


When FastText Pays Attention: Efficient Estimation of Word Representations using Constrained Positional Weighting


Vít Novotný

(Faculty of Informatics Masaryk University, Brno, Czech Republic
 <https://orcid.org/0000-0002-3303-4130>, witiko@mail.muni.cz)

Michal Štefánek

(Faculty of Informatics Masaryk University, Brno, Czech Republic
 <https://orcid.org/0000-0003-1766-5538>, stefanik.m@mail.muni.cz)

Eniafe Festus Ayetiran

(Department of Mathematical Sciences, Achievers University, Owo, Nigeria
Faculty of Informatics Masaryk University, Brno, Czech Republic
 <https://orcid.org/0000-0002-6816-2781>, ayetiran@mail.muni.cz)

Petr Sojka

(Faculty of Informatics Masaryk University, Brno, Czech Republic
 <https://orcid.org/0000-0002-5768-4007>, sojka@fi.muni.cz)

Radim Řehůřek

(RARE Technologies Ltd., radim@rare-technologies.com)

Abstract: In 2018, Mikolov et al. [Mik+18] introduced the positional language model, which has characteristics of attention-based neural machine translation models and which achieved state-of-the-art performance on the intrinsic word analogy task. However, the positional model is not practically fast and it has never been evaluated on qualitative criteria or extrinsic tasks. We propose a *constrained positional model*, which adapts the sparse attention mechanism from neural machine translation to improve the speed of the positional model. We evaluate the positional and constrained positional models on three novel qualitative criteria and on language modeling. We show that the positional and constrained positional models contain interpretable information about the grammatical properties of words and outperform other shallow models on language modeling. We also show that our constrained model outperforms the positional model on language modeling and trains twice as fast.

Keywords: Word embeddings, Positional embeddings, Language modeling, Attention

Categories: G.1.3, G.3, I.2, I.2.7

DOI: 10.3897/jucs.69619

1 Introduction

“Words do not mean, people do.” [Wri08]

Word representations of shallow log-bilinear language models (LBLs) such as word2vec [Mik+13a; Mik+13b] and fastText [Boj+17] have found many applications in natural language processing (NLP) including word similarity, word analogy, and language

modeling [Boj+17; ASN21] as well as word sense disambiguation [CLS14; ASN21], text classification [Kus+15; Nov+20a], semantic text similarity [CD17], and information retrieval [Nov+20b, Section 4]. Recently, Devlin et al. [Dev+18] introduced the deep attention-based language model BERT, which has redefined the state of the art for eleven NLP tasks and turned LBLs to a baseline. Independently, Mikolov et al. [Mik+18] have introduced the positional LBL, which resembles attention-based language models and which has reached state-of-the-art performance on the intrinsic word analogy task.

Clark et al. [Cla+19] also showed that ensembling LBLs with BERT improves performance on the dependency parsing task compared to either LBLs or BERT alone, which has reinvigorated the fading interest in LBLs. Although surprising, the results of Clark et al. [Cla+19] are supported by cognitive psychology: Kahneman [Kah11] describes the human mind as an interplay of two systems: the fast, intuitive, and emotional System 1, and the slow, effortful and logical System 2. Peters et al. [Pet+06] have shown that systems 1 and 2 are mutually supportive and that System 1 adapts to new and more challenging tasks by coaching System 2 to take over its more menial tasks. If we treat Kahnemann’s systems 1 and 2 as a metaphor for LBLs and BERT, the results of Clark et al. [Cla+19] seems natural.

In our paper, we describe the relationship between the attention mechanism and the positional LBL of Mikolov et al. [Mik+18], and we propose our constrained positional LBL that adapts the attention sparsification techniques of Dai et al. [Dai+19], Child et al. [Chi+19], Beltagy, Peters, Cohan [BPC20], and Zaheer et al. [Zah+20] for LBLs. We also develop three novel qualitative criteria, which we use to evaluate the positional and constrained positional LBLs in addition to the extrinsic language modeling task.

The rest of our paper is structured as follows: In Section 2, we describe the dense and sparse attention mechanisms, and we relate them to the positional LBL of Mikolov et al. [Mik+18] and our proposed constrained positional LBL. In Section 3, we describe our experimental setup and propose three novel qualitative evaluation measures. In Section 4, we discuss the results of our experiments. We conclude in Section 5 by summarizing our results and suggesting directions for future work.

2 Models

In this section, we describe the dense attention mechanism and we relate it to the positional LBL. Additionally, we describe attention sparsification techniques and we use them to develop our proposed constrained positional LBL.

2.1 Attention

In this section, we describe the purpose of attention in neural machine translation (NMT), and we describe sparsification techniques that make attention computationally tractable.

2.1.1 Dense attention

Early neural machine translation used *encoder-decoder models*, where an encoder recurrent neural network (RNN) would first read and encode a *source sequence* into a *fixed-length context vector* and a decoder RNN would then produce a *translated target sequence* from the context vector [SVL14; Cho+14b]. Due to the context vector’s fixed length, translation performance would deteriorate for longer sequences [Cho+14a]. To

enable the translation of longer source sequences, Bahdanau, Cho, Bengio [BCB16] equipped the decoder with a dense *attention* mechanism. Instead of having a single encoded vector for the entire source sequence, the dense attention would construct a different context vector for each target word. Here, the context vector would be a weighted average of the encoder's hidden states, where the weights would be trained to relate relevant source words to the target word.

Following the success of dense attention in NMT, Cheng, Dong, Lapata [CDL16] proposed to use the dense attention mechanism directly in the long-short-term memory (LSTM) cells of RNNs: Instead of computing the current *memory* and *hidden state* using the previous memory and hidden state alone, the current memory and hidden state would be computed as weighted averages of all previous memories and hidden states. Dense attention would act as a random-access memory mechanism, enabling the LSTM to recall long-range memories. Later, Vaswani et al. [Vas+17] proposed the *Transformer* architecture, which successfully replaced recurrence by the vertical stacking of dense attention, and which has been shown to be a *Turing-complete* [PMB19] *universal approximator* [Yun+20].

2.1.2 Sparse attention

Since the dense attention mechanism learns weights for all pairs of source and target words, its space complexity is $\mathcal{O}(n^2)$ in the source sequence length. Several *sparse attention* architectures have been proposed in literature to enable the translation of longer source sequences by making the space complexity $\mathcal{O}(n)$.

Child et al. [Chi+19] proposed the *Sparse Transformer* architecture, which factorized the dense attention using p separate attention heads to learn only $\mathcal{O}(n \cdot \sqrt[p]{n})$ weights. They showed that the resulting model could use larger context sizes and achieved significantly better results than Transformers on density modeling tasks.

Following the success of Sparse Transformers, Beltagy, Peters, Cohan [BPC20] proposed the *Longformer* architecture, which reduced the number of attention weights to $\mathcal{O}(n)$ and achieved significantly better results than Transformers on multiple long document tasks including question answering, coreference resolution, and classification.

Finally, Zaheer et al. [Zah+20] proposed the *BigBird* architecture. Like Longformers, BigBird also used $\mathcal{O}(n)$ weights. Unlike Longformers, BigBird has been shown to be a Turing-complete universal approximator. Although attention enabled the recollection of long-range memories, sparse attention made it computationally tractable to do so.

2.2 Log-bilinear language models

In this section, we propose our constrained positional model that learns word representations while taking into account morphology and the mutual positions of words. When modeling positions of words, we only use a sparse subset of word vector features, following the hypothesis of Bach [Bac12] that only a fraction of a word's meaning depends on the narrow context of a paragraph, whereas the rest of its meaning is either fixed or depends on a broader context.

