

On-Device Text Representations Robust To Misspellings via Projections

Chinnadhurai Sankar
Mila, Université de Montréal*
Montreal, QC, Canada
chinnadhurai@gmail.com

Sujith Ravi
Amazon Alexa†
Sunnyvale, CA, USA
sravi@sravi.org

Zornitsa Kozareva
Google
Mountain View, CA, USA
zornitsa@kozareva.com

Abstract

Recently, there has been a strong interest in developing natural language applications that live on personal devices such as mobile phones, watches and IoT with the objective to preserve user privacy and have low memory. Advances in Locality-Sensitive Hashing (LSH)-based projection networks have demonstrated state-of-the-art performance in various classification tasks without explicit word (or word-piece) embedding lookup tables by computing on-the-fly text representations.

In this paper, we show that the projection based neural classifiers are inherently robust to misspellings and perturbations of the input text. We empirically demonstrate that the LSH projection based classifiers are more robust to common misspellings compared to BiLSTMs (with both word-piece & word-only tokenization) and fine-tuned BERT based methods. When subject to misspelling attacks, LSH projection based classifiers had a small average accuracy drop of 2.94% across multiple classifications tasks, while the fine-tuned BERT model accuracy had a significant drop of 11.44%.

1 Introduction

At the core of Natural Language Processing (NLP) neural models are pre-trained word embeddings like Word2Vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014a) and ELMo (Peters et al., 2018). They help initialize the neural models, lead to faster convergence and have improved performance for numerous application such as Question Answering (Liu et al., 2018), Summarization (Cheng and Lapata, 2016), Sentiment Analysis (Yu et al., 2017). While word embeddings are powerful in unlimited constraints such as computation power

and compute resources, it becomes challenging to deploy them to on-device due to their huge size.




Figure 1: One-hot word embedding look-up vectors vs linear combination of LSH projection based vectors (Ravi, 2017) representing the same word.

This led to interesting research by (Ravi and Kozareva, 2018; Sankar et al., 2019), who showed that word embeddings can be replaced with lightweight binary Locality-Sensitive Hashing (LSH) based projections learned on-the-fly. The projection approach surmounts the need to store any embedding matrices, since the projections are dynamically computed. This further enables user *privacy* by performing inference directly on device without sending user data (e.g., personal information) to the server. The embedding memory size is reduced from $O(V)$ to $O(K)$, where V is the token vocabulary size and $K \ll V$, is the binary LSH projection size. The projection representations can operate on either word or character level, and can be used to represent a sentence or a word depending on the NLP application. For instance, recently the Projection Sequence Networks (ProSeq) (Kozareva and Ravi, 2019) used BiLSTMs over word-level projection representations to represent long sentences and achieved close to state-of-the-art results in both short and long text classification tasks with varying amounts of supervision and vocabulary sizes.

*Work done during internship at Google

†Work done while at Google AI




Figure 2: Memory for V look-up vectors for each token vs storing $K (\ll V)$ vectors and linearly combining them for token representation. We consider $K = 1120$ following (Ravi and Kozareva, 2018) in this paper.

Despite being successful, there are no existing systematic research efforts focusing on evaluating the capabilities of the LSH based projection for text representations. To that end, we empirically analyze the effectiveness and robustness of the LSH projection approach for text representation by conducting two types of studies in this paper.

1. *Classification with perturbed inputs*, where we show that Projection based networks 1) Projection Sequence Networks (ProSeqo) (Kozareva and Ravi, 2019) and 2) Self-Governing Neural Networks (SGNN) models (Ravi and Kozareva, 2019) evaluated with perturbed LSH projections are robust to misspellings and transformation attacks, while we observe significant drop in performance for BiLSTMs and fine-tuned BERT classifiers.
2. *Perturbation Analysis*, where we test the robustness of the projection approach by directly analyzing the changes in representations when the input words are subject to the char misspellings. The purpose of this study is to examine if the words or sentences with misspelling are nearby in the projection space instead of frequently colliding with the projection representations of other valid words.

Overall, our studies showcase the robustness of LSH projection representations and resistance to misspellings. Due to their effectiveness, we believe that in the future, text representations using LSH projections can go beyond memory constrained settings and even be exploited in large scale models like Transformers (Vaswani et al., 2017).

