

Incorporating Subword Information into Matrix Factorization Word Embeddings

Alexandre Salle Aline Villavicencio

Institute of Informatics

Universidade Federal do Rio Grande do Sul

Porto Alegre, Brazil

alex@alexsalle.com avillavicencio@inf.ufrgs.br

Abstract

The positive effect of adding subword information to word embeddings has been demonstrated for predictive models. In this paper we investigate whether similar benefits can also be derived from incorporating subwords into counting models. We evaluate the impact of different types of subwords (n-grams and unsupervised morphemes), with results confirming the importance of subword information in learning representations of rare and out-of-vocabulary words.

1 Introduction

Low dimensional word representations (embeddings) have become a key component in modern NLP systems for language modeling, parsing, sentiment classification, and many others. These embeddings are usually derived by employing the distributional hypothesis: that similar words appear in similar contexts (Harris, 1954).

The models that perform the word embedding can be divided into two classes: predictive, which learn a target or context word distribution, and counting, which use a raw, weighted, or factored word-context co-occurrence matrix (Baroni et al., 2014). The most well-known predictive model, which has become eponymous with word embedding, is word2vec (Mikolov et al., 2013a). Popular counting models include PPMI-SVD (Levy et al., 2014), GloVe (Pennington et al., 2014), and LexVec (Salle et al., 2016b).

These models all learn word-level representations, which presents two main problems: 1) Learned information is not explicitly shared among the representations as each word has an independent vector. 2) There is no clear way to represent out-of-vocabulary (OOV) words.

fastText (Bojanowski et al., 2017) addresses these issues in the Skip-gram word2vec model by

representing a word by the sum of a unique vector and a set of shared character n-grams (from hereon simply referred to as n-grams) vectors. This addresses both issues above as learned information is shared through the n-gram vectors and from these OOV word representations can be constructed.

In this paper we propose incorporating subword information into counting models using a strategy similar to fastText. We use LexVec as the counting model as it generally outperforms PPMI-SVD and GloVe on intrinsic and extrinsic evaluations (Salle et al., 2016a; Cer et al., 2017; Wohlgenannt et al., 2017; Konkol et al., 2017), but the method proposed here should transfer to GloVe unchanged.

The LexVec objective is modified¹ such that a word's vector is the sum of all its subword vectors. We compare 1) the use of n-gram subwords, like fastText, and 2) unsupervised morphemes identified using Morfessor (Virpioja, 2013) to learn whether more linguistically motivated subwords offer any advantage over simple n-grams.

To evaluate the impact subword information has on in-vocabulary (IV) word representations, we run intrinsic evaluations consisting of word similarity and word analogy tasks. The incorporation of subword information results in similar gains (and losses) to that of fastText over Skip-gram. Whereas incorporating n-gram subwords tends to capture more syntactic information, unsupervised morphemes better preserve semantics while also improving syntactic results. Given that intrinsic performance can correlate poorly with performance on downstream tasks (Tsvetkov et al., 2015), we also conduct evaluation using the VecEval suite of tasks (Nayak et al., 2016), in which all subword models, including fastText, show no significant improvement over word-level models.

We verify the model's ability to represent

¹Our implementation of subword LexVec is available at <https://github.com/alexandres/lexvec>

OOV words by quantitatively evaluating nearest-neighbors. Results show that, like fastText, both LexVec n-gram and (to a lesser degree) unsupervised morpheme models give coherent answers.

This paper discusses related word (§2), introduces the subword LexVec model (§3), describes experiments (§4), analyzes results (§5), and concludes with ideas for future works (§6).

2 Related Work

Word embeddings that leverage subword information were first introduced by Schütze (1993) which represented a word of as the sum of four-gram vectors obtained running an SVD of a four-gram to four-gram co-occurrence matrix. Our model differs by learning the subword vectors and resulting representation jointly as weighted factorization of a word-context co-occurrence matrix is performed.

