

Towards Lossless Encoding of Sentences

Gabriele Prato* Mathieu Duchesneau Sarath Chandar Alain Tapp

Mila, Université de Montréal

Abstract

A lot of work has been done in the field of image compression via machine learning, but not much attention has been given to the compression of natural language. Compressing text into lossless representations while making features easily retrievable is not a trivial task, yet has huge benefits. Most methods designed to produce feature rich sentence embeddings focus solely on performing well on downstream tasks and are unable to properly reconstruct the original sequence from the learned embedding. In this work, we propose a near lossless method for encoding long sequences of texts as well as all of their sub-sequences into feature rich representations. We test our method on sentiment analysis and show good performance across all sub-sentence and sentence embeddings.

1 Introduction

Compressing information by encoding it into a fixed size representation in such a way that perfect decoding is possible is challenging. Instead, most of the existing sentence encoding methods focus more on learning encoding such that the encoded representations are good enough for the downstream tasks. In this work, we focus on perfectly decodable encoding of sentences which will be very useful in designing good generative models that can generate longer sentences.

Early efforts such as (Hinton and Salakhutdinov, 2006) have shown autoencoders to effectively yield compressed input representations. Pollack (1990) was the first to propose using autoencoders recursively. Such models have been shown to be useful for a multitude of tasks. Luong et al. (2013) use recursive neural networks and neural language models to better represent rare words

via morphemes. Socher et al. (2011a) use recursive autoencoders for paraphrase detection, learning sentence embeddings (Socher et al., 2010) and syntactic parsing. Socher et al. (2011b) also use a recursive autoencoder to build a tree structure based on error reconstruction. Additionally, Socher et al. (2012) use a matrix-vector RNN to learn semantic relationships present in natural language and show good performance on such task as well as sentiment classification. Then, Socher et al. (2013) introduced the Recursive Neural Tensor Network, trained on a their proposed Sentiment Treebank corpus to better deal with negating sub-sequences for better sentiment classification. Recently, Kokkinos and Potamianos (2017) proposed Structural Attention to build syntactic trees and improve even further performance on SST. Parse trees do alleviate the burden of learning the syntactic structure of text, but these methods limit the number of generated embeddings to the number of nodes in the parse tree. Our proposed method does not have such a restriction as all possible syntactic tree can be simultaneously represented by the architecture.

Convolutional Neural Networks (LeCun et al., 1989) have been used in natural language processing as well. Convolutions work well for extracting low and high level text features and building sequence representations. Lai et al. (2015) proposed to use CNNs recurrently and show good performance on various language tasks. Zhang et al. (2015); Dos Santos and Gatti de Bayser (2014) both train CNNs on character level for sentiment analysis, while Johnson and Zhang (2014) work on word level. Kalchbrenner et al. (2014) propose a Dynamic Convolutional Neural Network for semantic modelling of sentences and apply their model to sentiment prediction. Our proposed model is very similar to 1D CNNs. In our case though, we use a multilayer perceptron in parallel

*Corresponding author: gabriele.prato@umontreal.ca

instead of a kernel to extract meaningful information out of the layer’s input.

Much progress has been made in recent years in the field of general purpose sentence embeddings. Fixed length representations of sentence wide context are learned with the objective of serving for a wide range of downstream tasks. [Conneau et al. \(2017\)](#) trained a bidirectional LSTM on the AllNLI natural language inference corpus ([Bowman et al., 2015](#); [Williams et al., 2017](#)) producing embeddings that generalized well on the SentEval ([Conneau and Kiela, 2018](#)) benchmark. Following this trend, [Subramanian et al. \(2018\)](#) trained a GRU ([Cho et al., 2014](#)) on Skip-thought vectors ([Kiros et al., 2015](#)), neural machine translation, parsing and natural language inference to get even better downstream task results. More recently, [Devlin et al. \(2018\)](#); [Liu et al. \(2019b,a\)](#) use Transformers ([Vaswani et al., 2017](#)) to produce sentence wide context embeddings for each input token and get state-of-the-art results on multiple natural language processing tasks. [Dai et al. \(2019\)](#) improve the Transformer method by recursively applying it to fixed length segments of text while using a hidden state to model long dependencies. One downside to these sentence embedding generation methods is that the context is always sequence wide. Our proposed model computes a sentence embedding as well as an embedding for all possible sub-sentences of the sequence with sub-sentence wide context only. All embeddings generated throughout our architecture are constructed the same way and thus share the same properties.

