative nouns and adjectives). However, traditional terminological models (and even less traditional ones, such as ontologies) are not properly equipped to describe the terms that denote these concepts and account for their specific linguistic properties, namely the fact that they require arguments (*X changes Y*; *impact of X on Y*). Frame Semantics (Fillmore, 1982; Fillmore and Baker, 2010) presents itself as a suitable alternative to these models since it is designed to connect linguistic properties to an abstract conceptual structure. In addition, it is well equipped to represent predicative lexical units and their argument structure.

2.1 Discovering frames in the field of the environment

L'Homme et al. (2014) describe a method to discover semantic frames based on an existing terminological resource called DiCoEnviro[2], that contains English and French terms related to the field of the environment. Each entry in DiCoEnviro is devoted to a lexical unit (LU), i.e. a lexical item that conveys a specific meaning, and states the argument structure of the LU, as in the following examples:

- warm_{1a}, vi: $\text{climate}_{[\text{Patient}]}$ warms
- warm_{1b}, vt: $\text{gas}_{[\text{Agent}]}$ or $\text{change}_{[\text{Cause}]}$ warms $\text{climate}_{[\text{Patient}]}$
- warm, adj.: warm $\text{climate}_{[\text{Patient}]}$

Argument structures state the number of obligatory participants, and two different systems are used to label them: the first one accounts for the semantic roles of arguments (**Agent**, **Patient**, **Cause**); the second one gives a typical term, i.e. a term that is representative of what can appear in that position.

Many entries – especially entries that describe predicative terms – come with annotated contexts that show how arguments[3] are realized in sentences extracted from an environmental corpus. For example, annotated contexts for $warm_{1b}$ are shown in Table 1.

The primary radiative effect of CO2 and water vapour$_{[CAUSE]}$ *is to* ***WARM*** *the surface climate*$_{[PATIENT]}$ *but cool the stratosphere.*
As increases in other greenhouse gases$_{[CAUSE]}$ ***WARM*** *the atmosphere and surface*$_{[PATIENT]}$*, the amount of water vapour also increases, amplifying the initial warming effect of the other greenhouse gases.*
The simulations of this assessment report (for example, Figure 5) indicate that the estimated net effect of these perturbations$_{[CAUSE]}$ *is to* ***HAVE WARMED*** *the global climate*$_{[PATIENT]}$ *since 1750*$_{[TIME]}$.

Table 1: Annotated contexts for $warm_{1b}$

Argument structures and annotations were used to discover frames using two different methods. A semantic frame is a knowledge structure that represents specific situations (e.g. a teaching situation, a selling situation, a driving situation). A frame includes participants (called frame elements or FEs), some of which are obligatory (core FEs) and some of which are optional (non-core FEs). For instance, the Operate_vehicle frame describes a situation in which a Vehicle is set in motion by a Driver and includes the following core FEs: **Area**, **Driver**, **Goal**, **Path**, **Source**, and **Vehicle**. Lexical units such as *cycle*, *cruise*, *drive*, *pedal*, and *ride* evoke this frame (FrameNet, 2015). In this previous work, it was assumed that terms that share similarities with regard to their argument structures (number and semantic roles of arguments) and that share similarities with regard to the non-obligatory participants annotated in contexts are likely to evoke the same frame.

The first method consisted in comparing the argument structures and non-obligatory participants of terms already encoded in the terminological resource. This method shows that the verbs $cool_{1a}$, $warm_{1a}$ and the nouns $cooling_1$ and $warming_1$ share many features. They all have a single argument (a **Patient**) and share some non-obligatory participants (**Degree**, **Duration**, **Location**).

The second method – which was applied only to the English terms – consisted in comparing the contents of the terminological resource to that of

[2] See http://olst.ling.umontreal.ca/cgi-bin/dicoenviro/search_enviro.cgi.

[3] Non-obligatory participants are also annotated, as shown in the last sentence in Table 1, in which the phrase *since 1750*, which expresses Time, is annotated.