There are many models that use character-level subword information to form word representations (Ling et al., 2015; Cao and Rei, 2016; Kim et al., 2016; Wieting et al., 2016; Verwimp et al., 2017), as well as fastText (the model on which we base our work). Closely related are models that use morphological segmentation in learning word representations (Luong et al., 2013; Botha and Blunsom, 2014; Qiu et al., 2014; Mitchell and Steedman, 2015; Cotterell and Schütze, 2015; Bhatia et al., 2016). Our model also uses n-grams and morphological segmentation, but it performs explicit matrix factorization to learn subword and word representations, unlike these related models which mostly use neural networks.

Finally, Cotterell et al. (2016) and Vúlic et al. (2017) retrofit morphological information onto pre-trained models. These differ from our work in that we incorporate morphological information at training time, and that only Cotterell et al. (2016) is able to generate embeddings for OOV words.

3 Subword LexVec

The LexVec (Salle et al., 2016a) model factorizes the PPMI-weighted word-context co-occurrence matrix using stochastic gradient descent.

$$PPMI_{wc} = \max(0, \log \frac{M_{wc} M_{**}}{M_{w*} M_{*c}}) \quad (1)$$

where M is the word-context co-occurrence matrix constructed by sliding a window of fixed size

centered over every *target* word w in the *subsampling* (Mikolov et al., 2013a) training corpus and incrementing cell M_{wc} for every *context* word c appearing within this window (forming a (w, c) pair). LexVec adjusts the PPMI matrix using *context distribution smoothing* (Levy et al., 2014).

With the PPMI matrix calculated, the sliding window process is repeated and the following loss functions are minimized for every observed (w, c) pair and target word w :

$$L_{wc} = \frac{1}{2} (u_w^\top v_c - PPMI_{wc})^2 \quad (2)$$

$$L_w = \frac{1}{2} \sum_{i=1}^k \mathbf{E}_{c_i \sim P_n(c)} (u_w^\top v_{c_i} - PPMI_{wc_i})^2 \quad (3)$$

where u_w and v_c are d -dimensional word and context vectors. The second loss function describes how, for each target word, k *negative samples* (Mikolov et al., 2013a) are drawn from the smoothed context unigram distribution.

Given a set of subwords S_w for a word w , we follow fastText and replace u_w in eqs. (2) and (3) by u'_w such that:

$$u'_w = \frac{1}{|S_w| + 1} (u_w + \sum_{s \in S_w} q_{hash(s)}) \quad (4)$$

such that a word is the sum of its word vector and its d -dimensional subword vectors q_x . The number of possible subwords is very large so the function $hash(s)$ ² hashes a subword to the interval $[1, buckets]$. For OOV words,

$$u'_w = \frac{1}{|S_w|} \sum_{s \in S_w} q_{hash(s)} \quad (5)$$

We compare two types of subwords: simple n-grams (like fastText) and unsupervised morphemes. For example, given the word “cat”, we mark beginning and end with angled brackets and use all n-grams of length 3 to 6 as subwords, yielding $S_{cat} = \{\langle ca, at \rangle, cat\}$. Morfessor (Virtioja, 2013) is used to probabilistically segment words into morphemes. The Morfessor model is trained using raw text so it is entirely unsupervised. For the word “subsequent”, we get $S_{subsequent} = \{\langle sub, sequent \rangle\}$.

²<http://www.isthe.com/chongo/tech/comp/fnv/>

4 Materials

Our experiments aim to measure if the incorporation of subword information into LexVec results in similar improvements as observed in moving from Skip-gram to fastText, and whether unsupervised morphemes offer any advantage over n-grams. For IV words, we perform intrinsic evaluation via word similarity and word analogy tasks, as well as downstream tasks. OOV word representation is tested through qualitative nearest-neighbor analysis.

All models are trained using a 2015 dump of Wikipedia, lowercased and using only alphanumeric characters. Vocabulary is limited to words that appear at least 100 times for a total of 303517 words. Morfessor is trained on this vocabulary list.

We train the standard LexVec (LV), LexVec using n-grams (LV-N), and LexVec using unsupervised morphemes (LV-M) using the same hyper-parameters as [Salle et al. \(2016a\)](#) (window = 2, initial learning rate = .025, subsampling = 10^{-5} , negative samples = 5, context distribution smoothing = .75, positional contexts = True).

