A Predictive Approach to Semantic Change Modelling

Mohamed Boukhaled, Benjamin Fagard, Thierry Poibeau

A Predictive Approach to Semantic Change Modelling

Mohamed Amine Boukhaled (amine.boukhaled@ens.fr), LATTICE (CNRS, ENS/PSL, U. Sorbonne nouvelle/USPC, France)

Benjamin Fagard (benjamin.fagard@ens.fr), LATTICE (CNRS, ENS/PSL, U. Sorbonne nouvelle/USPC, France)

Thierry Poibeau (thierry.poibeau@ens.fr), LATTICE (CNRS, ENS/PSL, U. Sorbonne nouvelle/USPC, France)

Author information – this field cannot be edited

Main body – enter your text and images here:

1   Introduction

Although it is well known that word meaning evolves over time, the cause and the pace of change is still largely unknown. In this context, computational modelling can shed new light on the problem by considering at the same time a large number of variables that are supposed to interact in the process. This field has already given birth to a large number of publications ranging from early work involving statistical and mathematical formalism (Bailey, 1973 ; Kroch, 1989) to more recent work involving robotics and large-scale simulations (Steels, 2011).

We consider that semantic change includes all kinds of change in the meanings of lexical items happening over the years. For example, the word awful has dramatically changed in meaning, moving away from a rather positive perspective equivalent to impressive or majestic at the beginning of the nineteenth century to a negative one equivalent to disgusting and messy nowadays (Wijaya and Yeniterzi, 2011).

In this work, we address the question of semantic change from a computational point of view. Our aim is to capture the systemic change of words meanings in an empirical model that is also predictive, contrary to most previous approaches that try to model and account for past data. We will first describe our methodology, then the experiment and our results.

2   Proposed methodology

Our goal is to train a model representing semantic change over a certain period and, from that, predict future semantic changes. The evaluation can thus be based on the observation of the gap between actual data and predicted data. Our model is based on two main components:

1- Diachronic word embedding representing the meaning of words over time-periods, following (Turney and Pantel, 2010). Word embeddings are known to effectively represent the meaning of words thanks to their surrounding contexts. The representation can be extended to consider a diachronic perspective (word embeddings are first trained for each time-period and then aligned temporally, so as to be able to track semantic change over time, see Fig. 1). For our study, we used the pre-trained diachronic word embeddings released by Hamilton et al. (2016): for each decade from 1800 to 1990, a specific word embedding is built using the word2vec skipgram algorithm. The training corpus used to produce these word embeddings was derived from the English Google Books N-gram datasets (Lin et al., 2012), which contain large amounts of historical texts in many languages (we used 5-grams with no part-of-speech tags). Each word in the corpus appearing from 1800 to 1999 is represented by a set of twenty continuous 300-dimensional vectors; one vector for each decade.
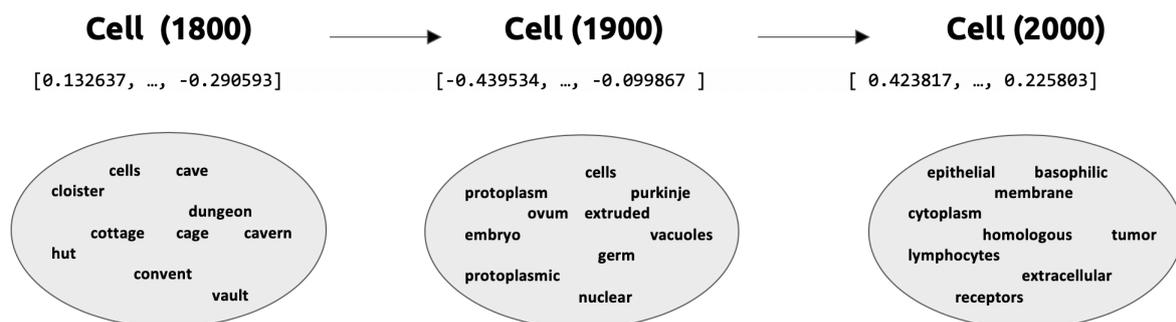

Figure 1. Two-dimensional visualization of the semantic change in the English word cell using diachronic word embedding. In the early 19th century cell referred to "cage" or "dungeon", whereas in the late 20th century its meaning shifted toward a scientific perspective.

