



HAL
open science

Evaluation of Word Embeddings from Large-Scale French Web Content

Hadi Abdine, Christos Xypolopoulos, Moussa Kamal Eddine, Michalis Vazirgiannis

► **To cite this version:**

Hadi Abdine, Christos Xypolopoulos, Moussa Kamal Eddine, Michalis Vazirgiannis. Evaluation of Word Embeddings from Large-Scale French Web Content. Conférence Nationale en Intelligence Artificielle 2022 (CNIA 2022), Jun 2022, Saint-etienne, France. hal-04454559

HAL Id: hal-04454559

<https://hal.science/hal-04454559v1>

Submitted on 13 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Evaluation of Word Embeddings from Large-Scale French Web Content

Hadi Abdine

LIX, École Polytechnique
hadi.abdine@polytechnique.edu

Christos Xypolopoulos

LIX, École Polytechnique
christos.xypolopoulos@polytechnique.edu

Moussa Kamal Eddine

LIX, École Polytechnique
moussa.kamal-eddine@polytechnique.edu

Michalis Vazirgiannis

LIX, École Polytechnique
AUEB
mvazirg@lix.polytechnique.fr

Abstract

Distributed word representations are popularly used in many tasks in natural language processing. Adding that pretrained word vectors on huge text corpus achieved high performance in many different NLP tasks. This paper introduces multiple high-quality word vectors for the French language where two of them are trained on massive crawled French data during this study and the others are trained on an already existing French corpus. We also evaluate the quality of our proposed word vectors and the existing French word vectors on the French word analogy task. In addition, we do the evaluation on multiple real NLP tasks that shows the important performance enhancement of the pre-trained word vectors compared to the existing and random ones. Finally, we created a demo web application to test and visualize the obtained word embeddings¹. The produced French word embeddings are available to the public, along with the finetuning code on the NLU tasks² and the demo code³.

1 Introduction

Word embeddings are a type of word representation that has become very important and widely used in different applications in natural language processing. For instance, pre-trained word embeddings played an essential role in achieving an impressive performance with deep learning models on challenging natural language understanding problems. These word vectors are produced using unsupervised neural network models based on the idea that a word's meaning is related to the context where it appears. Thus, similar words in meaning have a similar representations. In result, the quality of word vectors is directly related to the amount and

the cleanness of the corpus they were trained on. Many techniques and algorithms have been proposed since 2013 to learn these word distributed representations. Mainly, we focus on Word2Vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), and FastText (Bojanowski et al., 2017) that make use of two approaches: CBoW and Skip-Gram. These methods were used to learn word vectors using huge corpora. Most of the publicly available word vectors are pretrained on massive English raw text. However, the ones for other languages are few, nonexistent, or trained using a very small corpus that cannot produce good quality word vectors.

In this paper, we are interested in creating static French word embeddings with a benchmark for all French static word embeddings. Thus, we do not include contextual word embeddings such as FlauBERT (Le et al., 2020), CamemBERT (Martin et al., 2020) and BARThez (Eddine et al., 2021) in the comparison. These attention based pretrained models, even if they have better performance, their cost in terms of both memory and speed is more significant. This work presents French word embeddings trained on a large corpus collected/crawled from the French web with more than 1M domains. In addition, we train Word2Vec embedding on other French corpora and resources in order to compare with the ones trained on the collected French corpus. Then, we evaluate these word vectors on the test set created in (Grave et al., 2018) for the word analogy task. Finally, we evaluate the French word embeddings on some tasks from the FLUE benchmark (Le et al., 2020) and compare the results with the ones of the existing French FastText (Grave et al., 2018) and Word2Vec (Fauconnier, 2015) vectors produced using Common Crawl, Wikipedia and FrWac corpus. In the end, we discuss the significance of the word analogy task on word embeddings quality.