Both Skip-gram (SG) and fastText (FT) are trained using the reference implementation³ of fastText with the hyper-parameters given by [Bojanowski et al. \(2017\)](#) (window = 5, initial learning rate = .025, subsampling = 10^{-4} , negative samples = 5).

All five models are run for 5 iterations over the training corpus and generate 300 dimensional word representations. LV-N, LV-M, and FT use 2000000 buckets when hashing subwords.

For word similarity evaluations, we use the WordSim-353 Similarity (WS-Sim) and Relatedness (WS-Rel) ([Finkelstein et al., 2001](#)) and SimLex-999 (SimLex) ([Hill et al., 2015](#)) datasets, and the Rare Word (RW) ([Luong et al., 2013](#)) dataset to verify if subword information improves rare word representation. Relationships are measured using the Google semantic (GSem) and syntactic (GSyn) analogies ([Mikolov et al., 2013a](#)) and the Microsoft syntactic analogies (MSR) dataset ([Mikolov et al., 2013b](#)).

We also evaluate all five models on downstream tasks from the VecEval suite ([Nayak et al., 2016](#))⁴, using only the tasks for which training and evaluation data is freely available: chunking, sentiment

| Evaluation | LV | LV-N | LV-M | SG | FT |
|------------|-------------|-------------|-------------|-------------|-------------|
| WS-Sim | .749 | .748 | .746 | .783 | .778 |
| WS-Rel | .627 | .627 | .625 | .683 | .672 |
| SimLex | .359 | .374 | .366 | .371 | .367 |
| RW | .461 | .522 | .479 | .481 | .500 |
| GSem | 80.7 | 73.8 | 80.7 | 78.9 | 77.0 |
| GSyn | 62.8 | 68.6 | 63.8 | 68.2 | 71.1 |
| MSR | 49.6 | 55.0 | 53.8 | 57.8 | 59.6 |
| Chunk | 90.4 | 90.6 | 90.5 | 90.4 | 90.4 |
| Sentiment | 77.0 | 77.0 | 77.6 | 75.3 | 77.9 |
| Questions | 87.4 | 87.4 | 87.3 | 86.6 | 85.1 |
| NLI | 43.3 | 43.4 | 43.3 | 43.4 | 43.8 |

Table 1: Word similarity (*Spearman’s rho*), analogy (% accuracy), and downstream task (% accuracy) results. In downstream tasks, for the same model accuracy varies over different runs, so we report the mean over 20 runs, in which the only significantly ($p < .05$ under a random permutation test) different result is in chunking.

and question classification, and natural language identification (NLI). The default settings from the suite are used, but we run only the *fixed* settings, where the embeddings themselves are not tunable parameters of the models, forcing the system to use only the information already in the embeddings.

Finally, we use LV-N, LV-M, and FT to generate OOV word representations for the following words: 1) “hellooo”: a greeting commonly used in instant messaging which emphasizes a syllable. 2) “marvelicious”: a made-up word obtained by merging “marvelous” and “delicious”. 3) “louisana”: a misspelling of the proper name “Louisiana”. 4) “rereread”: recursive use of prefix “re”. 5) “tuzread”: made-up prefix “tuz”.

5 Results

Results for IV evaluation are shown in table 1, and for OOV in table 2. Like in FT, the use of subword information in both LV-N and LV-M results in 1) better representation of rare words, as evidenced by the increase in RW correlation, and 2) significant improvement on the GSyn and MSR tasks, in evidence of subwords encoding information about a word’s syntactic function (the suffix “ly”, for example, suggests an adverb). There seems to be a trade-off between capturing semantics and syntax as in both LV-N and FT there is an accompanying decrease on the GSem tasks in exchange for gains on the GSyn and MSR tasks. Morphological