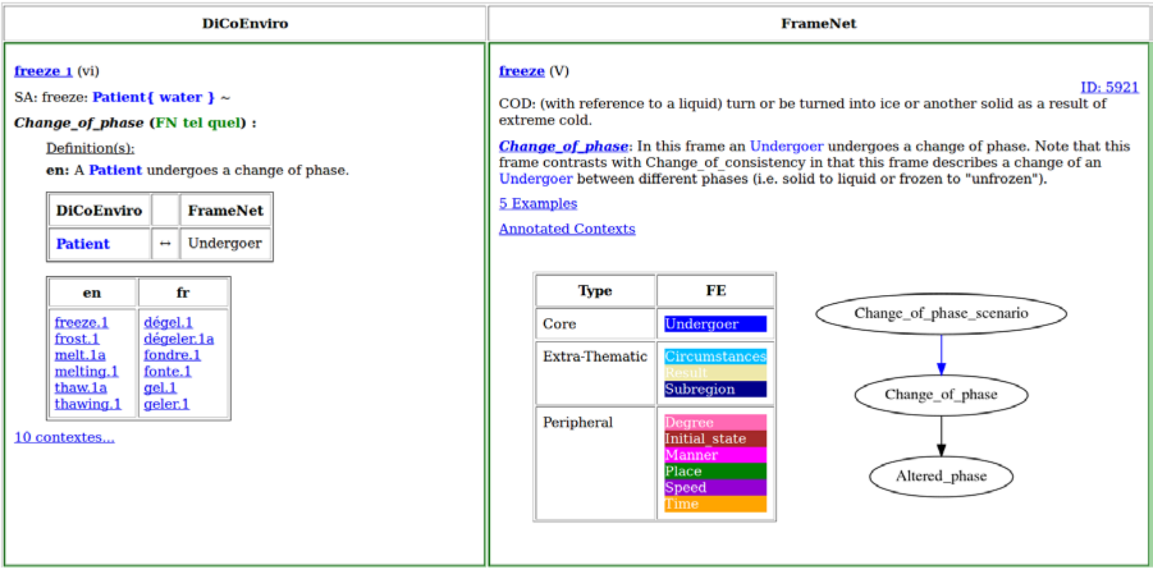

Figure 1: Comparison between the terminological database and FrameNet

FrameNet.[4] Relevant data were extracted from the FrameNet database for terms that were recorded in the terminological resource, as shown in Figure 1. This figure shows an example in which a correspondence between FrameNet and the terminological database could be established. However, in many instances, matches could not be made as nicely. In various cases, specific frames needed to be defined for the environmental terms (for instance, a new frame was created to capture adjectives such as *clean*, *environmental* and *green*, whose meaning can be loosely described as "that does not harm the environment"). In other cases, existing frames in FrameNet needed to be adapted to the data extracted from the terminological database for different reasons (slightly more specific meanings, different number of arguments, etc.).

2.2 A "framed" representation of the terminology of the environment

It soon became obvious that some of the frames identified based on the methods described in Section 2.1 could be linked. For instance, all processes related to changes affecting the environment appeared to be somehow related.

Again, using FrameNet (2015) as a reference, relations were established between some of the frames defined for environmental terms. Two relations not found in FrameNet (2015) were added (*Is opposed to* and *Is a property of*). This work led to the development of a resource called the Framed DiCoEnviro,[5] in which users can navigate through frames and relations between frames, and access the terms that evoke these frames along with their annotations. Figure 2 shows some of the relations identified between the frame Change_of_temperature (COT) (that contains verbs such as $cool_{1a}$, $warm_{1a}$ and the nouns $cooling_{1}$ and $warming_{1}$) and other frames.

3 Method for discovering related LUs

The methods described above allowed us to define a first subset of frames that are relevant for the field of the environment, link part of these frames and assign lexical units (LUs) to them. Based on this preliminary data, we explored the potential of a semi-automatic method to enrich our resource by adding new LUs to existing frames or discovering new frames. This method exploited distributional information obtained from a much larger corpus than the one used in the work described above.