We first present the general word2vec model of Mikolov et al. [Mik+13a] and Mikolov et al. [Mik+13b], followed by the subword fastText model of Bojanowski et al. [Boj+17], and the positional model of Mikolov et al. [Mik+18]. Finally, we propose our constrained positional model together with its theoretical foundations, computational benefits, and its close relation to the sparse attention mechanism described in Section 2.1.2.

2.2.1 General model

Mikolov et al. [Mik+13a] introduced the continuous bag of words (CBOW) model, which learns word representations by predicting a masked word w_t from its context $C_t = w_{t-c}, \dots, w_{t-1}, w_{t+1}, \dots, w_{t+c}$, where c is *window size* and w_1, \dots, w_T is the *training corpus*:

$$\arg \min_{\theta} \left[L(\theta) = - \sum_{t=1}^T \log \Pr(w_t \mid C_t; \theta) \right]. \quad (1)$$

To estimate $\Pr(w_t \mid C_t)$, Mikolov et al. [Mik+13b] used a simplified variant of the noise contrastive approximation [GH12], which they called *negative sampling*:

$$\Pr(w_t \mid C_t) = \sigma(s(w_t, C_t)) \prod_{n \in N_{C_t}} \sigma(-s(n, C_t)), \quad (2)$$

where σ is the logistic function $x \mapsto 1/(1+e^{-x})$, N_{C_t} is a set of negative examples n for context C_t , and $s(w_t, C_t)$ is a scoring function that measures how well the masked word w_t matches the context C_t :

$$s(w_t, C_t) = \mathbf{u}_{C_t}^\top \cdot \mathbf{v}_{w_t}, \mathbf{u}_{C_t} = \frac{1}{|C_t|} \sum_{w \in C_t} \mathbf{u}_w. \quad (3)$$

Here, $\mathbf{u}_w \in \mathbb{R}^D$ is the *input vector* of the context word w , $\mathbf{v}_{w_t} \in \mathbb{R}^D$ is the *output vector* of the masked word w_t , \mathbf{u}_{C_t} is the *context vector* and D is the number of word vector features, which is usually in the low hundreds.

Li [Li92] has shown that if we order words $w_{(i)}$ by their decreasing relative frequencies $f_{w_{(i)}}$ in a corpus, then $f_{w_{(i)}}$ exhibits a power law:

$$f_{w_{(i)}} = \frac{c}{i^\alpha}, \text{ where } c \approx 0.1 \text{ and } \alpha \approx 1. \quad (4)$$

This law, originally proposed for English by Zipf [Zip32], shows that most words in our training corpus will only represent a small subset of our vocabulary. By the end of the training, CBOW will have overfit the input and output vectors of the few most frequent words, whereas it will have underfit the input and output vectors of most other words.

To equalize the number of training samples for vocabulary words, Mikolov et al. [Mik+13b] discard corpus words with the following probability:

$$\Pr_{\text{discard}}(w_t) = \max \left(0, 1 - \sqrt{\frac{r}{f_{w_t}}} \right), \quad (5)$$

where the low-pass threshold r ensures that rare words w_t with $f_{w_t} \leq r$ are never discarded.

2.2.2 Subword model

The CBOW model only learns representations for words that are present in the training corpus. Additionally, vectors for different inflectional forms of a word share no weights, which delays training convergence for morphologically rich languages.

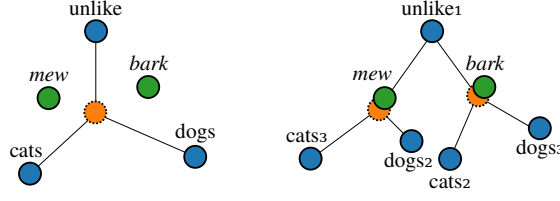


Figure 1: The subword (left) and positional (right) models with their input (blue), context (orange), and output (green) vectors for the sentences “Unlike dogs, cats \langle masked word \rangle .” and “Unlike cats, dogs \langle masked word \rangle .” with two masked words: mew and bark.

In response, Bojanowski et al. [Boj+17] have extended CBOW by modeling subwords instead of words: The input vector \mathbf{u}_w for a word w become a sum of the input vectors \mathbf{u}_g for the subwords $g \in G_w$ of w :

$$\mathbf{u}_w = \sum_{g \in G_w} \mathbf{u}_g. \quad (6)$$

2.2.3 Positional model

In many sentences, the position of context words influences their syntactic function, which is important for predicting the masked word. Consider the following two sentences, which produce an identical context vector \mathbf{u}_{C_t} despite the different masked words:

1. Unlike dogs, cats \langle masked word \rangle .
2. Unlike cats, dogs \langle masked word \rangle .

If the context C_t is large, distant words only introduce noise to the context vector \mathbf{u}_{C_t} .

To better adapt to these situations, we would like to have separate input vectors $\mathbf{u}_{w,p}$ for different positions $p \in P$ of a context word w :

$$\mathbf{u}_{C_t} = \frac{1}{|P|} \sum_{p \in P} \mathbf{u}_{w_{t+p},p}. \quad (7)$$

See also Figure 1. Since this would increase the size of the vocabulary by a factor of $|P| = 2c$, Mikolov et al. [Mik+18] proposed the *positional weighting*:

$$\mathbf{u}_{w_{t+p},p} = \mathbf{u}_{w_{t+p}} \odot \mathbf{d}_p, \quad (8)$$

Here, $\mathbf{d}_p \in \mathbb{R}^{D'}$, $p \in P$ are *positional vectors* with $D' = D$ features and $\odot : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}^D$ is the Hadamard vector product.

Technically, the positional model may seem different from dense attention: Positional vectors change the input vectors of context words, so that the vectors better reflect the relationship between the masked word and the context words given their relative positions. In contrast, dense attention relates each target word in the translated sequence to relevant context words in the source sentence. However, they both serve the same purpose: to make the context vector for a given masked/target word more meaningful.

Compared to the subword model, the positional model more than doubles the training time, since we need to compute the Hadamard product and for each gradient update of an input vector $\mathbf{u}_w \in \mathbb{R}^D$, we also need to update the weights of a positional vector $\mathbf{d}_p \in \mathbb{R}^D$. The model can benefit from larger contexts C_t , but the training time scales linearly with the context window size c , which makes the model even more expensive.

Mikolov et al. [Mik+18] used the positional model to improve the state-of-the-art accuracy on the English word analogy task by 5%. This demonstrates the importance of relating different positions of a sentence when creating its representation.

2.2.4 Constrained positional model

According to the hypothesis of Bach [Bac12], the meaning of most words is partially *fixed* and partially dependent on the *narrow context* of a paragraph as well as the *broad context* that includes the conversational setting, the time and location of an utterance, and salient common ground, which may or may not be captured in the text:

To hold that certain terms are context sensitive is not to deny that they have dictionary meanings. The claim is not that their meanings vary with the context. It is that their (standing) meanings determine their contents as a function of contexts of their use. After all, we wouldn't look words up in the dictionary if they didn't have (fairly) stable meanings. [...] Indeed, two different kinds of context are involved. Narrow context consists of matters of objective fact to which the determination of the semantic contents of certain expressions are sensitive. Broad context is the conversational setting, the mutual cognitive context or salient common ground. It includes the current state of the conversation (what has just been said, what has just been referred to, etc.), the physical situation (if the parties are face to face), salient personal knowledge, and relevant broader common knowledge.