³<https://github.com/facebookresearch/fastText>

⁴<https://github.com/NehaNayak/veceval>

| Word | Model | 5 Nearest Neighbors |
|----------------|-------|---|
| “helooo” | LV-N | hellogoodbye, hello, helloworld, halloween, helluva |
| | LV-M | kitsos, finos, neros, nonono, theodoroi |
| | FT | hello, helloworld, hellogoodbye, halloween, joegazz |
| “marvelicious” | LV-N | delicious, marveled, marveling, licious, marvellous |
| | LV-M | marveling, marvelously, marveled, marvelled, loquacious |
| | FT | delicious, deliciously, marveling, licious, marvelman |
| “louisana” | LV-N | luisana, pisana, belisana, chiisana, rosana |
| | LV-M | louisy, louises, louison, louiseville, louisade |
| | FT | luisana, louisa, belisana, anabella, rosana |
| “rereread” | LV-N | reread, rereading, read, writeread, rerecord |
| | LV-M | alread, carreer, whiteread, unremarked, oread |
| | FT | reread, rereading, read, reiterate, writeread |
| “tuzread” | LV-N | tuzi, tuz, tuzla, prizren, momchilgrad, studenica |
| | LV-M | tuzluca, paczk, goldsztajn, belzberg, yizkor |
| | FT | pazaryeri, tufanbeyli, yenipazar, leskovac, berovo |

Table 2: We generate vectors for OOV using subword information and search for the nearest (cosine distance) words in the embedding space. The LV-M segmentation for each word is: $\{\langle \text{hell}, \text{o}, \text{o}, \text{o} \rangle\}$, $\{\langle \text{marvel}, \text{i}, \text{ciou} \rangle\}$, $\{\langle \text{louis}, \text{ana} \rangle\}$, $\{\langle \text{re}, \text{re}, \text{read} \rangle\}$, $\{\langle \text{tu}, \text{z}, \text{read} \rangle\}$. We omit the LV-N and FT n-grams as they are trivial and too numerous to list.

segmentation in LV-M appears to favor syntax less strongly than do simple n-grams.

On the downstream tasks, we only observe statistically significant ($p < .05$ under a random permutation test) improvement on the chunking task, and it is a very small gain. We attribute this to both regular and subword models having very similar quality on frequent IV word representation. Statistically, these are the words that are most likely to appear in the downstream task instances, and so the superior representation of rare words has, due to their nature, little impact on overall accuracy. Because in all tasks OOV words are mapped to the “ $\langle \text{unk} \rangle$ ” token, the subword models are not being used to the fullest, and in future work we will investigate whether generating representations for all words improves task performance.

In OOV representation (table 2), LV-N and FT work almost identically, as is to be expected. Both find highly coherent neighbors for the words “helooo”, “marvelicious”, and “rereread”. Interestingly, the misspelling of “louisana” leads to coherent name-like neighbors, although none is the expected correct spelling “louisiana”. All models stumble on the made-up prefix “tuz”. A possible fix would be to down-weight very rare subwords in the vector summation. LV-M is less robust than LV-N and FT on this task as it is highly sensitive to incorrect segmentation, exemplified in the “helooo” example.

Finally, we see that nearest-neighbors are a mixture of similarly pre/suffixed words. If these pre/suffixes are semantic, the neighbors are semantically related, else if syntactic they have similar syntactic function. This suggests that it should be possible to get *tunable* representations which are more driven by semantics or syntax by a *weighted* summation of subword vectors, given we can identify whether a pre/suffix is semantic or syntactic in nature and weigh them accordingly. This might be possible without supervision using corpus statistics as syntactic subwords are likely to be more frequent, and so could be down-weighted for more semantic representations. This is something we will pursue in future work.

6 Conclusion and Future Work

In this paper, we incorporated subword information (simple n-grams and unsupervised morphemes) into the LexVec word embedding model and evaluated its impact on the resulting IV and OOV word vectors. Like fastText, subword LexVec learns better representations for rare words than its word-level counterpart. All models generated coherent representations for OOV words, with simple n-grams demonstrating more robustness than unsupervised morphemes. In future work, we will verify whether using OOV representations in downstream tasks improves performance. We will also explore the trade-off between

semantics and syntax when subword information is used.