The method we tested to discover related LUs is based on the neighbourhood graph induced by a distributional model of semantics. Distributional semantic models are commonly used to estimate semantic similarity, the underlying hypothesis be-

[4] The FrameNet team releases an XML version of the database (Baker and Hong, 2010).

[5] See `http://olst.ling.umontreal.ca/dicoenviro/framed/index.php` (in development).

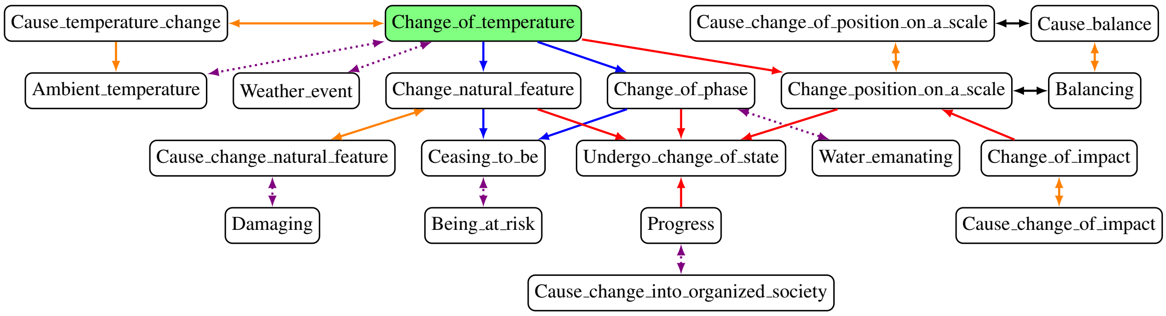

Figure 2: Change_of_temperature and related frames

ing that words that appear in similar contexts tend to be semantically related (Harris, 1954). The usual method of querying a distributional model is simply to compute, given a particular word, a sorted list of similar words. This method has several drawbacks, as has been pointed out recently by Gyllensten and Sahlgren (2015), who use a relative neighbourhood graph to query distributional models in a way that accounts for the fact that the query can have multiple senses. The method used here is similar in that it exploits a distributional neighbourhood graph. This allows us to take a list of terms and visualize their semantic neighbourhood, in order to identify related terms that can be encoded as frame-evoking LUs, either in existing frames or in new ones.

Various kinds of graphs could be used to compute and visualize the distributional neighbourhood of a particular word or set of words. We use a k-nearest-neighbour (k-NN) graph, two examples of which are the symmetric k-NN graph and the mutual k-NN graph (Maier et al., 2007). In a symmetric k-NN graph, two words w_i and w_j are connected if w_i is among the k nearest neighbours (NNs) of w_j or if w_j is among the k NNs of w_i. In a mutual k-NN graph, the two words are connected only if both conditions are true: w_i is among the k NNs of w_j and w_j is among the k NNs of w_i. In this work, we chose to use a mutual k-NN graph[6], the intuition behind this decision being that if two words are mutual NNs, there is a better chance that they actually do have similar meanings. This principle has been exploited elsewhere (Ferret, 2012; Claveau et al., 2014).

The graph construction procedure can be summarized as follows. Given a distributional semantic model, we compute the pairwise similarity between all words. For each word, we compute its k NNs by sorting all other words in decreasing order of similarity to that word and keeping the k most similar. Then, for each word w_i and each neighbour w_j in the k NNs of w_i, we add an edge in the graph between w_i and w_j if w_i is also among the k NNs of w_j. The resulting graph can be used to visualize the distributional neighbourhood of a term or set of terms.

4 Model selection

Any model that allows us to estimate the semantic similarity of two words can be used to build a semantic neighbourhood graph such as the one described in Section 3. We tested two different distributional semantic models for this purpose. Both models have several parameters which must be set and which can have a significant impact on the accuracy of the model in a given application. We therefore used an automatic evaluation procedure to tune the models' parameters and select a model for manual evaluation.