For example, consider the following sentence:

- Fruit flies like $\langle \text{masked word} \rangle$.

The sentence admits at least two interpretations:

1. what a fly likes (adj-noun-verb- $\langle \text{mask} \rangle$), 2. how fruit flies (noun-verb-prep- $\langle \text{mask} \rangle$).

Some masked words, such as “moisture”, satisfy only the first interpretation. Others, such as “a vegetable”, satisfy both interpretations.

Let us now rearrange the sentence as follows:

- $\langle \text{Masked word} \rangle$ flies like fruit.

The rearranged sentence only admits the second interpretation. The masked words still include “a vegetable” but no longer “moisture”.

To better adapt to these situations, context vectors \mathbf{u}_{C_t} should contain two types of features:

1. D' narrow-context-dependent features that take the positions of context words into account and inhibit the prediction of *moisture* in the rearranged sentence, and
2. $D - D'$ fixed and broader-context-dependent features that disregard the positions of context words and encourage the prediction of “a vegetable” in both sentences.

Since CBOW does not model the broader context, we cannot distinguish between fixed and broader-context-dependent features. However, we can reduce broader-context-dependence with diachronic CBOW [Yao+18].

In the positional model, $D = D'$. Therefore, no word vector features are either fixed or broader-context-dependent. To represent the parts of a word's meaning that are fixed or dependent on the broader context, we propose to constrain the number of positionally-dependent features as follows:

$$0 < D' \ll D. \quad (9)$$

We define the constrained Hadamard vector product $\odot: \mathbb{R}^D \times \mathbb{R}^{D'} \rightarrow \mathbb{R}^D, D' < D$:

$$\mathbf{u}_{w_{t+p}} \odot \mathbf{d}_p = \mathbf{u}_{w_{t+p}} \odot [\mathbf{d}_p \overbrace{1 \ 1 \ \dots \ 1}^{D-D' \text{ times}}]. \quad (10)$$

When D' is small, the constrained positional model can reach the speed of the subword model while modeling both the fixed and the context-dependent parts of a word's meaning. The model can benefit from larger contexts C_t without making the computational complexity of training impractical, which is also the purpose of the sparse attention mechanism.

3 Experimental setup

In this section, we describe our baseline, the initialization of weights, the hyperparameter and parameter optimization, the qualitative evaluation measures, the extrinsic NLP tasks used for performance estimation, and our training corpora.

3.1 Baseline

In our experiments, we compare our constrained positional model against the subword and positional models described in Section 2.2. For the subword model, we use the implementation in Gensim 3.8.3 [RS10]. Since no public implementation of the positional model exists, we release our own implementation as a free open-source software library.¹

3.2 Initialization

For the general model, we follow the implementation of Bojanowski et al. [Boj+17] and we initialize the features u_i of the input word vectors \mathbf{u}_w as i.i.d. r.v.'s with continuous uniform distribution:

$$\mathbf{u}_w = (u_1, \dots, u_D), u_i \sim \mathcal{U}\left(\pm \frac{1}{D}\right). \quad (11)$$

We initialize the output word vectors \mathbf{v}_{w_t} to zero. For the subword model, we initialize the input subword vectors \mathbf{u}_g as in (11) and we also initialize the output subword vectors \mathbf{v}_g to zero.

For the positional model, Mikolov et al. [Mik+18] do not describe the initialization of either the input subword vectors \mathbf{u}_g or the positional vectors \mathbf{d}_p . Since no public implementation exists either, we initialize the features u_i of \mathbf{u}_g and the features d_j of \mathbf{d}_p as i.i.d. r.v.'s with the square-root normal distribution $\mathcal{N}^{0.5}(\mu, \sigma^2)$ of Pinelis [Pin18]:

$$\begin{aligned} \mathbf{u}_g &= (u_1, \dots, u_D), \mathbf{d}_p = (d_1, \dots, d_{D'}), \\ u_i \sim d_j &\sim \mathcal{N}^{0.5}(\mu, \sigma^2), \mu = 0, \sigma^2 = \frac{1}{3D^2}. \end{aligned} \quad (12)$$

See also Appendix A, where we show two other initialization options for the positional model, discuss their properties, and show their practical effect on the training of the model.

For the constrained positional model, we initialize the first D' features of \mathbf{u}_g and \mathbf{d}_p as in (12) and the other $D - D'$ features of \mathbf{u}_g and \mathbf{d}_p as in (11).

¹ <https://github.com/MIR-MU/pine>

3.3 Optimization

In this section, we describe which hyperparameters of the subword, positional, and constrained positional models we set according to previous work and which hyperparameters we optimized using the English word analogy task. We also describe how we train the model parameters θ .

3.3.1 Hyperparameters

For the subword, positional, and constrained positional models, we use the following hyperparameter values of Mikolov et al. [Mik+18], which give state-of-the-art performance on the English word analogy task: We store subwords of size 3–6 in a vocabulary backed by a hash table with bucket size $2 \cdot 10^6$. We discard words with less than 5 occurrences in the corpus and we equalize the number of training samples with the low-pass threshold $r = 10^{-5}$. We use $D = 300$ features in the input and output subword vectors. For the negative sampling loss, we use $|N_{C_t}| = 10$ negative samples. For the backpropagation of the loss function L , we use the initial learning rate $\gamma_0 = 0.05$.

For the subword, positional, and constrained positional models, we optimize the context window size c , because unlike the subword model, the positional model should benefit from larger contexts. For the constrained positional model, we optimize the number of positional features D' to find the proper ratio between the fixed, narrow-context-dependent, and broader-context-dependent parts of a word's meaning. To find the optimal hyperparameter values, we maximize a model's accuracy on the English word analogy task [Mik+13b] using Sequential Model-Based Optimization with the Tree-structured Parzen Estimator. Like Grave et al. [Gra+18], we restrict the vocabulary for word analogies to the $2 \cdot 10^5$ most frequent words in the training corpus.

3.3.2 Parameters

Following Bojanowski et al. [Boj+17], we optimize the model's parameters θ by stochastic gradient descent over one epoch with the loss function $L(\theta)$ presented in (1) and with a linear decay of the learning rate γ_t from γ_0 to zero:

$$\gamma_t = \gamma_0 \cdot \left(1 - \frac{t}{T}\right). \quad (13)$$

We optimize the parameters in parallel using the HogWild lock-free approach of Recht et al. [Rec+11] with 8 Intel Xeon X7560 2.26 GHz CPU cores. Since the optimization problem (1) is not sparse w.r.t. the positional vectors d_p , most of which are updated at each training step, HogWild is less appropriate for the positional and constrained positional models than for the general and subword models. For all models, we report training times.

3.4 Qualitative evaluation

In this section, we propose qualitative evaluation measures, which we use to show the properties of the positional and constrained positional models. See also Appendix B, where we show how our proposed measures relate to the conditional probability $\Pr(w_t \mid C_t)$ from (2).

3.4.1 Masked word prediction

For the example sentences C_t of the positional and constrained positional models from Section 2.2, we show masked words w_t in the descending order of the conditional probabilities $\Pr(w_t \mid C_t)$ from (2).

3.4.2 Importance of positions

For each position p , we show the min-max-scaled ℓ_2 -norm $\|\mathbf{d}_p\|$ of the positional vector \mathbf{d}_p , which measures the importance of position p for predicting masked words.

Additionally, we cluster the D' features $d_{p,j}$ of the positional vectors \mathbf{d}_p . For each cluster J and a position p , we show the mean absolute value $1/|J| \cdot \sum_{j \in J} |d_{p,j}|$, which measures the importance of position p according to cluster J .