2 Recursive Autoencoder

We introduce our recursive autoencoding approach in this section. First we define our model’s architecture and how each encoding and decoding recursion is performed. We then describe how the model keeps track of the recursion steps, followed by a description of how the input is represented. We also explain the advantages of using the mean squared error loss for our method. Finally, we dive into the implementation details.

2.1 Model Architecture

Our model is a recursive auto-encoder. Figure 1 shows an example of our architecture for a sequence of length three.

The encoder takes an input sequence




Figure 1: Example of our recursive autoencoder with an input sequence of length three. The encoder recursively takes two embeddings and outputs one until a single one is left and the decoder takes one embedding and outputs two until there are as many as in the original sequence.

$\{x_1, \dots, x_n\}$, where n is the sequence length of the layer’s input, and outputs a sequence $\{y_1, \dots, y_{n-1}\}$. The same $\{y_1, \dots, y_{n-1}\}$ is then used as input for the next recursion until the output sequence contains only a single element y_1 , the sentence embedding. The recursion performs the following operation:

$$y_i = \text{MLP}_{enc}([x_i; x_{i+1}]) \forall i \in \{1, \dots, n-1\} \quad (1)$$

where MLP_{enc} is a shared multilayer perceptron and $[x_i; x_{i+1}]$ is the concatenation of the embeddings x_i and x_{i+1} . MLP_{enc} is shared throughout all of the encoding recursion steps.

For decoding, it is the inverse procedure of recursively transforming an input sequence $\{x_1, \dots, x_n\}$ into an output sequence $\{y_1, \dots, y_{n+1}\}$:

$$[y_i; y'_{i+1}] = \text{MLP}_{dec}(x_i) \forall i \in \{1, \dots, n\} \quad (2)$$

where MLP_{dec} is the shared multilayer perceptron used by all decoding recursive steps and $[y_i; y'_{i+1}]$ is an embedding twice the size of x_i , which we then split into two embeddings y_i and y'_{i+1} , each of the same size as x_i . Since we obtain two embeddings y_i and y'_{i+1} for each x_i , we will have the following embeddings: $y_1, \{y_2, \dots, y_n\}, \{y'_2, \dots, y'_n\}$ and y'_{n+1} . We merge the overlapping sets by computing the mean:

$$y_i = \frac{y_i + y'_i}{2} \forall i \in \{2, \dots, n\} \quad (3)$$

and set $y_{n+1} = y'_{n+1}$. We now have a single set of embeddings $\{y_1, \dots, y_{n+1}\}$. Both *max* and *mean* functions gave similar results, hence we stick with *mean* throughout all experiments. The output embeddings are then used as input for the next decoding recursion until we get as many elements as the original input sequence.

2.2 Step Encoding

To help the recursive autoencoder keep track of the number of recursive steps which were applied to an embedding, we concatenate to the input of MLP_{enc} the number of the current recursive step as a scalar, starting from 1 for the first recursion, as well as a one-hot of that scalar with custom bucket sizes: $\{1, 2, 3-4, 5-7, \dots\}$. All buckets after 5-7 are also of size 3. We found this combination of both scalar and one-hot to give best results. When decoding, we also concatenate to the input of MLP_{dec} this scalar and one-hot, but instead of increasing our recursive step count, we subtract one to it after each recursive decoding step.

2.3 Input Representation

We use uncased GloVe embeddings (Pennington et al., 2014) of size 300 to represent the initial input sequence words, which are then passed through a learned resizing multilayer perceptron (MLP_{in}) before given as input to the encoder. The output of the decoder is also passed through a different learned resizing multilayer perceptron (MLP_{out}) to get back to the GloVe embedding size. We use a vocabulary of 337k words throughout all tasks.