¹nlp.polytechnique.fr/

²github.com/hadi-abdine/WordEmbeddingsEvalFLUE/

³github.com/hadi-abdine/LinguisticResourcesDemo/

2 Related Work

Since distributed word representations have been introduced, many pretrained word vectors for many languages are produced. For example, English word vectors were learned using a portion of the google news dataset and published alongside Word2Vec (Mikolov et al., 2013). After that, word representations were trained for 157 languages including **French** using FastText (Grave et al., 2018) to learn word vectors on Common Crawl corpus and Wikipedia. In 2015, (Fauconnier, 2015) pretrained multiple Word2Vec models for the French language using FrWac corpus (Baroni et al., 2009) and FrWiki dump (French data from Wikipedia). Since Word2Vec first appeared, the evaluation of the word vectors was based on the word analogy task introduced in (Mikolov et al., 2013). This task evaluates the quality of word embeddings based on a linear relation between words. Following this idea, a list of French analogy questions to evaluate French word embeddings is introduced in (Grave et al., 2018).

Finally, (Le et al., 2020) introduced a general benchmark to evaluate French NLP systems named FLUE containing multiple and diverse French tasks to evaluate French natural understanding models. Our main objective is to evaluate the quality of static word embeddings for the French language and provide new ones trained on a large, diverse dataset that perform well on word analogy tasks, as well on language understanding tasks. Mentioning that, there are many papers showing that pretrained attention based models like BERT (Devlin et al., 2019) and ELMo (Peters et al., 2018) out-perform static word embeddings. However, our intuition is that static word embeddings are still relatively crucial for low resources and small applications where using expensive models that require huge loads of GPU resources is not applicable.

3 Crawling The French Web

Seeds Our initial seeds were selected by finding the most popular web pages under the top level domain ".fr". Our intuition is that web pages constitute the core of the french web graph. To achieve this, we used a publicly available list⁴. The frontier was then updated with the new links that were discovered during crawling.

Crawler For crawling we chose Heritrix⁵, an

⁴Most popular websites with fr domain

⁵<https://github.com/internetarchive/heritrix3>

open-source web-scale crawler that is supported by the Internet Archive. Our setup constitutes of a single node with 25 threads crawling for a period of 1.5 months.

Output The generated output follows the WARC file format, with files split to 1GB each. This format has been traditionally used to store "web crawls" as sequences of content blocks harvested from the World Wide Web. The content blocks in a WARC file may contain resources in any format; examples include the binary images or audiovisual files that may be embedded or linked to HTML pages.

Extraction To extract the embedded text we used a warc-extractor tool.⁶ This tool is used in order to parse the records with the WARC and then parse the HTML pages. During this step we also integrated the FastText language detection module⁷ that allows us to filter the HTML text by language.

Deduplication To eliminate redundant data in the corpus, we used Isaac Whitfield's deduplication tool⁸, the same used on Common Crawl corpus (Ortiz Suarez et al., 2019) based on a very fast and collisions resistant hash algorithm.

Ethics Heritrix is designed to respect the robots.txt (file written by the website owners to give instructions about their site to web crawlers), exclusion directives and META nofollow tags.

4 Word Embeddings Training

In order to learn the French word embeddings, we used Gensim's Word2Vec to produce four models of CBoW (Continuous Bag of Words) French word vectors.

Two of these models were trained on a 33GB shuffled portion of the French corpus used to train FlauBERT (Le et al., 2020). This corpus, consists of 24 sub-corpora collected from different sources with diverse subjects and writing styles. The first model is noted as fr_w2v_fl_w5 trained using Word2Vec CBoW with a window size of 5 and min_count of 60, in other words, we only consider training an embedding for a word, only if it appears at least 60 times in the corpus. The second one is fr_w2v_fl_w20 trained using Word2Vec CBoW with a window size of 20 and min_count of 5.