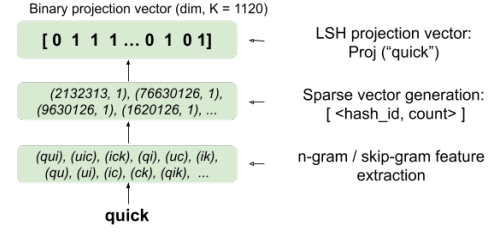


Figure 3: Binary Locality-Sensitive Hashing (LSH) projection representation for text.

2 Binary LSH projections for text representations

The dependency on vocabulary size V , is one of the primary reasons for the huge memory footprint of embedding matrices. It is common to represent a token, x by one-hot representation, $\mathbb{Y}(x) \in [0, 1]^V$ and a distributed representation of the token is obtained by multiplying the one-hot representation with the embedding matrix, $W_V \in \mathbb{R}^{d \times V}$ as in

$$U_V(x) = W_V * \mathbb{Y}(x)^\top \in \mathbb{R}^d$$

One way to remove the dependency on the vocabulary size is to learn a smaller matrix, $W_K \in \mathbb{R}^{d \times K}$ ($K \ll V$), as shown in Figure 2. For instance, 300-dimensional Glove embeddings, W_V (Pennington et al., 2014b) with 400k vocabulary size occupies > 1 GB while the W_K occupies only ≈ 1.2 MB for $K = 1000$ yielding a $1000\times$ reduction in size. Instead of learning a unique vector for each token in the vocabulary, we can think of the columns of this W_K matrix as a set of basis vectors and each token can be represented as a linear combination of basis vectors in W_K as in Figure 1. We select the basis vectors from W_K for each token with a fixed K -bit binary vector instead of a V -bit one-hot vector.

The LSH *Projection function*, \mathbb{P} (Figure 3)(Ravi, 2017, 2019) used in SGNN (Ravi and Kozareva, 2018) and ProSeqo (Kozareva and Ravi, 2019) does exactly this as it dynamically generates a fixed binary projection representation, $\mathbb{P}(x) \in [0, 1]^K$ for any token, x by extracting morphological input features like char (or token) n-gram & skip-gram features, parts of speech tags etc. from x and a modified Locality-Sensitive Hashing (LSH) based transformation, \mathbb{L} as in

$$x \xrightarrow{\mathbb{F}} [f_1, \dots, f_n] \xrightarrow{\mathbb{L}} \mathbb{P}(x) \in [0, 1]^K$$

where \mathbb{F} extracts n-grams (or skip-grams), $[f_1, \dots, f_n]$ from the input text. Here, $[f_1, \dots, f_n]$

could refer to either character level or token level n-grams(or skip-grams) features. Given the LSH projection representation, $\mathbb{P}(x)$, the distributed representation of the token, x is represented as in

$$U_W(x) = W_K * \mathbb{P}(x)^\top \in \mathbb{R}^d$$

. It is worth noting that projection operation, \mathbb{P} can also be used map an entire sentence directly to the $[0, 1]^K$ space.

As for the projection based classifiers, the ProSeqo model (Kozareva and Ravi, 2019) runs a BiLSTM over word-level binary LSH projection representations to predict the correct classes, while the SGNN model (Ravi and Kozareva, 2018) computes a binary LSH projection representation for the entire input text, followed by a 2-layer MLP and a softmax layer on top of it for class prediction. SGNN was designed for short text, while ProSeqo is also suitable for long text classification tasks.

There have been a number of research efforts (Sakaguchi et al., 2016; Edizel et al., 2019; Pruthi et al., 2019) to improve the robustness of neural classifiers to misspelling attacks and other text transformations. Recently, Pruthi et al. (2019) observe that fine-tuned BERT and BiLSTM based models are very brittle (for e.g., accuracy drops from 90.3% to 45.8% in the SST (Socher et al., 2013) classification task) to adversarial misspelling attacks. Contrary to intuition, they observe that word-piece and character-level models are more susceptible to spelling attacks compared to the word-level models.

The LSH projection operation, \mathbb{P} , is a function of n-grams (and skip-grams) of the input text x and usually the fraction of n-grams affected by spelling attacks tend to be minimal resulting in insignificant changes to the projection representation, $\mathbb{P}(x)$. Therefore, we hypothesize that the projection based models like ProSeqo, SGNN, etc. are *inherently* robust to commonly occurring spelling attacks. In the following sections, we investigate the robustness of projection based classifiers by subjecting them to common misspellings, followed by an analysis of changes in the binary LSH projections of input text under such transformations.