2.4 Mean Squared Error

To compute the loss between input GloVe embeddings and the output embeddings, we use the mean squared error (MSE) loss. Obtaining an MSE of 0 would mean our method is lossless, which would not necessarily be the case with the cross entropy loss. MSE also allows us to work with a vocabulary far larger than what is usually the case, as the common classification layer plus cross entropy loss setup tends to have issues with large vocabularies.

2.5 Implementation Details

The two embeddings given as input to MLP_{enc} are each of size d_{emb} , as is also its output embedding. Same for MLP_{dec} , the input embedding is of size d_{emb} and the two output embeddings are each of

size d_{emb} . Both multilayer perceptrons have one hidden layer of size $\frac{2}{3}d_{emb}$, halfway between the input and output size. We apply LayerNorm (Lei Ba et al., 2016) on the output of each layers of the MLPs, followed by a ReLU activation. The input and output resizing modules MLP_{in} and MLP_{out} also have one hidden layer halfway the size of their input and output. They also use ReLU activations, except for MLP_{out} 's last layer. No LayerNorm is used in these resizing components. We test four different d_{emb} embedding sizes in section 3.1.

3 Experiments

In this section, we first present the autoencoding results. Then we present the results on sentiment analysis using our sentence encoding on the Stanford Sentiment Treebank dataset (Socher et al., 2013).

3.1 Autoencoding

As a first experiment, we tested our model on the autoencoding task. Training was done on the BookCorpus (Zhu et al., 2015) dataset, comprising eleven thousand books and almost one billion words. At test time, we measured accuracy by computing the MSE distance between an output embedding and the entire vocabulary. We count an output embedding as "correct" if the closest embedding out of all the vocabulary of size 337k is its corresponding input embedding.

For the autoencoder, we tried four embedding sizes: 300, 512, 1024 and 2048. In all cases, models are given GloVe embeddings of size 300 as input. They also all output embeddings of size 300. Reconstruction accuracy is shown for different sequence lengths in Figure 2. With an embedding size of 2048, the model is able to reproduce near perfectly sequences of up to 40 tokens. Longer sentences aren't able to do better and have on average 39 correct tokens. This results in model accuracy linearly going down after a certain threshold, as can be seen in Figure 2.

To demonstrate how good our model is at reconstruction, we trained a stacked LSTM on the same autoencoding task. Figure 2 shows performance of LSTM models for embedding sizes 300, 512 and 1024. All LSTMs have two encoder and two decoder layers. The 1024 variant seems to have reached a saturation point, as it performs similarly to the 512 version. All RAEs and LSTMs were trained for 20 epochs and models with same




Figure 2: Accuracy comparison of different embedding sizes (300, 512, 1024 and 2048) for different sequence lengths. Left is our recursive autoencoder and right a stacked LSTM. An output embedding is counted as correct if the closest embedding out of all the vocabulary is its corresponding input embedding.




Figure 3: Accuracy comparison of our RAE model versus a stacked LSTM for embedding sizes 512 and 1024. Models of same embedding size have the same capacity.

Model	d_{emb}	MSE (dev)
LSTM	300	0.0274
	512	0.0231
	1024	0.0191
RAE	300	0.0208
	512	0.0124
	1024	0.0075
	2048	0.0019

Table 1: Mean squared error loss of stacked LSTMs and our RAE model for different embedding sizes. All models are trained on the autoencoding task for 20 epochs and models of same embedding size have the same capacity. MSE is computed on the BookCorpus dev set (Zhu et al., 2015), between the input GloVe embeddings (Pennington et al., 2014) and output embeddings.

embedding size have the same capacity. Figure 3 shows a better side by side comparison of the RAE and the LSTM for embedding sizes 512 and 1024. Table 1 shows the MSE loss of all models on the dev set after 20 epochs. The LSTM with an embedding size of 1024 only slightly achieves lower MSE than the RAE with embedding size 300.

When the output and input embeddings don't match as nearest, they are usually close. Figure 4 shows the gain in accuracy for the 1024 and 2048 variants when considering an output embedding as correct if the input embedding is in the five closest to the output, out of all the vocabulary. For the 1024 version, we see on average an increase in accuracy of 2.7%, while for the 2048 variant, the gain only starts to get noticeable for sequences longer than 30, with an overall average increase of 1.4%.