The other two models were trained on the 33GB

⁶<https://github.com/alexeygrigorev/warc-extractor>

⁷<https://github.com/vinhkhuc/JFastText>

⁸<https://github.com/whitfin/runiq>

Embeddings	Vocab	Tool	Method	Corpus	Dimension
fr_w2v_web_w5	0.8M	Word2Vec	cbow	fr_web	300
fr_w2v_web_w20	4.4M	Word2Vec	cbow	fr_web	300
fr_w2v_fl_w5	1M	Word2Vec	cbow	flaubert_corpus	300
fr_w2v_fl_w20	6M	Word2Vec	cbow	flaubert_corpus	300
cc.fr.300(Grave et al., 2018)	2M	FastText	skip-gram	CC+wikipedia	300
frWac_200_cbow(Fauconnier, 2015)	3.6M	Word2Vec	cbow	frWac	200
frWac_500_cbow(Fauconnier, 2015)	1M	Word2Vec	cbow	frWac	500
frWac_700_cbow(Fauconnier, 2015)	184K	Word2Vec	skip-gram	frWac	700
frWiki_1000_cbow(Fauconnier, 2015)	66K	Word2Vec	cbow	wikipedia dump	1000

Table 1: Summary of the models used in our experiments

deduplicated French corpus collected from the web (section 3). The first is noted as fr_w2v_web_w5 trained on Word2Vec CBoW with a window size of 5 and min_count of 60. The second is fr_w2v_web_w20 trained on Word2Vec CBoW with a window size of 20 and min_count of 5. All the models examined in this paper have an embedding size of 300. Table 1 contains the details of each word embeddings used in this paper.

5 Word Analogy Evaluation

In this work, the first evaluation method for the word vectors is the analogy task (Mikolov et al., 2013). In this task, given three words, A, B and C, we can predict a word D such that the relation between A and B is the same between C and D with an assumption that a linear relation between word pairs indicates the quality of word embeddings. For example, if the relation between word representations of **king** and **man** is similar to the relation between the representations of **queen** and **woman**, this will imply good word embeddings. For instance, in this evaluation we use the word vectors x_A , x_B and x_C of three words A, B and C to compute the vector $x_B - x_A + x_C$, and the closest vector in the dictionary to the resulted vector will be considered the vector of the word D.

The performance of word vectors is finally computed using the average accuracy over the evaluation set.

To evaluate our French word embeddings presented in section 3, we use the French analogy dataset created in (Grave et al., 2018). This dataset contains 31 688 questions. We also compare the results of our word vectors with the one of the French FastText (cc.fr.300), also produced in (Grave et al., 2018) and the French wacky word vectors.

Results Table 2 shows the accuracy on the

Embeddings	Accuracy
fr_w2v_web_w5	41.95
fr_w2v_web_w20	52.50
fr_w2v_fl_w5	43.02
fr_w2v_fl_w20	45.88
cc.fr.300	63.91
frWac_500_cbow	67.98
frWac_200_cbow	54.45
frWac_700_sg	55.52
frWiki_1000_cbow	0.87

Table 2: Accuracy on the analogy task

analogy task for all the models. The results indicate that the FastText word vectors and the French wacky word embeddings are by far better than our CBoW Word2Vec word vectors in the analogy task. However, the authors in (Khalife et al., 2019) proved that the property of geometrical relation with respect to analogies does not generally hold and is likely to be incidental rather than systematic. In conclusion, we decided to evaluate the word vectors on natural language understanding (NLU) tasks such as text classification, paraphrasing and language inference that are presented in (Le et al., 2020).

Demo and web app To visualize and examine our word embeddings, we created a web app that contains the following tools:

1. Word analogy examiner: it takes as input three words, A, B and C, and it computes the closest ten vectors to $x_B - x_A + x_C$, and displays their corresponding words as we can see in figure 1.
2. A cosine similarity calculator: it takes two words and computes the cosine similarity be-



Figure 1: Example of word analogy using the analogy tool in our web app (with fr_w2v_web_w20).




Figure 2: 2 of 8 clusters of the query: "Top 200 closest words to **paris**" using the visualization tool in our web app (with fr_w2v_web_w20)

tween the corresponding vectors.

3. A similar words tool: find the top ten similar words to an input word.
4. A visualization tool: using T-SNE and k-means, this tool displays in a 2-D vector space the top n closest words to the word w distributed in k clusters, where n , w and k are chosen by the user. For example, in figure 2 we see a part of the graph for $n=200$, $w=paris$ and $k=8$.