3 Effect of Misspellings on Text Classification

We study the robustness of two types of projection based models – ProSeqo and SGNN. On the other hand, we fine-tune the pretrained BERT-base model

(with word-piece tokenization) (Devlin et al., 2018) and train two-layer BiLSTMs (with both word-only and word-piece tokenization) for comparable accuracies with respect to the projection based models for a fair comparison. By word-only tokenization, we mean that models encode input words using a lookup table for each word. In our setup, we test the robustness of the neural classifiers by subjecting the corresponding test sets to common misspellings and omissions. We consider the following perturbation operations: randomly dropping, inserting, and swapping internal characters within words of the input sentences (Gao et al., 2018; Pruthi et al., 2019)¹. We decide to perturb each word in a sentence with a fixed probability, $P_{perturb}$. Following (Ravi and Kozareva, 2018), we fix the projection dimension to $K = 1120$.

3.1 Datasets

For evaluation purposes, we use the following text classification datasets for dialog act classification MRDA (Shriberg et al., 2004) and SWDA (Godfrey et al., 1992; Jurafsky et al., 1997), for intent prediction ATIS (Tür et al., 2010) and long text classification Amazon Reviews (Zhang et al., 2015) and Yahoo! Answers (Zhang et al., 2015). Table 1 shows the characteristics of each dataset.

Tasks	# Classes	Avg-len	Train	Test
ATIS (Dialog act)	21	11	4.4k	0.89k
MRDA (Dialog act)	6	8	78k	15k
SWDA (Intent Prediction)	42	7	193k	5k
YAHOO (Answers Categorization)	10	108	1400k	60k
AMAZON (Review Prediction)	5	92	3000k	650k

Table 1: Classification Dataset Characteristics

Datasets → Models ↓	Accuracy drop (%) ± std-deviation (over 5 runs)			
	MRDA	ATIS	YAHOO	AMAZON
BERT-base	8.25±3.4	15.57±1.2	9.06±5.4	12.88±3.8
BiLSTM-wp	8.91±3.9	20.11±3.1	11.32±4.5	9.46±1.3
BiLSTM-w	15.14±4.3	16.23±2.6	9.32±2.3	8.88±2.4
SGNN	1.91±0.5	2.80±0.3	-	-
ProSeqo	2.11±0.4	2.84±0.6	3.11±0.3	3.91±0.5

Table 2: Each entry in the table denotes the average drop in accuracy(%) when classifiers are subject to test inputs with misspellings, $P_{perturb} = 0.2$. ATIS and MRDA are smaller datasets, while Yahoo! Answers and Amazon reviews are larger. BERT-base refers to the BERT-base (Devlin et al., 2018) word-piece model fine-tuned to individual tasks and it converged to state-of-the-art for all tasks.

¹Further details on the perturbation operations and training details necessary for reproducibility are presented in the supplementary material