4.1 Corpus and reference data

The corpus used to build the models is the PANACEA Environment English monolingual corpus (Catalog Reference ELRA-W0063), a corpus containing 28071 web pages related to the environment (approximately 50 million tokens). The corpus was compiled automatically using a focused web crawler developed within the PANACEA project, and is freely distributed by ELDA for research purposes.[7] The corpus was

[6] We also tested the symmetric k-NN graph, but we only report results obtained with the mutual graph. We achieved higher F-scores using the mutual graph.

[7] See http://catalog.elra.info/product_info.php?products_id=1184.

converted from XML to raw text and lemmatized using TreeTagger (Schmid, 1994).

Reference data were extracted from the Framed DiCoEnviro.[8] The reference data are sets of LUs that evoke the same semantic frame. The list of English LUs was extracted from each of the frames included in the Change_of_temperature (COT) scenario[9] (cf. Figure 2). Two LUs (*thawing* and *thinning*) were excluded because they were not in the vocabulary used to construct the models, which contains the 10,000 most frequent lemmatized words in the corpus, excluding stop words. We obtained 13 sets containing a total of 53 LUs, each frame containing between 2 and 7 LUs. The number of unique LUs is 45, several LUs evoking more than one frame.[10]

4.2 Models tested

Two different distributional semantic models were tested. The first is a bag-of-words (BOW) model (Schütze, 1992; Lund et al., 1995), which is based on a word-word cooccurrence matrix computed using a sliding context window. The second is word2vec (Mikolov et al., 2013a; Mikolov et al., 2013b), a neural language model that has been used in many NLP applications in the past few years. Word2vec (W2V) learns distributed word representations that can be used in the same way as BOW vectors to estimate semantic similarity.

The models' parameters were tuned by testing various combinations of parameter values, building neighbourhood graphs from each resulting model, and computing evaluation metrics on these graphs based on the reference data described in Section 4.1.

Some of the main choices that must be made when training a model using word2vec pertain to the architecture of the model (continuous skip-gram or continuous bag-of-words), the training algorithm (hierarchical softmax or negative sampling), the use of subsampling of frequent words, the size (dimensionality) of the word vectors and the size of the context window. We tested various values for each of these parameters, including the recommended values[11] when available. A total of 160 models were tested. In the case of the BOW model, important parameters[12] include the type, shape and size of the context window, the weighting scheme applied to the cooccurrence frequencies, and the use of dimensionality reduction. Again, we tested different values for these parameters. Each model was tested with and without dimensionality reduction, for which we used singular value decomposition (SVD). A total of 320 BOW models were built and evaluated (160 unreduced and 160 reduced using SVD).

For both models, we used the cosine similarity to estimate the similarity between words.

4.3 Evaluation metrics for model selection

For each model tested, we constructed multiple k-NN graphs, using different values of k. For each of these graphs, we computed evaluation metrics using the reference data described in Section 4.1. We used precision and recall to check to what extent LUs belonging to the same frame were connected in the graph. These metrics are computed for each of the 45 unique LUs in the reference data. Let w_i be an LU, $R(w_i)$ the set of related LUs that evoke at least one of the frames evoked by w_i, and $NN(w_i)$ the set of words that are adjacent to w_i in the graph. Furthermore, let TP_i (true positives) be the number of words in $NN(w_i)$ that are one of the related LUs in $R(w_i)$, FP_i (false positives) the number of words in $NN(w_i)$ that are not in $R(w_i)$ and FN_i (false negatives) the number of words in $R(w_i)$ that are not in $NN(w_i)$. The evaluation metrics are then calculated as usual:

$$\text{precision}_i = \frac{TP_i}{TP_i + FP_i}$$

$$\text{recall}_i = \frac{TP_i}{TP_i + FN_i}$$

$$\text{F-score}_i = \frac{2 \times \text{precision}_i \times \text{recall}_i}{\text{precision}_i + \text{recall}_i}$$

[8] Data extracted on 2015-05-22. Data has been added since then, as the resource is in development.

[9] Frames related to the scenario only through a *See also* relation were excluded.