6 Evaluation On FLUE Benchmark Tasks

FLUE is a French language understanding evaluation benchmark (Le et al., 2020), created to evaluate the performance of French NLP models. It contains many datasets that vary in subjects, difficulty level, size, and degree of formality. This section presents the results of finetuning of our different Word2Vec embeddings and compares them with the finetuning results of the FastText and the French wacky Word2Vec embeddings on some FLUE benchmark tasks. We also include a non-pretrained word embedding for comparison.

6.1 Text Classification

Dataset description In this task, the used dataset is the French portion of CLS (Cross Lingual Sentiment) (Prettenhofer and Stein, 2010). It consists of Amazon reviews divided into three subsets: books, DVD, and music. Following (Le et al., 2020), each subset is binarized by considering all reviews above three stars as positive and the ones below three stars as negative, while eliminating the ones with three stars. In addition, each subset contains a balanced train and test set that has around 1000 positive and 1000 negative reviews each. Finally, a validation set is formed for each subset by taking a random split of 20% of the training data.

Model description The model uses word embedding layer that varies in initial weights and in dimension according to the evaluated word embedding. This layer is followed by a two-layer 1500D (per direction) Bi-LSTM. Finally, we use the classification head used in (Devlin et al., 2019) that is composed of the following layers: a dropout layer with a dropout rate of 0.1, a linear layer, a hyperbolic tangent activation layer, a second dropout layer with the same dropout rate and a second linear layer with output dimension of two (the same as the number of classes).

We train the model for 30 epochs for all the different word embeddings while performing a grid search over three different learning rates: 5e-5, 2e-5 and 5e-6. The best-performing model on the validation set is chosen for the evaluation on the test set.

Embeddings	Books	Music	DVD
w2v_0	67.20	67.20	62.15
fr_w2v_web_w5	79.38	79.49	80.29
fr_w2v_web_w20	81.55	79.75	80.75
fr_w2v_fl_w5	82.16	81.00	79.64
fr_w2v_fl_w20	82.38	82.58	82.43
cc.fr.300	75.30	74.35	72.43
frWac_500_cbow	81.30	80.80	79.59
frWac_200_cbow	79.65	75.40	80.75
frWac_700_sg	77.55	75.20	78.43
frWiki_1000_cbow	66.30	70.05	60.94

Table 3: Accuracy on the French CLS

Results Table 3 presents the accuracy on the test set for each word embeddings. The results demonstrate the importance of pretrained models. As we see, all pre-trained word embeddings surpass by far the non pretrained embeddings with a difference that can go to more than 20% in accuracy. In addition, we see clearly the contradiction in results between the analogy task and the sentiment analysis.

6.2 Paraphrasing

Dataset description This task uses the French portion of PAWS-X (Yang et al., 2019) which extends the cross-lingual adversarial dataset for paraphrase identification. This task aims to identify if there is a semantic relation between a pair of sentences or not. The dataset is obtained from translating collected pairs of sentences in English from Wikipedia and Quora with a high lexical overlap ratio and judged by humans. The used dataset contains 49 401 training samples, 1 992 validation samples and 1 985 test samples.

Model description To finetune the word vectors on PAWS-X, we use Enhanced Sequential Inference Model (ESIM) (Chen et al., 2017). This model is formed from three essential components: input encoding, local inference modeling and inference composition.

After the word embedding layer, the input encoding layer is composed of 1-layer 1500D BiLSTM that encodes the local inference information of sen-

tences A and B with a length of respectively l_a and l_b to give the output \bar{A} and \bar{B} .