3.4.3 Importance of context words

For clusters J and context words w , we use the mean absolute value $1/|J| \cdot \sum_{j \in J} |u_{w,j}|$ to measure the importance of context words w using cluster J , where $u_{w,j}$ are the features of the input vector \mathbf{u}_w for w . For each cluster J , we show context words whose importance is maximized by J .

3.5 Performance estimation

In this section, we describe the extrinsic language modeling task, which we used to estimate the performance of the input word vectors \mathbf{u}_w produced by the subword, positional, and constrained positional models.

3.5.1 Language modeling

For language modeling, we use a recurrent neural network (RNN) with the following architecture:

1. an input layer mapping a vocabulary V of words w to their *frozen* input vectors \mathbf{u}_w ,
2. two hidden layers with $D = 300$ LSTM units,
3. a fully-connected linear layer of size $|V|$, and
4. a softmax output layer that computes a probability distribution over the vocabulary V using tied weights [IKS17].

We evaluate our language model on the English datasets² introduced by Botha, Blunsom [BB14] and we report the validation and test perplexities. We use the same preprocessing and data splits as Botha, Blunsom [BB14].

To train the RNN, we use stochastic gradient descent over 50 epochs, negative log-likelihood loss, dropout 0.5, batch size 40, and an initial learning rate 20 that is divided by 4 after each epoch with no decrease of validation loss. We clip gradients with ℓ_2 -norm above 0.25.

² <http://bothameister.github.io/>

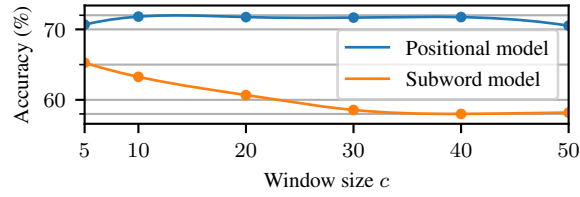


Figure 2: Word analogy accuracy of the subword and positional models trained with different context window sizes c .

3.6 Datasets

For hyperparameter optimization, parameter optimization, qualitative evaluation, and performance estimation, we use the 2017 English Wikipedia dataset³ over a single epoch as our training corpus. The dataset contains 14.3 GiB of raw text.

We preprocess our dataset by lower-casing and by tokenizing to longest sequences of Unicode characters with the *word* property. After tokenization, our dataset contains 2,423,655,228 words.

4 Results

In this section, we show and discuss the results of hyperparameter and parameter optimization, qualitative evaluation, and performance estimation.

4.1 Optimization

Table 1 shows that the positional and constrained positional models benefit from larger contexts compared to the subword model. This is further evidenced by Figure 2, which shows that the accuracy of the subword model steadily declines as the window size increases, whereas the positional model can cope with window sizes up to 40.

Table 1 also shows that the reduction of positional dimensionality D' halves the training time of the constrained positional model compared to the positional model. Figure 3 shows that the reduction of positional dimensionality also improves the accuracy of the constrained positional model compared to the positional model.

³ <https://github.com/RaRe-Technologies/gensim-data> (release wiki-english-20171001)

Model	c	D'	Training Time
Subword	5		1 hour and 11 minutes
Positional	15	300	4 hours and 12 minutes
Constrained positional	15	60	2 hours and 5 minutes

Table 1: The optimal context window sizes c and numbers of positional features D' , and training times in hours for the subword, positional, and constrained positional models.

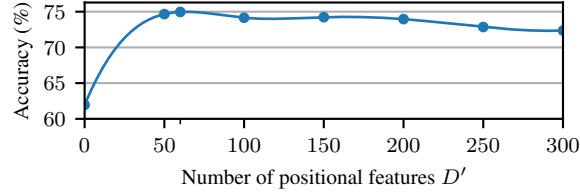


Figure 3: Word analogy accuracy of the constrained positional model trained with different numbers of positional features D' .

4.2 Masked word prediction

Table 2 shows that the positional model predicts:

$$\Pr(\text{mew} \mid C_t^1) > \Pr(\text{bark} \mid C_t^1), \quad (14)$$

$$\Pr(\text{mew} \mid C_t^2) < \Pr(\text{bark} \mid C_t^2). \quad (15)$$

This matches our expectations and indicates that the model’s context vectors contain narrow-context-dependent features that take the positions of context words “dogs” and “cats” into account.

Table 2 also shows that the constrained positional model predicts:

$$\Pr(\text{moisture} \mid C_t^3) > \Pr(\text{moisture} \mid C_t^4), \quad (16)$$

$$\Pr(\text{vegetable} \mid C_t^3) \approx \Pr(\text{vegetable} \mid C_t^4). \quad (17)$$

This indicates that the model contains not only narrow-context-dependent features that take the position of “moisture” into account, but also fixed and broader-context-dependent features that disregard the position of “a vegetable”.

4.3 Importance of positions and context words

Figure 4 shows that in the positional and constrained positional models, the importance of positions $p \in [-2; 2]$ sharply decreases with their distance from the masked word. This shows that one of the basic functions of positional weighting is the attenuation of distant context words.

In the positional model, the importance of positions $p \notin [-2; 2]$ increases with their distance from the masked word and even exceeds the importance of position $p = -2$ in the distant left context $p < -12$. In the constrained positional model, the importance of positions $p \notin [-2; 2]$ is almost constant. Below, we will explain the cause of this difference using cluster analysis.

Figure 5 shows that in the positional model, the features of positional vectors fall into three main clusters: The two smaller clusters, which we call *antepositional* and *postpositional*, and the bigger cluster, which is missing from the constrained positional model and which we call *informational*.

The antepositional and postpositional features increase the importance of positions p in anteposition $(-2, -1)$ and in postposition $(1, 2)$ of the masked word, respectively. Context words whose importance is maximized by antepositional features include “in”, “for”, and “coca”. Context words whose importance is maximized by postpositional features include

$C_t^1 = \text{"Unlike dogs, cats } \langle \text{masked word} \rangle \text{"}$	$C_t^2 = \text{"Unlike cats, dogs } \langle \text{masked word} \rangle \text{"}$	$C_t^3 = \text{"Fruit flies like } \langle \text{masked word} \rangle \text{"}$	$C_t^4 = \text{"} \langle \text{Masked word} \rangle \text{ flies like fruit."}$
# Prediction	# Prediction	# Prediction	# Prediction
1 cats	1 kennels	1 fruit	1 fruit
2 spayed	2 cats	2 flies	2 insects
3 kennels	3 puppies	3 insects	3 flies
⋮	⋮	⋮	⋮
1820 mew (100%)	⋮	246 vegetable (99.9%)	⋮
⋮	4065 bark (99.9%)	⋮	259 vegetable (99.9%)
5581 bark (99.7%)	⋮	9036 moisture (69.6%)	⋮
⋮	5623 mew (99.8%)	⋮	33465 moisture (42.8%)

(a) Positional model
(b) Constrained positional model

Table 2: Masked words w_t predicted by the positional and constrained positional models for four example sentences. For selected words, we also show the conditional probability $P(w_t | C_t)$ in parentheses.

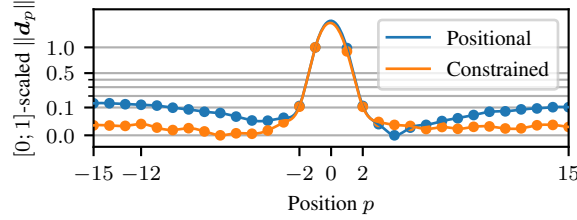


Figure 4: The importance of different positions p for predicting masked words in the positional and constrained positional models.