References

- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. Dont count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 238–247.
- Parminder Bhatia, Robert Guthrie, and Jacob Eisenstein. 2016. Morphological priors for probabilistic neural word embeddings. In *EMNLP*.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *TACL*, 5:135–146.
- Jan A. Botha and Phil Blunsom. 2014. Compositional morphology for word representations and language modelling. In *ICML*.
- Kris Cao and Marek Rei. 2016. A joint model for word embedding and word morphology. In *Rep4NLP@ACL*.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14.
- Ryan Cotterell and Hinrich Schütze. 2015. Morphological word-embeddings. In *HLT-NAACL*.
- Ryan Cotterell, Hinrich Schütze, and Jason Eisner. 2016. Morphological smoothing and extrapolation of word embeddings. In *ACL*.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppín. 2001. Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414. ACM.
- Zellig S Harris. 1954. Distributional structure. *Word*.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. Character-aware neural language models. In *AAAI*.
- Michal Konkol, Tomas Brychcin, Michal Nykl, and Tomáš Hercig. 2017. Geographical evaluation of word embeddings. In *IJCNLP*.
- Omer Levy, Yoav Goldberg, and Israel Ramat-Gan. 2014. Linguistic regularities in sparse and explicit word representations. *CoNLL-2014*, page 171.
- Wang Ling, Chris Dyer, Alan W. Black, Isabel Trancoso, Ramon Fernandez, Silvio Amir, Luís Marujo, and Tiago Luís. 2015. Finding function in form: Compositional character models for open vocabulary word representation. In *EMNLP*.
- Minh-Thang Luong, Richard Socher, and Christopher D Manning. 2013. Better word representations with recursive neural networks for morphology. *CoNLL-2013*, 104.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013b. Linguistic regularities in continuous space word representations. In *HLT-NAACL*, pages 746–751.
- Jeff Mitchell and Mark Steedman. 2015. Orthogonality of syntax and semantics within distributional spaces. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1301–1310.
- Neha Nayak, Gabor Angeli, and Christopher D. Manning. 2016. [Evaluating word embeddings using a representative suite of practical tasks](#). In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP, RepEval@ACL 2016, Berlin, Germany, August 2016*, pages 19–23. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. *Proceedings of the Empirical Methods in Natural Language Processing (EMNLP 2014)*, 12.
- Siyu Qiu, Qing Cui, Jiang Bian, Bin Gao, and Tie-Yan Liu. 2014. Co-learning of word representations and morpheme representations. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 141–150.
- Alexandre Salle, Marco Idiart, and Aline Villavicencio. 2016a. [Enhancing the lexvec distributed word representation model using positional contexts and external memory](#). *CoRR*, abs/1606.01283.
- Alexandre Salle, Aline Villavicencio, and Marco Idiart. 2016b. [Matrix factorization using window sampling and negative sampling for improved word representations](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*. The Association for Computer Linguistics.

- Hinrich Schütze. 1993. Word space. In *Advances in neural information processing systems*, pages 895–902.
- Yulia Tsvetkov, Manaal Faruqui, Wang Ling, Guillaume Lample, and Chris Dyer. 2015. Evaluation of word vector representations by subspace alignment. In *EMNLP*.
- Lyan Verwimp, Joris Pelemans, Hugo Van hamme, and Patrick Wambacq. 2017. Character-word lstm language models. In *EACL*.
- Peter; Grnroos Stig-Arne; Kurimo Mikko Virpioja, Sami; Smit. 2013. [Morfessor 2.0: Python implementation and extensions for morfessor baseline](#). D4 julkaistu kehittmis- tai tutkimusraportti tai -selvitys.
- Ivan Vúlic, Nikola Mrksic, Roi Reichart, Diarmuid Ó Séaghdha, Steve J. Young, and Anna Korhonen. 2017. [Morph-fitting: Fine-tuning word vector spaces with simple language-specific rules](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 56–68. Association for Computational Linguistics.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2016. Charagram: Embedding words and sentences via character n-grams. In *EMNLP*.
- Gerhard Wohlgenannt, Nikolay Klimov, Dmitry Mouromtsev, Daniil Razdyakonov, Dmitry Pavlov, and Yury Emelyanov. 2017. Using word embeddings for visual data exploration with ontodia and wikidata. In *BLINK/NLIWoD3@ISWC*.