The following component is the local inference modeling where:

1. We compute the similarity matrix E between the two sentences representing the attention weights, such that $e_{ij} = \bar{a}_i^T \bar{b}_j$.
2. We compute the local relevance between the two sentences where each word will be represented as a weighted summation of the relevant information to this word. For instance, $\tilde{a}_i = \sum_{j=1}^{l_b} softmax(e_{ij}) \bar{b}_j$ and $\tilde{b}_j = \sum_{i=1}^{l_a} softmax(e_{ij}) \bar{a}_i$.
3. We find the enhancement of local inference information by a concatenation in a way that the authors expect would help to enhance the local inference. For instance, the resulted vectors will be $m_a = [\bar{a}; \tilde{a}; \bar{a} - \tilde{a}; \bar{a} \odot \tilde{a}]$ and $m_b = [\bar{b}; \tilde{b}; \bar{b} - \tilde{b}; \bar{b} \odot \tilde{b}]$.

The fourth layer is a projection layer that is formed from a linear layer with a hidden dimension of 1500D, ReLU activation and a dropout layer with a rate of 0.1

Finally, the inference composition layer is composed of 1-layer 1500D BiLSTM layer that gives the vectors V_a and V_b as output. We used the same classification head as in the text classification task. Where the input to the classification head is the concatenated vector of average pooling and max pooling of V_a and V_b .

Embeddings	Accuracy
w2v_0	73.49
fr_w2v_web_w5	77.12
fr_w2v_web_w20	78.02
fr_w2v_fl_w5	75.2
fr_w2v_fl_w20	74.24
cc.fr.300	62.80
frWac_500_cbow	71.47
frWac_200_cbow	74.80
frWac_700_sg	67.39
frWiki_1000_cbow	69.76

Table 4: Accuracy on the French PAWS-X using ESIM

Results The final accuracy of each word vectors is reported in Table 4. Again, one can see that not only there are opposed results between NLU tasks and the analogy task, but also that random word

vectors outperform the word vectors that achieved the second-highest accuracy in the analogy task.

6.3 Natural Language Inference

Dataset and model description The French natural language inference task in FLUE uses the French portion of XNLI dataset (Conneau et al., 2018). It contains NLI data for 15 languages. Each pair-sentences in the test and validation sets are annotated manually by humans with three classes of inference: entailment, neutral, and no entailment (contradiction). Where the French portion in the training dataset of XNLI is obtained by machine translation. In this task, the goal is to determine whether there is entailment, contradiction, or neutral relation from a sentence p called premise to another sentence h called hypothesis. Note that the same sample can be used twice with a reverse order between the two sentences with a different label. The dataset consists of 92 702 training samples, 2 491 validation samples and 5 010 test samples. The used model to finetune the word vectors on this task is ESIM with the same configurations and parameters described in section 6.2.

Embeddings	Accuracy
w2v_0	61.37
fr_w2v_web_w5	67.71
fr_w2v_web_w20	68.27
fr_w2v_fl_w5	69.41
fr_w2v_fl_w20	69.57
cc.fr.300	64.70
frWac_500_cbow	63.82
frWac_200_cbow	63.74
frWac_700_sg	60.78
frWiki_1000_cbow	61.34

Table 5: Accuracy on the French XNLI using ESIM

Results We report the final accuracy of finetuning the various word vectors on the French XNLI dataset in Table 5. The results keep showing the advantages of pretrained weights and the out-performance of our CBoW Word2Vec weights over the already existing word vectors.

6.4 Noun Sense Disambiguation

Dataset description This task is proposed by (Le et al., 2020) for the word sense disambiguation (WSD) of French that targets the nouns only. It is based on the French portion of the WSD task in (Navigli et al., 2013) to create the evaluation

set composed of 306 sentences and 1 445 French nouns annotated with WordNet sense keys and manually verified. The training set is obtained by using the best English-French machine translation system in fairseq tool (Ott et al., 2019) to translate the SemCor and the WordNet Glosses Corpus to French.

Model description We used the same classifiers presented by (Vial et al., 2019) that forward the output of our words vectors into a stack of 6 transformer encoder layers and predict the meaning of the word through a softmax layer at the end. We used the same parameters as in (Le et al., 2020).