“ago”, “else”, and “cola”. The number of antepositional and postpositional features in the positional model is 76, which is close to the $D' = 60$ positional features selected for the constrained positional model by hyperparameter optimization. This indicates that the highest task performance is reached when only antepositional and postpositional features remain.

The informational features increase the importance of positions $p \notin [-2; 2]$. Levy, Goldberg [LG14] showed that context words with large input vectors have high self-information. We believe that the purpose of informational features is to amplify distant self-informational context words that indicate the general topic of a sentence. Context words whose importance is maximized by informational features include “finance”, “sports”, and “politics”. In the constrained positional model, informational features $d_{p,j}$ are effectively replaced by ones, which is close to what the positional model has learnt.

4.4 Language modeling

Figure 6 shows that the positional and constrained positional models consistently outperform the subword model during the training of RNN language models. Figure 6 also shows that RNN language models have converged and that training for more epochs would not have improved their perplexity.

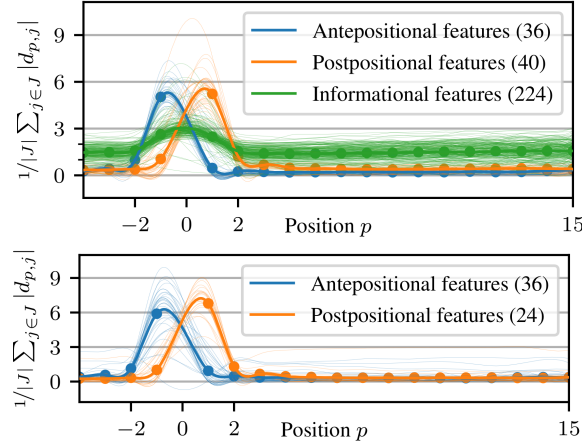


Figure 5: The importance of different positions p for predicting masked words in the positional (top) and constrained positional (bottom) models according to different clusters J of positional features. For each cluster J , we show its size $|J|$ in parentheses.

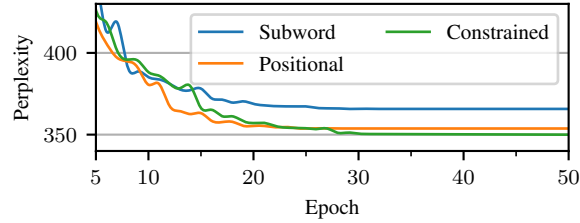


Figure 6: Validation perplexities at different epochs of RNN language models that use subword, positional, and constrained positional models as their lookup tables.

Table 3 shows that the constrained positional model produces word vectors that are better suited for initializing the lookup tables of RNN language models than the subword and positional models.

5 Conclusion and future work

In our work, we have related the attention mechanism from NMT to the positional language model of Mikolov et al. [Mik+18] and adapted the attention sparsification techniques of Zaheer et al. [Zah+20] to develop our constrained positional model.

	Subword	Positional	Constrained positional
Test perplexity	360.91	347.52	343.13

Table 3: Test perplexities of RNN language models that use subword, positional, and constrained positional models as lookup tables. Best result is **emphasized**.

We have shown that our constrained positional model is $2\times$ faster to train and more accurate at modeling English than the positional model. Future work should focus at the quantitative evaluation on language modeling in languages other than English and on other extrinsic NLP tasks, both alone and together with deep language models.

Furthermore, we have developed three novel qualitative evaluation measures and we used them to show that the positional vectors in English positional and constrained positional models serves two distinct roles: They allow the use of larger context sizes and they determine the grammatical properties of words. Future work should investigate the role of positional vectors in languages other than English.

Acknowledgements

The first author's work was graciously funded by the South Moravian Centre for International Mobility as a part of the Brno Ph.D. Talent project.

References

- [ASN21] Ayetiran, E. F., Sojka, P., Novotný, V. "EDS-MEMBEd: Multi-sense embeddings based on enhanced distributional semantic structures via a graph walk over word senses". In: *Knowledge-Based Systems* 219 (2021).
- [Bac12] Bach, K. "Context dependence (such as it is)". In: *The Continuum Companion to the Philosophy of Language*. Ed. by Garcia-Carpintero, M., Kolbel, M. 1st ed. Bloomsbury Publishing, 2012, pp. 153–184. ISBN: 9780826444066.
- [BB14] Botha, J. A., Blunsom, P. "Compositional Morphology for Word Representations and Language Modelling". In: *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32*. ICML'14. Beijing, China: JMLR.org, 2014, II–1899–II–1907.
- [BCB16] Bahdanau, D., Cho, K., Bengio, Y. *Neural machine translation by jointly learning to align and translate*. 2016. URL: <https://arxiv.org/pdf/1409.0473v7.pdf> (visited on 10/15/2021).
- [Boj+17] Bojanowski, P. et al. "Enriching word vectors with subword information". In: *Transactions of the Association for Computational Linguistics* 5 (2017), pp. 135–146. URL: <https://www.aclweb.org/anthology/Q17-1010.pdf>.
- [BPC20] Beltagy, I., Peters, M. E., Cohan, A. *Longformer: The long-document transformer*. 2020. URL: <https://arxiv.org/pdf/2004.05150v2.pdf> (visited on 10/15/2021).
- [CD17] Charlet, D., Damnati, G. "Simbow at SemEval-2017 task 3: Soft-cosine semantic similarity between questions for community question answering". In: *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. 2017, pp. 315–319. URL: <https://www.aclweb.org/anthology/S17-2051.pdf>.
- [CDL16] Cheng, J., Dong, L., Lapata, M. *Long short-term memory-networks for machine reading*. 2016. URL: <https://arxiv.org/pdf/1601.06733v7.pdf> (visited on 10/15/2021).
- [Chi+19] Child, R., Gray, S., Radford, A., Sutskever, I. *Generating long sequences with sparse transformers*. 2019. URL: <https://arxiv.org/pdf/1904.10509v1.pdf> (visited on 10/15/2021).