[10] Polysemous LUs evoke different frames. For instance, warm$_{1a}$ (intransitive verb) evokes the Change_of_temperature frame; warm$_{1b}$ (transitive verb) evokes the Cause_temperature_change frame; and warm$_{2}$ (adjective) evokes the Ambient_temperature frame.

[11] See `https://code.google.com/p/word2vec/#Performance`.

[12] Several studies have assessed the influence of this model's parameters. The relative importance of several parameters was quantified using analysis of variance by Lapesa et al. (2014).

The average precision, recall and F-score for a particular graph are then computed by taking the mean scores over all LUs in the reference data.

4.4 Results

Table 2 shows how precision, recall and F-score vary with respect to k. As the density of the graph increases, recall increases and precision decreases, the average F-score peaking around $k = 10$. The table also shows the number of nodes in the subgraph corresponding to the 45 LUs and their adjacent nodes in the graph. Based on these results, we selected 10 as an appropriate value of k.

k	nb nodes	precision	recall	F1
5	125	0.2120	0.2125	0.1915
10	206	0.1681	0.3005	**0.1971**
15	284	0.1429	0.3560	0.1858
20	359	0.1253	0.3999	0.1730
25	431	0.1108	0.4339	0.1594

Table 2: Evaluation metrics and number of nodes in the subgraph wrt k (averaged over all models)

Table 3 shows the average and maximum scores of each model (BOW, BOW reduced using SVD, and W2V) with $k = 10$. These results suggest that the BOW model performs best for this application.

Model	Avg prec. (max)	Avg rec. (max)	Avg F1 (max)
BOW	**0.1960** (**0.2775**)	**0.3153** (**0.4268**)	**0.2184** (**0.3016**)
SVD	0.1567 (0.2007)	0.2987 (0.3830)	0.1903 (0.2412)
W2V	0.1517 (0.2245)	0.2875 (0.4206)	0.1826 (0.2727)

Table 3: Evaluation metrics wrt model (with $k = 10$)

By analyzing how precision and recall varied with respect to the BOW model's parameters, we determined the optimal parameter values for this application. For example, the optimal window size was determined to be 3 words. The corresponding graph was then evaluated manually.

5 Evaluation

Once the model had been selected, the corresponding neighbourhood graph was evaluated manually. The evaluation was carried out by one of the co-authors of this paper, who is responsible for the development of the Framed DiCoEnviro. The 45 unique LUs in the reference data had 137 unique neighbours (adjacent nodes in the graph). These 137 words were evaluated manually in order to determine to what extent the graph can serve to discover frame-evoking LUs that can be added to the database.

The evaluation was carried out one frame at a time by observing the subgraph corresponding to that frame's LUs and their neighbours (adjacent nodes in the neighbourhood graph). For example, the subgraph for the frame Cause_change_of_impact is shown in Figure 3. In each subgraph, the LUs already encoded in that frame were highlighted in green, and those encoded in other frames in the COT scenario were highlighted in blue. One or more numbers were appended to the label of each LU to indicate which frame(s) it evokes.

For each word that was not already encoded as an LU in the COT scenario (i.e. for each white node), the evaluator was asked to choose one of the following categories:

1. The word should be encoded as an LU in the COT scenario
   (a) in an existing frame;
   (b) in a new frame.
2. The word should be encoded as an LU in another scenario
   (a) in an existing frame that is related to the COT scenario (by a *See also* relation);
   (b) in an existing frame that is not related to the COT scenario;
   (c) in a new frame.
3. The word should not be encoded as an LU in the database, but it is the realization of a core FE of one of the frames in the COT scenario.
4. The word should not be encoded as an LU in the database, nor is it the realization of a core FE of one of the frames in the COT scenario.