3.2 Sentiment Analysis

With strong autoencoding performance, one would think that features get deeply encoded into the representation, making it difficult to easily extract them back, which is crucial for a great number of tasks. To this end, we test our architecture on the sentiment analysis task.

The Stanford Sentiment Treebank (Socher et al., 2013) is a sentiment classification task where each sample in the dataset is a sentence with its corresponding sentiment tree. Each node in the tree is human annotated, with the leaves representing the sentiment of the words, all the way up to the root node, representing the whole sequence. Comparison is usually done on a binary or five label classification task, ranging from negative to positive. Most models are usually by design only able to classify the root node, while our architecture al-




Figure 4: Difference in accuracy when counting an output embedding as correct if the corresponding input embedding is in the five closest versus the closest. Comparison is done on our RAE model with embedding sizes 1024 and 2048.

lows classification of every node in the tree. We use a linear layer on top of each embedding in the encoder to classify sentiment.

We present in Table 2 results for fine-grained sentiment analysis on all nodes as well as comparison with recent state-of-the-art methods on binary sentiment classification of the root node. For the five class sentiment task, we compare our model with the original Sentiment Treebank results and beat all the models. In order to compare our approach with state-of-the-art methods, we also trained our model on the binary classification task with sole classification of the root node. Other presented models are GenSen (Subramanian et al., 2018) and BERT_{BASE} (Devlin et al., 2018). Both these recent methods perform extremely well on multiple natural language processing tasks. We set the RAE embedding size d_{emb} to 1024. Larger embedding sizes did not improve the accuracy of our model for this task. In this setting, the RAE has 11M parameters, while the models we compare with, GenSen and BERT_{BASE}, have respectively 100M and 110M parameters. Both our model and GenSen fail to beat the RNTN model for the SST-2 task. We see an improvement in accuracy when combining both methods’ embeddings, surpassing every model in the SST paper, while being close to BERT_{BASE}’s performance.

Training solely on sentiment classification had same performance as jointly training on the autoencoding task, as the latter had no impact on the sentiment analysis performance. Joint training though had a small impact on reconstruction.

Model	SST-5 (All)	SST-2 (Root)
NB	67.2	81.8
SVM	64.3	79.4
BiNB	71.0	83.1
VecAvg	73.3	80.1
RNN	79.0	82.4
MV-RNN	78.7	82.9
RNTN	80.7	85.4
RAE	81.07	83
GenSen	-	84.5
RAE + GenSen	-	86.43
BERT _{BASE}	-	93.5

Table 2: SST-5 and SST-2 performance on all and root nodes respectively. Model results in the first section are from the Stanford Treebank paper (2013). GenSen and BERT_{BASE} results are from (Subramanian et al., 2018) and (Devlin et al., 2018) respectively.

4 Conclusion & Future Work

In this paper, we introduced a recursive autoencoder method for generating sentence and sub-sentence representations. Decoding from a single embedding and working with a 337k vocabulary, we manage to get near perfect reconstruction for sequences of up to 40 length and very good reconstruction for longer sequences. Capitalizing on our model’s architecture, we showed our method to perform well on sentiment analysis and more precisely its advantage when classifying sentiment trees.

Continuing in the direction of training our model on different NLP tasks, we would like our representations to generalize well on downstream tasks while maintaining their reconstruction property. We would also like to further explore the usage of sub-sentence representations in natural language processing. Finally, we would like to learn our sentence embeddings’ latent space, similarly to Subramanian et al. (2018)’s method, so as to leverage our autoencoder’s strong reconstruction ability and generate very long sequences of text.

Acknowledgments

This research was enabled in part by support provided by Compute Canada (www.computecanada.ca). We would also like to thank Tom Bosc, Sandeep Subramanian, Sai Rajeswar and Chinnadhurai Sankar for their invaluable feedback.