Accuracy(%) (Averaged over 5 runs)									
Datasets	MRDA			ATIS			SWDA		
Perturb(%)	BiLSTM-wp	BiLSTM-w	SGNN	BiLSTM-wp	BiLSTM-w	SGNN	BiLSTM-wp	BiLSTM-w	SGNN
0	79.23	78.14	87.22	91.73	92.04	93.51	72.92	72.94	76.21
Perturbation operation: <i>drop</i>									
20	69.46 \pm 1.1	74.96 \pm 0.7	85.43 \pm 0.3	81.95 \pm 2.2	80.15 \pm 2.0	91.05 \pm 0.2	65.05 \pm 4.2	64.95 \pm 4.1	70.76 \pm 1.1
40	72.34 \pm 2.9	55.17 \pm 3.1	84.62 \pm 0.22	71.69 \pm 2.9	65.88 \pm 2.1	91.86 \pm 0.3	61.24 \pm 4.1	64.95 \pm 5.2	67.79 \pm 1.2
60	69.81 \pm 4.7	42.25 \pm 4.2	83.27 \pm 0.25	59.16 \pm 4.8	56.25 \pm 3.9	90.12 \pm 0.4	57.48 \pm 5.7	58.77 \pm 5.8	63.21 \pm 1.3
Perturbation operation: <i>swap</i>									
20	78.25 \pm 1.1	71.34 \pm 2.2	86.74 \pm 0.1	86.12 \pm 1.7	85.05 \pm 1.8	92.05 \pm 0.3	66.27 \pm 3.3	64.52 \pm 2.1	70.84 \pm 0.3
40	75.91 \pm 3.9	69.22 \pm 2.1	86.39 \pm 0.2	82.06 \pm 2.7	78.04 \pm 2.9	91.15 \pm 0.2	62.67 \pm 5.8	54.93 \pm 2.4	67.22 \pm 0.3
60	69.22 \pm 3.8	66.91 \pm 4.0	85.99 \pm 0.2	72.34 \pm 3.3	68.54 \pm 4.1	91.27 \pm 0.3	59.20 \pm 4.8	47.93 \pm 4.3	64.48 \pm 0.3
Perturbation operation: <i>all</i>									
20	72.96 \pm 1.3	73.39 \pm 2.3	86.71 \pm 0.4	80.40 \pm 1.7	83.55 \pm 2.1	92.83 \pm 0.2	60.49 \pm 4.3	64.28 \pm 3.3	68.96 \pm 0.2
40	70.32 \pm 3.4	63.04 \pm 4.3	85.31 \pm 0.5	71.62 \pm 3.1	75.81 \pm 2.6	90.71 \pm 0.3	54.96 \pm 4.1	59.46 \pm 4.6	65.44 \pm 0.4
60	67.64 \pm 5.7	55.50 \pm 5.3	84.21 \pm 0.5	61.10 \pm 6.1	66.58 \pm 5.6	88.35 \pm 0.3	49.62 \pm 6.7	51.85 \pm 6.3	64.97 \pm 0.5

Table 3: Comparison of projection based models vs BiLSTMs subject to various types and amounts of perturbations. BiLSTM-wp and BiLSTM-w refer to models with word-piece and word-only tokenization respectively.

3.2 Experiments and Results

Table 2 reports the average classifier accuracy drops when all the models are subject to all types of perturbations (*swap*, *drop*, & *add*) on multiple classification tasks (two short text and two long text). We see that the accuracy drop for the projection based models is significantly lower across all datasets. It is also worth noting that the standard deviations across the 5 runs are also minimal for the projection based models further showcasing the stability of projection representations.

In another experiment shown in Table 3, we subject different models to varying types and amounts of perturbations. Similarly, we see that the accuracy drop for the projection based models is the smallest across all datasets and amounts of perturbation. Compared to the word-only models, we observe that the word-piece models are also comparably susceptible to character perturbations which agrees with the findings in (Pruthi et al., 2019).

4 Perturbation Analysis

Apart from the classification experiments, we also directly analyze the changes in the binary LSH projection representations by subjecting input text to different types and amount of perturbations. To that end, we take a large corpus *enwik9*² (vocabulary size of 500k and 129M words) to analyze the average Hamming distance between LSH projections of the words in the corpus. Next, we compute the average changes in the projection representations by subjecting them to the character perturbations from Section 3. Table 4 shows the results. We make the

²*enwik9* is a byte-level dataset consisting of the first 10⁹ bytes of the English Wikipedia XML dump, <http://mattmahoney.net/dc/textdata.html>.

LSH Proj.Dim (K)	Character Perturbations	
	5%	10%
840	10.08	24.33
980	15.48	31.24
1120	18.83	33.65
1260	19.71	39.01

Table 4: Avg. changes in word projections (bits) for different Char Perturbation % in *enwik9* corpus.

following observations from our experiments:

1. Average Hamming distance between LSH projections of words is $\approx K/2$, where K is the projection dimension which implies that the words are more or less uniformly spread out from each other indicating that there are no bias issues in the $[0, 1]^K$ representation space.
2. Assuming $P_{perturb} = 0.2$, we observe that LSH projection changes only by $\approx 11\%$ w.r.t the average Hamming distance between the words in the corpus when subject to misspellings. For instance, if the average Hamming distance between LSH projections of words is 100 bits, misspellings change the projections by only 11 bits on average. Intuitively, this suggests that neural layers on top of the LSH projection tend to rarely confuse a misspelled word for another valid word.

Also from Table 4, we found that the changes in the LSH-projection, $\Delta_{\mathbb{P}(x)}$ due to perturbations is directly proportional to LSH projection dimension, K and perturbation probability, $P_{perturb}$ as in, $\Delta_{\mathbb{P}(x)} \propto K \cdot P_{perturb}$.