Embedding	Single		
	Mean	Std	Ens
w2v_0	47.85	± 1.17	52.87
fr_w2v_web_w5	50.76	± 1.4	53.77
fr_w2v_web_w20	50.28	± 0.92	53.22
fr_w2v_fl_w5	50.16	± 1.41	53.36
fr_w2v_fl_w20	50.65	± 1.62	52.46
cc.fr.300	49.28	± 1.5	52.39

Table 6: F1 score (%) on the NSD task

Results For every word vectors, we report in Table 6 the mean and standard deviation values of F1 scores (%) of 8 individual models. For instance, each model predicts the word with the highest probability in the softmax layer’s output and the F1 score (%) of the ensemble of models which averages out the output of the softmax layer. We observe in this task that, even though fr_w2v_web_w5 slightly outperforms the other word vectors, we have a very similar F1 score (%) for all the models. We can say that, for this task, pretrained word vectors do not improve the final score. We think that the nature of these vectors being static despite the context while the goal of the task is to investigate the different meanings of a word is the reason for the similar score between these models.

7 Conclusion

In this work, we contribute word embeddings learned on data crawled from the French web and other word vectors trained on a mixed dataset formed from different and diverse sources and subjects in French including common crawl and Wikipedia. These word vectors can be used as initial weights for various deep learning models which can bring important performance boost. In addition, these general domain word embeddings

could be further pretrained on specific domain data such as health and legal domain in order to adapt the weights to a suitable context.

Moreover, we created a benchmark of four NLP tasks to compare the quality of four produced French word embeddings and five existing ones. All the resources are presented and available to the research community via our web app.

Finally, we showed that the word analogy task could not be trusted to judge on word embedding's quality. The reason behind these results on the word analogy task could be a potential future study.

Acknowledgment

This research was supported by the ANR chair AML/HELAS (ANR-CHIA-0020-01).

References

- Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. 2009. [The wacky wide web: A collection of very large linguistically processed web-crawled corpora](#). *Language Resources and Evaluation*, 43:209–226.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [Xnli: Evaluating cross-lingual sentence representations](#). *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Moussa Kamal Eddine, Antoine J. P. Tixier, and Michalis Vazirgiannis. 2021. [Barthez: a skilled pre-trained french sequence-to-sequence model](#).
- Jean-Philippe Fauconnier. 2015. [French word embeddings](#).
- Edouard Grave, Piotr Bojanowski, Prakhya Gupta, Armand Joulin, and Tomas Mikolov. 2018. [Learning word vectors for 157 languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Sammy Khalife, Leo Liberti, and Michalis Vazirgiannis. 2019. [Geometry and analogies: A study and propagation method for word representations](#). In *Statistical Language and Speech Processing*, pages 100–111, Cham. Springer International Publishing.
- Hang Le, Loïc Vial, Jibril Frej, Vincent Segonne, Maximin Coavoux, Benjamin Lecouteux, Alexandre Alauzen, Benoit Crabbe, Laurent Besacier, and Didier Schwab. 2020. [FlauBERT: Unsupervised Language Model Pre-training for French](#). In *LREC*, Marseille, France.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Roberto Navigli, David Jurgens, and Daniele Vannella. 2013. [SemEval-2013 task 12: Multilingual word sense disambiguation](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 222–231, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Pedro Javier Ortiz Suarez, Benoit Sagot, and Laurent Romary. 2019. [Asynchronous pipelines for processing huge corpora on medium to low resource infrastructures](#). In *Proceedings of the Workshop on Challenges in the Management of Large Corpora (CMLC-7) 2019, Cardiff, 22nd July 2019*, pages 9 – 16, Mannheim. Leibniz-Institut für Deutsche Sprache.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.

- 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, Doha, Qatar. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Peter Prettenhofer and Benno Stein. 2010. [Cross-language text classification using structural correspondence learning](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1118–1127, Uppsala, Sweden. Association for Computational Linguistics.
- Loïc Vial, Benjamin Lecouteux, and Didier Schwab. 2019. [Sense Vocabulary Compression through the Semantic Knowledge of WordNet for Neural Word Sense Disambiguation](#). In *Global Wordnet Conference*, Wrocław, Poland.
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. [Paws-x: A cross-lingual adversarial dataset for paraphrase identification](#). *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.