- [Cho+14a] Cho, K., Van Merriënboer, B., Bahdanau, D., Bengio, Y. *On the properties of neural machine translation: Encoder-decoder approaches*. 2014. URL: <https://arxiv.org/pdf/1409.1259v2.pdf> (visited on 10/15/2021).
- [Cho+14b] Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y. *Learning phrase representations using RNN encoder-decoder for statistical machine translation*. 2014. URL: <https://arxiv.org/pdf/1406.1078v3.pdf> (visited on 10/15/2021).
- [Cla+19] Clark, K., Khandelwal, U., Levy, O., Manning, C. D. *What does BERT look at? An analysis of BERT's attention*. 2019. URL: <https://arxiv.org/pdf/1906.04341v1.pdf> (visited on 10/15/2021).
- [CLS14] Chen, X., Liu, Z., Sun, M. “A unified model for word sense representation and disambiguation”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2014, pp. 1025–1035. URL: <https://www.aclweb.org/anthology/D14-1110.pdf>.
- [Dai+19] Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q., Salakhutdinov, R. “Transformer-XL: Attentive Language Models beyond a Fixed-Length Context”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 2978–2988. doi: 10.18653/v1/P19-1285. URL: <https://www.aclweb.org/anthology/P19-1285>.
- [Dev+18] Devlin, J., Chang, M.-W., Lee, K., Toutanova, K. *BERT: Pre-training of deep bidirectional transformers for language understanding*. 2018. URL: <https://arxiv.org/pdf/1810.04805v2.pdf> (visited on 10/15/2021).
- [GH12] Gutmann, M. U., Hyvärinen, A. “Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics”. In: *The Journal of Machine Learning Research, JMLR* 13.1 (2012), pp. 307–361. URL: <https://www.jmlr.org/papers/volume13/gutmann12a/gutmann12a.pdf>.
- [Gra+18] Grave, E., Bojanowski, P., Gupta, P., Joulin, A., Mikolov, T. *Learning word vectors for 157 languages*. 2018. URL: <https://arxiv.org/pdf/1802.06893v2.pdf> (visited on 10/15/2021).
- [IKS17] Inan, H., Khosravi, K., Socher, R. *Tying word vectors and word classifiers: A loss framework for language modeling*. 2017. URL: <https://arxiv.org/pdf/1611.01462v3.pdf> (visited on 10/15/2021).
- [Kah11] Kahneman, D. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011. 512 pp. ISBN: 978-0-374-27563-1.
- [Kus+15] Kusner, M., Sun, Y., Kolkin, N., Weinberger, K. “From Word Embeddings To Document Distances”. In: *International Conference on Machine Learning*. Ed. by Bach, F., Blei, D. Vol. 37. Proceedings of Machine Learning Research. Lille, France: PMLR, July 2015, pp. 957–966. URL: <http://proceedings.mlr.press/v37/kusnerb15.html>.
- [LG14] Levy, O., Goldberg, Y. “Neural word embedding as implicit matrix factorization”. In: *Advances in Neural Information Processing Systems* 27 (2014), pp. 2177–2185. URL: http://www.cs.columbia.edu/~blei/seminar/2016_discrete_data/readings/LevyGoldberg2014.pdf.
- [Li92] Li, W. “Random texts exhibit Zipf’s-law-like word frequency distribution”. In: *IEEE Transactions on information theory* 38.6 (1992), pp. 1842–1845. URL: <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=165464>.

- [Mik+13a] Mikolov, T., Chen, K., Corrado, G., Dean, J. *Efficient estimation of word representations in vector space*. 2013. URL: <https://arxiv.org/pdf/1301.3781v3.pdf> (visited on 10/15/2021).
- [Mik+13b] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., Dean, J. “Distributed Representations of Words and Phrases and their Compositionality”. In: *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2013, pp. 3111–3119.
- [Mik+18] Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., Joulin, A. “Advances in Pre-Training Distributed Word Representations”. In: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. 2018. URL: <http://www.lrec-conf.org/proceedings/lrec2018/pdf/721.pdf>.
- [Nov+20a] Novotný, V., Ayetiran, E. F., Štefánik, M., Sojka, P. *Text classification with word embedding regularization and soft similarity measure*. 2020. URL: <https://arxiv.org/abs/2003.05019v1> (visited on 10/15/2021).
- [Nov+20b] Novotný, V., Sojka, P., Štefánik, M., Lupták, D. “Three is Better than One. Ensembling Math Information Retrieval Systems”. In: *CEUR Workshop Proceedings*. Thessaloniki, Greece, 2020, p. 30. URL: http://ceur-ws.org/Vol-2696/paper_235.pdf.
- [Pet+06] Peters, E., Västfjäll, D., Slovic, P., Mertz, C., Mazzocco, K., Dickert, S. “Numeracy and decision making”. In: *Psychological science* 17.5 (2006), pp. 407–413. DOI: 10.1111/j.1467-9280.2006.01720.x.
- [Pin18] Pinelis, I. *The exp-normal distribution is infinitely divisible*. 2018. URL: <https://arxiv.org/pdf/1803.09838v1.pdf> (visited on 10/15/2021).
- [PMB19] Pérez, J., Marinković, J., Barceló, P. *On the turing completeness of modern neural network architectures*. 2019. URL: <https://arxiv.org/pdf/1901.03429v1.pdf> (visited on 10/15/2021).
- [Rav18] Ravshan, S. K. *Factor Analysis and Uniform distributions*. Visited on 2020-12-05. 2018. URL: <https://ravshansk.com/articles/uniform-distribution.html>.
- [Rec+11] Recht, B., Re, C., Wright, S., Niu, F. “Hogwild!: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent”. In: *Advances in Neural Information Processing Systems*. Ed. by Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K. Q. Vol. 24. Curran Associates, Inc., 2011, pp. 693–701.
- [ŘS10] Řehůřek, R., Sojka, P. “Software Framework for Topic Modelling with Large Corpora”. In: *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. Valletta, Malta: ELRA, May 2010, pp. 45–50. URL: <https://www.fi.muni.cz/usr/sojka/papers/lrec2010-rehurek-sojka.pdf>.
- [SVL14] Sutskever, I., Vinyals, O., Le, Q. V. “Sequence to sequence learning with neural networks”. In: *Advances in neural information processing systems* 27 (2014), pp. 3104–3112.
- [Vas+17] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., Polosukhin, I. “Attention is All you Need”. In: *Advances in Neural Information Processing Systems*. Ed. by Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. Vol. 30. Curran Associates, Inc., 2017, pp. 5998–6008. URL: <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.

[Wri08] Wright, D. E. “Do Words Have Inherent Meaning?” In: *ETC: A Review of General Semantics* 65.2 (2008), pp. 177–190. URL: <https://www.jstor.org/stable/42578827>.

[Yao+18] Yao, Z., Sun, Y., Ding, W., Rao, N., Xiong, H. “Dynamic word embeddings for evolving semantic discovery”. In: *Proceedings of the eleventh acm international conference on web search and data mining*. 2018, pp. 673–681. URL: <https://doi.org/10.1145/3159652.3159703>.

[Yun+20] Yun, C., Bhojanapalli, S., Rawat, A. S., Reddi, S. J., Kumar, S. *Are Transformers universal approximators of sequence-to-sequence functions?* 2020. URL: <https://arxiv.org/pdf/1912.10077v2.pdf> (visited on 10/15/2021).

[Zah+20] Zaheer, M., Guruganesh, G., Dubey, A., Ainslie, J., Alberti, C., Ontanon, S., Pham, P., Ravula, A., Wang, Q., Yang, L., Ahmed, A. *Big bird: Transformers for longer sequences*. 2020. URL: <https://arxiv.org/pdf/2007.14062v1.pdf> (visited on 10/15/2021).

[Zip32] Zipf, G. K. *Selective Studies and the Principle of Relative Frequency in Language*. Cambridge MA: Harvard University Press, 1932.

A Initialization of the positional model

In this appendix, we expand on Section 3.2 by describing three initialization options for the positional model. We discuss the properties of the initialization options and their practical effect on the training of the model.

A.1 Identity positions and vanilla subwords

To keep the effective learning rate of the positional model the same as in the subword model, it is sufficient to keep the distribution of the context vector \mathbf{u}_{C_t} the same as in the subword model. To achieve this, we initialize the input subword vectors \mathbf{u}_g as in (11) and the positional vectors \mathbf{d}_p to one. Intuitively, the training starts with no positional weighting and the positional vectors are learnt later. In practice, $\mathbf{d}_p \gg \mathbf{u}_g$, causing the gradient $\nabla_{\mathbf{u}_g} L$ to explode for $D > 600$ soon after the training has begun. This leads to numerical instability and the model parameters θ tend to NaN as the training continues.