Table 4 shows the results of this evaluation. As these results show, most lexical items identified by the method (105 out of 137) can be encoded in a relevant frame in the field of the environment and

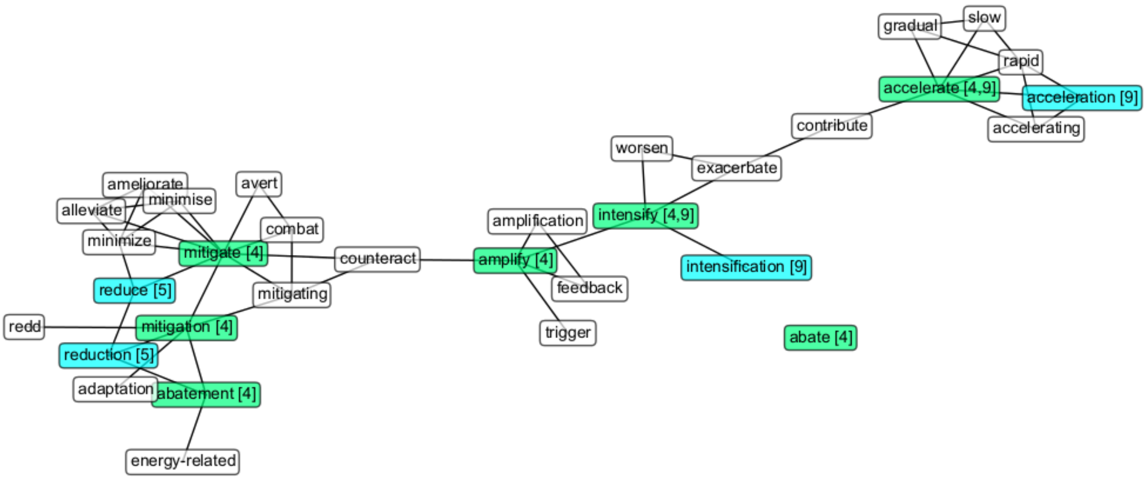

Figure 3: Subgraph for the frame Cause_change_of_impact

should be added to our resource. Among these, 88 would be frame-evoking LUs (categories 1 and 2) and 17 would be encoded as realizations of FEs (category 3). Interestingly, 48 lexical items are related to the COT scenario (categories 1a and 1b). The method allowed us to identify: 1. new frame-evoking LUs (such as *amplification*, *drop*, and *scarcity*) that had not been encoded in existing frames (category 1a); 2. LUs (such as *alteration* and *eliminate*) that evoke frames that had not been created (category 1b); and 3. variants (such as *cooler* for *cool* and *stabilise* for *stabilize*).

Category	Nb cases
1a	39
1b	9
1 (total)	**48**
2a	7
2b	3
2c	30
2 (total)	**40**
3	**17**
4	**32**
Total	**137**

Table 4: Summary of results.

The method also identified 40 items that would be encoded in environmentally relevant frames, but in a different scenario (category 2). It is worth pointing out that among these, 7 items correspond to LUs that would evoke a frame that is linked to the COT scenario (category 2a).

Finally, although 32 lexical items identified by the method would not be encoded in the resource and are thus considered false positives from the point of view of our application, further explanations are required. Some lexical items could evoke more general frames. For instance, *rapid* and *slow* would appear in the same frame if the general lexicon were considered. Other items identified are acronyms. *GW*, for instance, is the acronym for *global warming*. Technically, it could be defined as an LU evoking the COT frame, but multi-word terms and acronyms are not considered in the resource.

6 Concluding remarks

All in all the results obtained are quite interesting and show that the method can be used to assist lexicographers when defining frames and their lexical content, as the distributional neighbourhood of frame-evoking LUs often contain LUs that evoke the same frame or related frames. Distributional neighbourhood graphs provide information about the content of a specialized corpus that would be impossible to extract manually from such a large corpus. They are a very useful complement to other corpus tools, such as term extractors and concordancers, as they help lexicographers save time and locate relevant lexical units (near synonyms, variants) that they would otherwise miss.

In future work, we plan to integrate this methodology to assist lexicographers when defining new frames related to the field of the environment. It could be particularly useful to obtain a view on corpora that deal with new or more specific topics and unveil the lexical units used to convey the knowledge related to these topics. It would also be interesting to test the potential of the method in other fields of knowledge. Extensions of this work could also involve using a graph-based clustering method to discover sets of lexical units that evoke the same frame without using existing frames.