5 Conclusion

In this work, we perform a detailed study analyzing the robustness of recent LSH-based projection

neural networks for memory-efficient text representations. Based on multiple text classification tasks and perturbation studies, we find projection-based neural models to be robust to text transformations compared to BERT or BiLSTMs with embedding lookup tables for words and word-pieces.

References

- Jianpeng Cheng and Mirella Lapata. 2016. [Neural summarization by extracting sentences and words](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 484–494. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Bora Edizel, Aleksandra Piktus, Piotr Bojanowski, Rui Ferreira, Edouard Grave, and Fabrizio Silvestri. 2019. [Misspelling oblivious word embeddings](#). *CoRR*, abs/1905.09755.
- Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. [Black-box generation of adversarial text sequences to evade deep learning classifiers](#). In *2018 IEEE Security and Privacy Workshops, SP Workshops 2018, San Francisco, CA, USA, May 24, 2018*, pages 50–56.
- John J. Godfrey, Edward C. Holliman, and Jane McDaniel. 1992. [Switchboard: Telephone speech corpus for research and development](#). In *Proceedings of the 1992 IEEE International Conference on Acoustics, Speech and Signal Processing - Volume 1, ICASSP'92*, pages 517–520. IEEE Computer Society.
- Daniel Jurafsky, Rebecca Bates, Rachel Martin Noah Coccaro, Marie Meteer, Klaus Ries, Elizabeth Shriberg, Audreas Stolcke, Paul Taylor, and Van Ess-Dykema. 1997. [Automatic detection of discourse structure for speech recognition and understanding](#). In *Proceedings of IEEE Workshop on Automatic Speech Recognition and Understanding*, pages 88–95.
- Zornitsa Kozareva and Sujith Ravi. 2019. [ProSeqo: Projection sequence networks for on-device text classification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3892–3901.
- Xiaodong Liu, Yelong Shen, Kevin Duh, and Jianfeng Gao. 2018. [Stochastic answer networks for machine reading comprehension](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1694–1704. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). In *International Conference on Learning Representations, Workshop Track*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014a. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014b. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- 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. Association for Computational Linguistics.
- Danish Pruthi, Bhuwan Dhingra, and Zachary C. Lipton. 2019. [Combating adversarial misspellings with robust word recognition](#). *CoRR*, abs/1905.11268.
- Sujith Ravi. 2017. [Projectionnet: Learning efficient on-device deep networks using neural projections](#). *CoRR*, abs/1708.00630.
- Sujith Ravi. 2019. [Efficient on-device models using neural projections](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5370–5379.
- Sujith Ravi and Zornitsa Kozareva. 2018. [Self-governing neural networks for on-device short text classification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 804–810.
- Sujith Ravi and Zornitsa Kozareva. 2019. [On-device structured and context partitioned projection networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3784–3793, Florence, Italy. Association for Computational Linguistics.
- Keisuke Sakaguchi, Kevin Duh, Matt Post, and Benjamin Van Durme. 2016. [Robust word recognition via semi-character recurrent neural network](#). *CoRR*, abs/1608.02214.
- Chinnadhurai Sankar, Sujith Ravi, and Zornitsa Kozareva. 2019. [Transferable neural projection representations](#). 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 3355–3360, Minneapolis, Minnesota. Association for Computational Linguistics.

Elizabeth Shriberg, Rajdip Dhillon, Sonali Bhagat, Jeremy Ang, and Hannah Carvey. 2004. The ICSI meeting recorder dialog act (MRDA) corpus. In *Proceedings of the SIGDIAL 2004 Workshop, The 5th Annual Meeting of the Special Interest Group on Discourse and Dialogue, April 30 - May 1, 2004, Cambridge, Massachusetts, USA*, pages 97–100.

Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, Christopher Potts, et al. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*, volume 1631, page 1642.

Gökhan Tür, Dilek Hakkani-Tür, and Larry P. Heck. 2010. What is left to be understood in atis? In *Proceedings of 2010 IEEE Spoken Language Technology Workshop (SLT)*, pages 19–24.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.

Liang-Chih Yu, Jin Wang, K. Robert Lai, and Xuejie Zhang. 2017. [Refining word embeddings for sentiment analysis](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 534–539. Association for Computational Linguistics.

Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*, pages 649–657.