

# Predicting and interpreting embeddings for out of vocabulary words in downstream tasks

Nicolas Garneau\*, Jean-Samuel Leboeuf\*, Luc Lamontagne

Département d’informatique et de génie logiciel

Université Laval, Québec, Canada

{nicolas.garneau, jean-samuel.leboeuf}.1@ulaval.ca,  
luc.lamongtagne@ift.ulaval.ca

## Abstract

We propose a novel way to handle out of vocabulary (OOV) words in downstream natural language processing (NLP) tasks. We implement a network that predicts useful embeddings for OOV words based on their morphology and on the context in which they appear. Our model also incorporates an attention mechanism indicating the focus allocated to the left context words, the right context words or the word’s characters, hence making the prediction more interpretable. The model is a “drop-in” module that is jointly trained with the downstream task’s neural network, thus producing embeddings specialized for the task at hand. When the task is mostly syntactical, we observe that our model aims most of its attention on surface form characters. On the other hand, for tasks more semantical, the network allocates more attention to the surrounding words. In all our tests, the module helps the network to achieve better performances in comparison to the use of simple random embeddings.

## 1 Introduction and motivation

Goldberg (2017) emphasizes the fact that out of vocabulary (OOV) words represent a problem often underestimated for NLP tasks such as part of speech tagging (POS) or named entity recognition (NER) (Collobert et al., 2011; Turian et al., 2010). Due to the lack of proper ways to handle OOV words, researchers often resort to simply assign random embeddings to unknown words or to map them to a unique “unknown” embedding, hoping their model will generalize well nonetheless.

An interesting way to handle OOV words is the Mimick model (Pinter et al., 2017). This model aims to predict embeddings such as GloVe (Pennington et al., 2014) for OOV words by training a recurrent network on the characters of the words.

---

\*Authors contributed equally to this work.

While being simple, this model improves the accuracy of POS tagging as well as morphosyntactic attribute tagging on the Universal Dependencies corpus (De Marneffe et al., 2014).

We propose an extension to this model by taking into account not only the surface form of a word (i.e. its characters) but also the embeddings of its surrounding words. We hypothesize that context words provide useful semantic and syntactic information to model unknown word embeddings, hence complementing cues given by its characters. For this purpose, we introduce a module that can make, for the same word in different contexts, different predictions. It can also learn “specialized” embeddings for a specific downstream task which we evaluate for two sequence labeling tasks. Furthermore, we add to our model an attention/interpretation mechanism to determine which of the left context, right context or the surface form of a word receives more attention during prediction. Our experimental results are depicted in a quantitative and qualitative analysis.

## 2 Architecture

To test our ideas, we developed an OOV prediction module comprising the following components. First, the left context, right context and word characters are fed to three bi-LSTMs to produce separate encodings. These three hidden states are then passed to a linear layer on which a softmax is applied to determine their relative importance (i.e. their degree of attention). The output of this layer is then used to produce a weighted sum of the hidden states. Finally, a simple layer computes an embedding from this sum.

To evaluate the contribution of this OOV prediction scheme to sequence labeling tasks, we use a bi-LSTM architecture on the resulting word embeddings and apply a softmax on the hidden state of each word to predict tags.

Task	Metric	Random Emb.	Our module
NER	F1	77.56	<b>80.62</b>
POS	acc.	91.41	<b>92.58</b>

Table 1: Comparison of our OOV embeddings prediction scheme against random embeddings.

Task	Tag	Ex	Word	Left	Right
NER	O	1039	<b>0.81</b>	0.08	0.11
	B-PERS	63	0.21	0.31	<b>0.49</b>
	I-PER	119	0.16	<b>0.52</b>	0.32
	B-ORG	40	0.26	0.30	<b>0.44</b>
	I-ORG	3	0.27	0.31	<b>0.42</b>
	B-LOC	13	0.23	0.30	<b>0.47</b>
	I-LOC	2	0.16	<b>0.48</b>	0.36
	B-MISC	47	<b>0.40</b>	0.21	0.39
POS	I-MISC	5	<b>0.41</b>	0.26	0.33
	NNP	308	0.29	0.31	<b>0.40</b>
	NN	46	<b>0.45</b>	0.20	0.35
	CD	827	<b>0.86</b>	0.05	0.09
	NNS	23	<b>0.37</b>	0.24	<b>0.39</b>
	JJ	100	<b>0.49</b>	0.15	0.36

Table 2: Average weights assigned to word characters, left context and right context by the attention mechanism for NER and for POS tagging.

### 3 Experimental results and discussion

We evaluate the performance gain that our module can offer by solving two sequence labeling tasks, NER and POS tagging, using the *CoNLL 2003 shared task* dataset. We compare our module to a baseline where OOV words are assigned random embeddings. Table 1 shows the results we obtain. We can observe the clear advantage of proper handling of OOV words can provide. For both tasks, we gain a significant margin on the baseline, with

more than 3% of the F1 score for NER.

We can see from Table 2 that the network focuses more on the context for a semantic task such as NER. An interesting phenomenon is a focus on the right context when the entity is of type *B* and on the left context when the entity is of type *I*. We can also note that for the syntactic task (POS), the network tends to focus on the context for proper nouns (NNP), which corroborates our observations for the NER task. However, morphology plays a more important role to predict embeddings for other lexical categories. Embeddings for quantities (CD) are mostly predicted from their numerical characters.

We further qualitatively analyze the behavior of the network for a given OOV word appearing in different contexts in Table 3. When the target OOV word *langmore* is preceded by *john* or *australian*, the network gives high importance to these context words. However, an interesting phenomenon happens when a sentence begins with this word: the network shifts its attention from the left context to the right one and also assigns more importance to the morphology of the word, thus showing the network has truly learned where it can extract useful information.

### 4 Future works

In our future works, we plan to apply the attention mechanism specifically on the characters of the OOV word and the words that compose the context instead of using the hidden state of the respective elements only. We are also looking forward to testing our attention model in different languages and on other NLP tasks such as machine translation. We hope to present the full results and the architecture of our model in more details in a paper to be published relatively soon.

Word	Left	Right	Examples
0.24	<b>0.38</b>	<b>0.38</b>	<BOS> <i>langmore</i> , a persistent campaigner for interventionist economic
0.15	<b>0.59</b>	0.26	<BOS> <b>australian parliamentarian john</b> <i>langmore</i> has formally resigned from his lower house
0.15	<b>0.61</b>	0.24	<b>had received today from mr john vance</b> <i>langmore</i> , a letter resigning his place as
0.15	<b>0.69</b>	0.16	<BOS> <b>rtrs - australian mp john</b> <i>langmore</i> formally resigns . <EOS>
0.17	<b>0.40</b>	<b>0.43</b>	<BOS> <i>langmore</i> , <b>57</b> , announced in november that

Table 3: Qualitative example on the OOV word *langmore* which is an entity of type PER. We can clearly see that depending on the context, the weights may shift drastically.

## References

- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.
- Marie-Catherine De Marneffe, Timothy Dozat, Natalia Silveira, Katri Haverinen, Filip Ginter, Joakim Nivre, and Christopher D Manning. 2014. Universal stanford dependencies: A cross-linguistic typology. In *LREC*, volume 14, pages 4585–4592.
- Yoav Goldberg. 2017. Neural network methods for natural language processing. volume 10, chapter 8, pages 83–97.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Yuval Pinter, Robert Guthrie, and Jacob Eisenstein. 2017. Mimicking word embeddings using subword rnns. *arXiv preprint arXiv:1707.06961*.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pages 384–394. Association for Computational Linguistics.