Acknowledgments

This work was supported by the Social Sciences and Humanities Research Council (SSHRC) of Canada.

References

Collin Baker and Jisup Hong. 2010. Release 1.5 of the FrameNet data. International Computer Science Institute. Berkeley.

Vincent Claveau, Ewa Kijak, and Olivier Ferret. 2014. Explorer le graphe de voisinage pour améliorer les thésaurus distributionnels. In *Actes de la 21e conférence sur le traitement automatique des langues naturelles (TALN)*, p. 220–231.

Andrew Dolbey, Michael Ellsworth, and Jan Scheffczyk. 2006. BioFrameNet: A Domain-specific FrameNet Extension with Links to Biomedical Ontologies. *Proceedings of KR-MED 2006: Biomedical Ontology in Action*, Baltimore, Maryland.

Pamela Faber *et al.* 2006. Process-oriented terminology management in the domain of Coastal Engineering. *Terminology* 12(2): 189–213.

Olivier Ferret. 2012. Combining bootstrapping and feature selection for improving a distributional thesaurus. In *Proceeding of the 20th European Conference on Artificial Intelligence (ECAI)*, p. 336–341.

Charles J. Fillmore. 1982. Frame Semantics. In *Linguistics in the Morning Calm*, p. 111–137. Seoul: Hanshin Publishing Co.

Charles J. Fillmore and Collin Baker. 2010. A frames approach to semantic analysis. In B. Heine and H. Narrog (ed.), *The Oxford Handbook of Linguistic Analysis*, p. 313–339. Oxford: Oxford University Press.

FrameNet. 2015. `https://framenet.icsi.berkeley.edu/fndrupal/`. Accessed: 2015-09-24.

Amaru Cuba Gyllensten and Magnus Sahlgren. 2015. Navigating the semantic horizon using relative neighborhood graphs. *CoRR*, abs/1501.02670.

Zellig S. Harris. 1954. Distributional structure. *Word*, 10(2–3): 146–162.

Gabriella Lapesa, Stefan Evert, and Sabine Schulte im Walde. 2014. Contrasting syntagmatic and paradigmatic relations: Insights from distributional semantic models. In *Proceedings of the Third Joint Conference on Lexical and Computational Semantics (*SEM 2014)*, p. 160–170.

Marie-Claude L'Homme and Benoît Robichaud. 2014. Frames and terminology: Representing predicative units in the field of the environment. In Proceedings of the 4th Workshop on Cognitive Aspects of the Lexicon (CogALex-IV), p. 186–197.

Marie-Claude L'Homme, Benoît Robichaud, and Carlos Subirats. 2014. Discovering frames in specialized domains. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, p. 1364–1371.

Kevin Lund, Curt Burgess, and Ruth Ann Atchley. 1995. Semantic and associative priming in high-dimensional semantic space. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, p. 660–665.

Markus Maier, Matthias Hein, and Ulrike Von Luxburg. 2007. Cluster identification in nearest-neighbor graphs. In *Algorithmic Learning Theory*, p. 196–210. Springer.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. In *Proceedings of ICLR*.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Proceedings of NIPS*, p. 3111–3119.

Josef Ruppenhofer, Michael Ellsworth, Miriam R. L. Petruck, Christopher R. Johnson, and Jan Scheffczyk. 2010. FrameNet II: Extended theory and practice. `http://framenet2.icsi.berkeley.edu/docs/r1.5/book.pdf`. Accessed: 2015-09-24.

Helmut Schmid. 1994. Probabilistic part-of-speech tagging using decision trees. In *Proceedings of the International Conference on New Methods in Language Processing*.

Thomas Schmidt. 2009. The Kicktionary: A Multilingual Lexical Resource of Football Language. In H.C. Boas (ed.), *Multilingual FrameNets in Computational Lexicography. Methods and Applications*, p. 101–134. Berlin/NewYork: Mouton de Gruyter.

Hinrich Schütze. 1992. Dimensions of meaning. In *Proceedings of the 1992 ACM/IEEE Conference on Supercomputing (Supercomputing'92)*, p. 787–796.