References

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). *arXiv e-prints*, page arXiv:1508.05326.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. [On the Properties of Neural Machine Translation: Encoder-Decoder Approaches](#). *arXiv e-prints*, page arXiv:1409.1259.
- Alexis Conneau and Douwe Kiela. 2018. [Senteval: An evaluation toolkit for universal sentence representations](#). *arXiv preprint arXiv:1803.05449*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised Learning of Universal Sentence Representations from Natural Language Inference Data](#). *arXiv e-prints*, page arXiv:1705.02364.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive Language Models Beyond a Fixed-Length Context](#). *arXiv e-prints*, page arXiv:1901.02860.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). *arXiv e-prints*, page arXiv:1810.04805.
- Cicero Dos Santos and Maira Gatti de Bayser. 2014. [Deep convolutional neural networks for sentiment analysis of short texts](#).
- G. E. Hinton and R. R. Salakhutdinov. 2006. [Reducing the dimensionality of data with neural networks](#). *Science*, 313(5786):504–507.
- Rie Johnson and Tong Zhang. 2014. [Effective Use of Word Order for Text Categorization with Convolutional Neural Networks](#). *arXiv e-prints*, page arXiv:1412.1058.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. [A Convolutional Neural Network for Modelling Sentences](#). *arXiv e-prints*, page arXiv:1404.2188.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2015. [Skip-Thought Vectors](#). *arXiv e-prints*, page arXiv:1506.06726.
- Filippos Kokkinos and Alexandros Potamianos. 2017. [Structural Attention Neural Networks for improved sentiment analysis](#). *arXiv e-prints*, page arXiv:1701.01811.
- Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. 2015. [Recurrent convolutional neural networks for text classification](#).
- Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. 1989. [Back-propagation applied to handwritten zip code recognition](#). *Neural Comput.*, 1(4):541–551.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer Normalization](#). *arXiv e-prints*, page arXiv:1607.06450.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019a. [Improving Multi-Task Deep Neural Networks via Knowledge Distillation for Natural Language Understanding](#). *arXiv e-prints*, page arXiv:1904.09482.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019b. [Multi-Task Deep Neural Networks for Natural Language Understanding](#). *arXiv e-prints*, page arXiv:1901.11504.
- Thang Luong, Richard Socher, and Christopher Manning. 2013. [Better word representations with recursive neural networks for morphology](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 104–113. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Jordan B. Pollack. 1990. [Recursive distributed representations](#). *Artificial Intelligence*, 46(1):77 – 105.
- Richard Socher, Eric H. Huang, Jeffrey Pennington, Andrew Y. Ng, and Christopher D. Manning. 2011a. [Dynamic pooling and unfolding recursive autoencoders for paraphrase detection](#). In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, pages 801–809, USA. Curran Associates Inc.
- Richard Socher, Brody Huval, Christopher D. Manning, and Andrew Y. Ng. 2012. [Semantic compositionality through recursive matrix-vector spaces](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1201–1211. Association for Computational Linguistics.
- Richard Socher, Christopher D. Manning, and Andrew Y. Ng. 2010. [Learning continuous phrase representations and syntactic parsing with recursive neural networks](#). In *In Proceedings of the NIPS-2010 Deep Learning and Unsupervised Feature Learning Workshop*.
- Richard Socher, Jeffrey Pennington, Eric H. Huang, Andrew Y. Ng, and Christopher D. Manning. 2011b. [Semi-Supervised Recursive Autoencoders for Predicting Sentiment Distributions](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment tree-bank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642. Association for Computational Linguistics.
- Sandeep Subramanian, Sai Rajeswar Mudumba, Alessandro Sordani, Adam Trischler, Aaron C Courville, and Chris Pal. 2018. [Towards text generation with adversarially learned neural outlines](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 7551–7563. Curran Associates, Inc.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. [Learning General Purpose Distributed Sentence Representations via Large Scale Multi-task Learning](#). *arXiv e-prints*, page arXiv:1804.00079.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention Is All You Need](#). *arXiv e-prints*, page arXiv:1706.03762.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2017. [A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference](#). *arXiv e-prints*, page arXiv:1704.05426.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 649–657. Curran Associates, Inc.
- Yukun Zhu, Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *arXiv preprint arXiv:1506.06724*.