A.2 Positions same as vanilla subwords

The simplest option is to initialize both the input subword vectors \mathbf{u}_g and the positional vectors \mathbf{d}_p as in (11). In practice, this decreases the variance of the context vector \mathbf{u}_{C_t} in the positional model compared to the subword model:

$$\begin{aligned} \mathbb{V}\text{ar}\left[\frac{1}{|C_t|} \sum_{\substack{w \in C_t \\ g \in G_w}} \mathbf{u}_g \odot \mathbf{d}_p\right] &= \frac{1}{|C_t|^2} \sum_{\substack{w \in C_t \\ g \in G_w}} \left(\mathbb{V}\text{ar}[\mathbf{u}_g \odot \mathbf{d}_p] = \mathbb{E}[\mathbf{u}_g^2] \odot \mathbb{E}[\mathbf{d}_p^2]\right) \\ &= \frac{1}{|C_t|^2} \sum_{\substack{w \in C_t \\ g \in G_w}} \frac{1}{9D^4} \ll \frac{1}{|C_t|^2} \sum_{\substack{w \in C_t \\ g \in G_w}} \frac{1}{3D^2} = \frac{1}{|C_t|^2} \sum_{\substack{w \in C_t \\ g \in G_w}} \mathbb{V}\text{ar}[\mathbf{u}_g] = \mathbb{V}\text{ar}\left[\frac{1}{|C_t|} \sum_{\substack{w \in C_t \\ g \in G_w}} \mathbf{u}_g\right]. \end{aligned} \quad (18)$$

See also Figure 7. The decrease in $\mathbb{V}\text{ar}[\mathbf{u}_{C_t}]$ decreases $\mathbb{V}\text{ar}[\nabla L]$ and therefore it also decreases the effective learning rate of the positional model. As D increases, the context vector \mathbf{u}_{C_t} quickly tends to zero due to $\mathbb{V}\text{ar}[\mathbf{u}_g \odot \mathbf{d}_p] = 1/9D^4$.

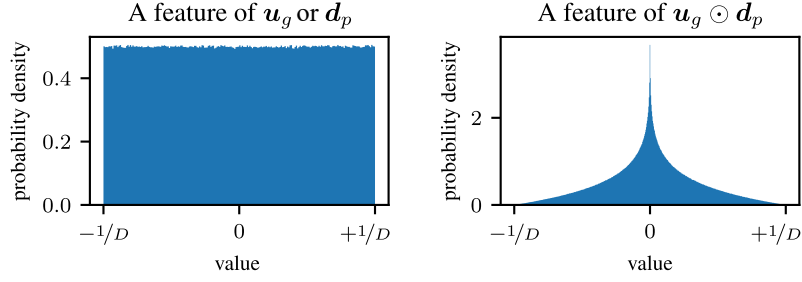


Figure 7: Probability densities of feature values in the input subword vectors \mathbf{u}_g , the positional vectors \mathbf{d}_p , and their product $\mathbf{u}_g \odot \mathbf{d}_p$ with the positions same as vanilla subwords initialization to $\mathcal{U}(\pm 1/D)$, $D = 1$. Since $\text{Var}[\mathbf{u}_g] \gg \text{Var}[\mathbf{u}_g \odot \mathbf{d}_p]$, the effective learning rate of the positional model is smaller than in the subword model.

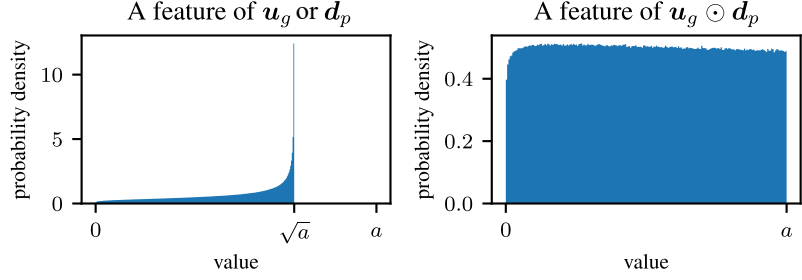


Figure 8: Probability densities of feature values in the input subword vectors \mathbf{u}_g , the positional vectors \mathbf{d}_p , and their product $\mathbf{u}_g \odot \mathbf{d}_p$ with the initialization to $\mathcal{U}^{0.5}(0, a)$, $a = 2$. For our use, we would need $\mathcal{U}^{0.5}(\pm 1/D)$ instead.

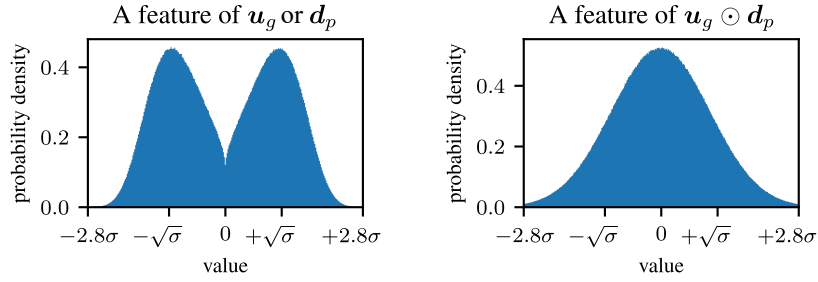


Figure 9: Probability densities of feature values in the input subword vectors \mathbf{u}_g , the positional vectors \mathbf{d}_p , and their product $\mathbf{u}_g \odot \mathbf{d}_p$ with the positions same as subwords initialization to the square-root normal distribution $\mathcal{N}^{0.5}(0, \sigma^2)$, $\sigma^2 = 1/3D^2$, $D = 1$.

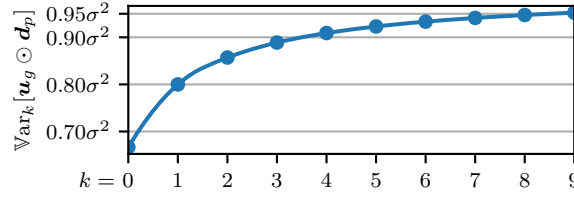


Figure 10: The variances $\text{Var}_k[\mathbf{u}_g \odot \mathbf{d}_p]$ for the feature values in the product $\mathbf{u}_g \odot \mathbf{d}_p$ of the input subword vectors \mathbf{u}_g and the positional vectors \mathbf{d}_p with the initialization to the square-root normal distribution $\mathcal{N}^{0.5}(0, \sigma^2)$, when we approximate the infinite sequence $(a_n)_{n=0}^\infty$ in the definition of $\mathcal{N}^{0.5}$ by first $k+1$ elements $(a_n)_{n=0}^k$.

A.3 Positions same as subwords

To keep the effective learning rate of the positional model the same as in the subword model and to avoid exploding gradients, it is sufficient to keep the distribution of the context vector \mathbf{u}_{C_t} the same as in the subword model and to initialize both the input subword vectors \mathbf{u}_g and the positional vectors \mathbf{d}_p from the same distribution. We could achieve this by initializing the features u_i of \mathbf{u}_g and the features d_j of \mathbf{d}_p as i.i.d. r.v.'s with the square-root distribution $\mathcal{U}^{0.5}(\pm 1/D)$ such that $u_i \cdot d_j \sim \mathcal{U}(\pm 1/D)$. Although an approximation of $\mathcal{U}^{0.5}(0, a)$ using the β -distribution is known [Rav18], see Figure 8, it does not extend to $\mathcal{U}^{0.5}(\pm 1/D)$, so we need another approach.

