

A Related Compression Methods

A.1 BERT Compression Methods

Many techniques have been proposed to compress BERT (Devlin et al., 2018). Ganesh et al. (2020) provide a survey on various compression methods for BERT. Most existing methods focus on alternative architectures in transformer layers or learning strategies.

In our work, we use DistilBERT and ALBERT-base as light pretrained language model encoders for semantic parsing. DistilBERT (Sanh et al., 2019) uses distillation to pretrain a model that is 40% smaller and 60% faster than BERT-base, while retaining 97% of its downstream performances. ALBERT (Lan et al., 2019) factorizes the embedding and shares parameters among the transformer layers in BERT and results in better scalability than BERT. ALBERT-xxlarge outperforms BERT-large on GLUE (Wang et al., 2018), RACE (Lai et al., 2017), and SQUAD (Rajpurkar et al., 2016) while using less parameters.

We use compositional code learning (Shu and Nakayama, 2017) to compress the model embeddings, which contain a substantial amount of model parameters. Previously ALBERT uses factorization to compress the embeddings. We find more compression possible with code embeddings.

A.2 Embedding Compression Methods

Varied techniques have been proposed to learn compressed versions of non-contextualized word embeddings, such as, Word2Vec (Mikolov et al., 2013) and GLoVe (Pennington et al., 2014). Subramanian et al. (2018) use denoising k-sparse autoencoders to achieve binary sparse interpretable word embeddings. Chen et al. (2016) achieve sparsity by representing the embeddings of uncommon words using sparse linear common combination of common words. Lam (2018) achieve compression by quantization of the word embeddings by using 1-2 bits per parameter. Faruqui et al. (2015) use sparse coding in a dictionary learning setting to obtain sparse, non-negative word embeddings. Raunak (2017) achieve dense compression of word embeddings using PCA combined with a post-processing algorithm. Shu and Nakayama (2017) propose to represent word embeddings using compositional codes learnt directly in end-to-end fashion using neural networks. Essentially few common basis vectors are learnt and embeddings are reconstructed using their composition via a discrete code vector

specific to each token embedding. This results in 98% compression rate in sentiment analysis task and 94% - 99% in machine translation tasks without performance loss while applied to LSTM based models. All the above techniques are applied to embeddings such as WordVec and Glove, or LSTM models.

We aim at learning space-efficient embeddings for transformer-based models. We focus on compositional code embeddings (Shu and Nakayama, 2017) since they maintain the vector dimensions, do not require special kernels for calculating in a sparse or quantized space, can be finetuned with transformer-based models end-to-end, and achieve extremely high compression rate. Chen et al. (2018) explores similar idea as Shu and Nakayama (2017) and experiment with more complex composition functions and guidances for training the discrete codes. Chen and Sun (2019) further show that end-to-end training from scratch of models with code embeddings is possible. Given various pretrained language models, we find that the method proposed by Shu and Nakayama (2017) is straightforward and perform well in our semantic parsing experiments.

References

- Ting Chen, Martin Renqiang Min, and Yizhou Sun. 2018. Learning k-way d-dimensional discrete codes for compact embedding representations. *arXiv preprint arXiv:1806.09464*.
- Ting Chen and Yizhou Sun. 2019. Differentiable product quantization for end-to-end embedding compression. *arXiv preprint arXiv:1908.09756*.
- Yunchuan Chen, Lili Mou, Yan Xu, Ge Li, and Zhi Jin. 2016. Compressing neural language models by sparse word representations. *arXiv preprint arXiv:1610.03950*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Manaal Faruqui, Yulia Tsvetkov, Dani Yogatama, Chris Dyer, and Noah Smith. 2015. Sparse overcomplete word vector representations. *arXiv preprint arXiv:1506.02004*.
- Prakhar Ganesh, Yao Chen, Xin Lou, Mohammad Ali Khan, Yin Yang, Deming Chen, Marianne Winslett, Hassan Sajjad, and Preslav Nakov. 2020. Compressing large-scale transformer-based models: A case study on bert. *arXiv preprint arXiv:2002.11985*.

100	Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang,	150
101	and Eduard Hovy. 2017. Race: Large-scale reading	151
102	comprehension dataset from examinations. <i>arXiv</i>	152
103	<i>preprint arXiv:1704.04683</i> .	153
104	Maximilian Lam. 2018. Word2bits-quantized word	154
105	vectors. <i>arXiv preprint arXiv:1803.05651</i> .	155
106	Zhenzhong Lan, Mingda Chen, Sebastian Goodman,	156
107	Kevin Gimpel, Piyush Sharma, and Radu Soricut.	157
108	2019. Albert: A lite bert for self-supervised learn-	158
109	ing of language representations. <i>arXiv preprint</i>	159
110	<i>arXiv:1909.11942</i> .	160
111	Tomas Mikolov, Kai Chen, Greg Corrado, and Jef-	161
112	frey Dean. 2013. Efficient estimation of word	162
113	representations in vector space. <i>arXiv preprint</i>	163
114	<i>arXiv:1301.3781</i> .	164
115	Jeffrey Pennington, Richard Socher, and Christopher D	165
116	Manning. 2014. Glove: Global vectors for word rep-	166
117	resentation. In <i>Proceedings of the 2014 conference</i>	167
118	<i>on empirical methods in natural language process-</i>	168
119	<i>ing (EMNLP)</i> , pages 1532–1543.	169
120	Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and	170
121	Percy Liang. 2016. Squad: 100,000+ questions	171
122	for machine comprehension of text. <i>arXiv preprint</i>	172
123	<i>arXiv:1606.05250</i> .	173
124	Vikas Raunak. 2017. Simple and effective dimension-	174
125	ality reduction for word embeddings. <i>arXiv preprint</i>	175
126	<i>arXiv:1708.03629</i> .	176
127	Victor Sanh, Lysandre Debut, Julien Chaumond, and	177
128	Thomas Wolf. 2019. Distilbert, a distilled version	178
129	of bert: smaller, faster, cheaper and lighter. <i>arXiv</i>	179
130	<i>preprint arXiv:1910.01108</i> .	180
131	Raphael Shu and Hideki Nakayama. 2017. Compress-	181
132	ing word embeddings via deep compositional code	182
133	learning. <i>arXiv preprint arXiv:1711.01068</i> .	183
134	Anant Subramanian, Danish Pruthi, Harsh Jhamtani,	184
135	Taylor Berg-Kirkpatrick, and Eduard Hovy. 2018.	185
136	Spine: Sparse interpretable neural embeddings. In	186
137	<i>Thirty-Second AAAI Conference on Artificial Intelli-</i>	187
138	<i>gence</i> .	188
139	Alex Wang, Amanpreet Singh, Julian Michael, Felix	189
140	Hill, Omer Levy, and Samuel R. Bowman. 2018.	190
141	GLUE: A multi-task benchmark and analysis plat-	191
142	form for natural language understanding. <i>CoRR</i> ,	192
143	abs/1804.07461.	193
144		194
145		195
146		196
147		197
148		198
149		199