2- Recurrent Neural Networks (RNNs) modelling the semantic change itself. RNNs are known to be powerful at recognizing dynamic temporal behaviour in diachronic data

(Medsker and Jain, 2001). In this experiment, we used the word embeddings representing the semantic space of each decade from 1800 to 1990 as input for the RNN, and from this we predicted the embedding corresponding to the 1990-1999 decade. Our RNN have a long short-term memory (LSTM) and are implemented through Tensorflow.

To explore different scenarios, we did several experiments with different vocabulary sizes (1000, 5000, 10000, 20000 and 50000 most frequent words). We used the stratified 10-fold cross-validation method to estimate the prediction error (i.e. 90% of the words were used for training, and 10% for testing). The overall prediction accuracy is taken as the average performance over these 10 runs.

3   Experiment, Results and Discussion

To get an overall estimation of the prediction accuracy, we compare each predicted embedding to the ground truth obtained from real data. Though it is impossible to predict exactly the exact vector corresponding to any word "w" as we are working in a continuous 300-dimensional space, one can assess the accuracy of the predicted meaning by extracting the closest vectors, i.e. the closest neighbours of a given word over time.

If the word "w" is actually the nearest semantic neighbour to the predicted vector then it is considered to be a correct prediction. Otherwise, it is considered to be an error (a false prediction). The results are summarized in Table 1.

Vocabulary Size	Accuracy
1000	91.7%
5000	86.1%
10000	71.4%
20000	52.2%
50000	25%

Table 1. Results of prediction accuracy measured for different vocabulary sizes. The training and the prediction using the RNNs model were performed on embedding derived from Google N-gram corpus.

The results show that the model can be highly effective at capturing semantic change and can achieve high accuracies when predicting the evolution of word meaning through distributional semantics. As one can see from Table 1, the model was able to achieve 71.4% accuracy when trained and tested exclusively on embeddings coming from the 10000 most frequent words of the corpus. The model was also able to correctly predict word embeddings for words that have radically changed their meaning over time such as awful, nice, call and record (Wijaya and Yeniterzi, 2011). The results also show better results when using smaller vocabulary sizes containing top frequent words. The decrease of performance with large vocabularies is simply due to the fact that infrequent words do not have enough occurrences to derive meaningful and stable enough contexts so as to observe reliable evolutions. It is thus fundamental to use large corpora for this kind of experiments, but also to adapt the size of the vocabulary to the size of the corpus.

Conclusion

We have proposed a new computational model of semantic change. Although this model is partially successful at representing this evolution, it can still appear as too simple compared to the complexity of language change in general and semantic change in particular. It is not obvious to see how this type of computational modelling can be combined with more traditional methods of linguistic analysis. However, we strongly believe that such empirical approaches based on diachronic vector-based representations can considerably help to refine and clarify theoretical insights on the foundations and mechanisms of semantic change, as well as provide an accurate empirical evaluation.

Acknowledgements

References

Enter your references here:

Bailey, C.-J.N. (1973). Variation and linguistic theory. Arlington: Centre for Applied Linguistics.

Hamilton, W.L., Leskovec, J. and Jurafsky, D. (2016). Diachronic word embeddings reveal statistical laws of semantic change. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 1489–1501.

Kroch, A.S. (1989). Reflexes of grammar in patterns of language change. *Language Variation and Change*, **1**: 199–244.

Lin, Y., Michel, J.-B., Aiden, E.L., Orwant, J., Brockman, W., Petrov, S. (2012). Syntactic annotations for the google books Ngram corpus. *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*: 169–174.

Medsker, L.R. and Jain, L.C. (2001). Recurrent neural networks. Boca Raton: CRC Press.

Steels, L. (2011). Modeling the cultural evolution of language. *Physical Life Review*, **8**: 339–356.

Turney, P.D. and Pantel, P. (2010). From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, **37**: 141–188.

Wijaya, D. and Yeniterzi, R. (2011). Understanding Semantic Change of Words Over Centuries. In the *Workshop on Detecting and Exploiting Cultural Diversity on the Social Web* (DETECT 2011) during CIKM 2011.