Assuming the context C_t is sufficiently large, then by the central limit theorem, the features of the context vector \mathbf{u}_{C_t} in the subword model have the normal distribution $\mathcal{N}(\mu, \sigma^2/|C_t|)$, where $\mu = \mathbb{E}[\mathcal{U}(\pm 1/D)] = 0$, $\sigma^2 = \text{Var}[\mathcal{U}(\pm 1/D)] = 1/3D^2$. To achieve the same distribution with the positional model, we initialize the features u_i of \mathbf{u}_g and the features d_j of \mathbf{d}_p as i.i.d. r.v.'s with some continuous distribution \mathcal{X} such that $\mathbb{E}[u_i \cdot d_j] = \mu$ and $\text{Var}[u_i \cdot d_j] = \sigma^2$. In our initialization, we use as \mathcal{X} the square-root normal distribution $\mathcal{N}^{0.5}(\mu, \sigma^2)$ of Pinelis [Pin18], see Figure 9. The continuous uniform $\mathcal{U}(\pm \sqrt[4]{3}/\sqrt{D})$ is also an option.

The definition of $\mathcal{N}^{0.5}$ by Pinelis [Pin18] contains an infinite sequence $(a_n)_{n=0}^\infty$:

$$\mathcal{N}^{0.5}(\mu, \sigma^2) = \epsilon \cdot e^{\sum_{n=0}^\infty a_n} \cdot \sqrt{\sigma} + \sqrt{\mu}, a_n = \frac{1}{4} \cdot \ln \left(1 + \frac{1}{\max(1, n)} \right) - \frac{G_n}{2n+1}, \quad (19)$$

where $(G_n)_{n=0}^\infty$ are i.i.d. r.v.'s with $\text{Gamma}(1/2, 1)$ distribution and ϵ is a Rademacher r.v. independent of all G_n . Since $\lim_{n \rightarrow \infty} a_n = 0$, we can approximate $\sum_{n=0}^\infty a_n$ by its first $k+1$ elements, but we need guarantees about $\mathbb{E}[u_i \cdot d_j]$ and $\text{Var}[u_i \cdot d_j]$. We can see that $\mathbb{E}[u_i \cdot d_j] = \mu$ for any k :

$$\begin{aligned} \mathbb{E}[u_i] &= \mathbb{E}[d_j] = \mathbb{E}[\epsilon] \cdot \mathbb{E}[e^{\sum_{n=0}^k a_n}] \cdot \mathbb{E}[\sqrt{\sigma}] + \mathbb{E}[\sqrt{\mu}] = 0 \cdot \mathbb{E}[e^{\sum_{n=0}^k a_n}] \\ &\quad \cdot \mathbb{E}[\sqrt{\sigma}] + \sqrt{\mu}, \mathbb{E}[u_i \cdot d_j] = \mathbb{E}[u_i] \cdot \mathbb{E}[d_j] = (\sqrt{\mu})^2 = \mu. \end{aligned} \quad (20)$$

We will denote the variance for a given k as Var_k :

$$\begin{aligned}\text{Var}_k[u_i] &= \text{Var}_k[d_j] = \text{Var}\left[\epsilon \cdot e^{\sum_{n=0}^k a_n} \cdot \sqrt{\sigma} + \sqrt{\mu}\right] \\ &= \left(\prod_{n=0}^k \mathbb{E}[e^{a_n}]^2 + \text{Var}\left[\prod_{n=0}^k e^{a_n}\right]\right) \cdot \sigma, \\ \text{Var}_k[u_i \cdot d_j] &= 2\mu \cdot \text{Var}_k[u_i] + \text{Var}_k[u_i]^2.\end{aligned}\tag{21}$$

In our initialization, we approximate $\sum_{n=0}^{\infty} a_n$ by $\sum_{n=0}^k a_n$, $k = 9$. As we show in Figure 10, this guarantees $\mathbb{E}[u_i \cdot d_j] = \mu$, $\text{Var}[u_i \cdot d_j]/\sigma^2 \in (0.95; 1]$.

B Qualitative evaluation measures

In this appendix, we expand on Section 3.4 by showing how the proposed qualitative evaluation measures relate to the conditional probability $\Pr(w_t \mid C_t)$ from (2). For a fixed set of negative samples N_{C_t} , $\Pr(w_t \mid C_t)$ is a strictly increasing transformation of the scoring function $s(w_t, C_t)$ from (3). Without loss of generality, our proofs will focus on $s(w_t, C_t)$ rather than on $\Pr(w_t \mid C_t)$.

B.1 Importance of positions

All else being constant, the ℓ_2 -norm $\|\mathbf{d}_p\|$ is an asymptotic upper bound on $|s(w_t, C_t)|$:

$$\begin{aligned}|s(w_t, C_t)| &= |\mathbf{u}_{C_t}^\top \cdot \mathbf{v}_{w_t}| = \frac{1}{|P|} \cdot |(\mathbf{u}_{w_{t+p}} \odot \mathbf{d}_p)^\top \mathbf{v}_{w_t} + \dots| \\ &\leq \frac{1}{|P|} \cdot (\|\mathbf{u}_{w_{t+p}}\| \cdot \|\mathbf{d}_p\| \cdot \|\mathbf{v}_{w_t}\| + \dots), |s(w_t, C_t)| \in \mathcal{O}(\|\mathbf{d}_p\|).\end{aligned}\tag{22}$$

All else being constant, $\sum_{j \in J} |d_{p,j}|$ is also an asymptotic upper bound on $|s(w_t, C_t)|$:

$$\begin{aligned}|s(w_t, C_t)| &= |\mathbf{u}_{C_t}^\top \cdot \mathbf{v}_{w_t}| = \frac{1}{|P|} \cdot \left| \sum_{j \in J} u_{w_{t+p},j} \cdot d_{p,j} \cdot v_{w_t,j} + \dots \right| \\ &\leq \frac{1}{|P|} \cdot \left(\sum_{j \in J} |u_{w_{t+p},j}| \cdot |d_{p,j}| \cdot |v_{w_t,j}| + \dots \right), |s(w_t, C_t)| \in \mathcal{O}\left(\sum_{j \in J} |d_{p,j}|\right),\end{aligned}\tag{23}$$

where $u_{w_{t+p},j}$ are features of the input vector $\mathbf{u}_{w_{t+p}}$ for the context word w_{t+p} at position p and $v_{w_t,j}$ are features of the output vector \mathbf{v}_{w_t} for the masked word w_t .

B.2 Importance of context words

All else being constant, $\sum_{j \in J} |u_{w,j}|$ is an asymptotic upper bound on the expected absolute difference $\mathbb{E}|s(w_t, C_t) - s(w_t, C'_t)|$ for a fixed context word w and random-valued masked words w_t and contexts C_t, C'_t , where w is at position p_1 in C_t and at

position p_2 in C'_t :

$$\begin{aligned}
& \mathbb{E}|s(w_t, C_t) - s(w_t, C'_t)| \\
&= \frac{1}{|P|} \cdot \sum_{w_t, C_t, C'_t} \left| \sum_{j \in J} u_{w,j} \cdot (d_{p_1,j} - d_{p_2,j}) \cdot v_{w_t,j} + \dots \right| \cdot \Pr(w_t, C_t, C'_t) \\
&\leq \frac{1}{|P|} \cdot \sum_{w_t, C_t, C'_t} |u_{w,j}| \cdot |(d_{p_1,j} - d_{p_2,j}) \cdot v_{w_t,j}| \cdot \Pr(w_t, C_t, C'_t) + \dots \cdot \Pr(w_t, C_t, C'_t), \\
&\mathbb{E}|s(w_t, C_t) - s(w_t, C'_t)| \in \mathcal{O}\left(\sum_{j \in J} |u_{w,j}|\right). \tag{24}
\end{aligned}$$