

Compressing Word Embeddings

Martin Andrews^(✉)

Red Cat Labs, Singapore, Singapore

Martin.Andrews@RedCatLabs.com

<http://www.RedCatLabs.com>

Abstract. Recent methods for learning vector space representations of words have succeeded in capturing fine-grained semantic and syntactic regularities using large-scale unlabelled text analysis. However, these representations typically consist of dense vectors that require a great deal of storage and cause the internal structure of the vector space to be opaque. A more ‘idealized’ representation of a vocabulary would be both compact and readily interpretable. With this goal, this paper first shows that Lloyd’s algorithm can compress the standard dense vector representation by a factor of 10 without much loss in performance. Then, using that compressed size as a ‘storage budget’, we describe a new GPU-friendly factorization procedure to obtain a representation which gains interpretability as a side-effect of being sparse and non-negative in each encoding dimension. Word similarity and word-analogy tests are used to demonstrate the effectiveness of the compressed representations obtained.

1 Introduction

Distributed representations of words have been shown to benefit NLP tasks like parsing, named entity recognition, and sentiment analysis, as well as being used as the raw material for other deep learning tasks.

Surprisingly, these word vector embeddings can be derived directly from raw, unannotated corpora. Once created, the vector embedding E can be expressed simply as a list of vocabulary words (of size V), and a matrix of size $V \times d$, where d is the dimensionality of the embedding space.

As argued in [1], while this dense matrix representation may be handled with ease by computers, there are cognitive arguments against such a representation being the basis of language (see [2] for broader discussion on this point). For instance, it seems unlikely that the same small set of features are sufficient and necessary to describe all semantic domains of a full adult vocabulary. It would also be uneconomical for people to store all negative properties of a concept, such as the fact that dogs do not have wheels, or that airplanes are not used for communication. Indeed, in feature-norming exercises (for example [3]) where participants are asked to list the properties of a word, the aggregate descriptions are typically limited to approximately 10–20 characteristics for a given concrete concept.

So, for cognitive plausibility, we claim that a feature set should have three characteristics: it should only store positive facts; it should have a wide range of feature types, to cover all semantic domains in the typical mental lexicon; and only a small number of these should be active to describe each word/concept.

2 Models

In this work, a 300-dimensional GloVe [4] embedding was used as a concrete baseline for a number of different lossy compression methods¹.

2.1 Simple Compression

A number of different compression methods were explored to establish a ‘storage budget’ for the subsequent experiments in sparse encodings. These initial methods focussed on discarding data, while resulting in an approximately equivalent level of performance in the ‘Google’ word analogy task (described later).

The methods that were explored included (i) directly discarding a fixed proportion of the vector dimensions; (ii) thresholding the data; (iii) quantising the data based on curves of the form $\pm|v_i|^\alpha$ for each dimension i of a given vector v ; and (iv) adaptive level encoding. This last approach, efficiently implemented via Lloyd’s algorithm [5], was shown to be capable of approximating each element v_i of the entire vocabulary with only 8 discrete levels (each dimension i being independently calibrated) while still performing acceptably on the analogy task.

Thus, for the 300-dimensional embedding, a 900-bit ‘storage budget’ was established, and used as the storage bound for the sparsification experiments.

2.2 Lloyd’s Algorithm

For a given number of quantisation levels n and an element $e \sim E$, we want to create a set of quantisation levels e_q that minimises:

$$\sum_{e \in E} \min_q (e - e_q)^2$$

This process can be performed iteratively, starting with the e_q placed uniformly within the values of E sorted numerically.

The algorithm then iteratively updates the e_q such that the centroid of each neighbouring cluster is found, and the e_q are then updated to these positions. The algorithm stops when no more changes are required.

This optimisation is done for each element/column of the embedding independently, and a list of the respective levels of e_q is stored, in addition to the index within each list for each element of the representation.

2.3 Sparse Compression

Non-negative Sparse Embedding (“NNSE”). A very promising approach to sparse embedding was taken in [1], where the following optimisation was solved iteratively to generate a sparse embedding A (which is V vectors each

¹ The techniques described in this paper can be applied to any embedding, since nothing specific to GloVe has been used.

with dimensionality k), and also creates a ‘dictionary’ D (which maps the sparse elements of A onto the real-valued embedding E):

$$\operatorname{argmin}_{A \in \mathbb{R}^{m \times k}, D \in \mathbb{R}^{k \times n}} \sum_{i=1}^V \|E_{i,:} - A_{i,:} \times D\|_2 + \lambda \|A_{i,:}\|_1$$

where : $D_{i,:}, D_{i,:}^T \leq 1, 1 \leq \forall i \leq k$
and $A_{i,j} \geq 0, 1 \leq \forall i \leq V, 1 \leq \forall j \leq k$

Although this produced excellent embedding results, one requirement is to vary the parameters k and λ such that the sparse embedding A had a ‘reasonable’ number of positive entries. Since one of the targets of this paper is compression, targeting a concrete compression level, this hyperparameter search is not acceptable.

Winner-Take-All Autoencoders. Another approach to sparse encoding, applied chiefly to visual tasks, was described in [6]. There, “Winner-Take-All Autoencoders” obtained strict sparsity targets by interposing a drop-out layer that only allows the top- α fraction of the layer’s inputs to be passed through, while the rest are set to zero. An important implementation detail is that this drop-out is done on a *per dimension* basis over each minibatch of training examples. This implicitly causes the training to balance the occurrences of non-zero entries across the different dimensions of the embedding space, while allowing the $A_{i,:}$ vectors to have varying numbers of non-zero entries.

Unfortunately, obtaining the top- α fraction of a set of numbers requires sorting the list, which means that including this layer in an neural network will move data-intensive computations off the GPU. In order to produce a sparse embedding in reasonable time, however, the GPU is required, so an algorithmic short-cut was implemented that can be efficiently executed inside the GPU.

GPU-Friendly Top- α . Clearly, determining the top- α percentile of a distribution is equivalent to determining the hurdle value h^* above which the fraction α of the realised values in the minibatch lie².

For a given embedding dimension j , and assuming that $\alpha \ll 50\%$ and that the distribution is not pathological, this hurdle can be bounded above by $h_j^+ = \max(A_{:,j})$ and below by $h_j^- = \text{mean}(A_{:,j})$. A binary section search can then be made for h^* , since the fractional percentile for any given h can be determined by taking the mean of a simple matrix $A_{:,j} > h$ indicator function³.

Although a thorough implementation of this binary section process would include stopping criteria, that would not make sense within the GPU’s cores’ execution flow. So a fixed number of bisection iterations was used (specifically 5 for our 16,384 minibatch size). This was found to give robust results, with a value of h being found that identified (approximately) the top- α percentile consistently. The net result of avoiding CPU operations was a 39 \times speed-up.

² A minibatch has 16,384 examples – large enough for distribution approximations.

³ Also, $\alpha_0^+ = (1/\text{batchsize})$ initially, since it is the maximum value in $A_{:,j}$.

Experiments were also conducted to try and optimise the placement of each h estimate, using distributional properties of $N()$, or linear interpolation. These were not successful, however, since it appears that the neural network learns to exploit the distribution assumptions being made, so as to win artificially low h (and thus a higher accepted α than required).

Sparse Autoencoding for Word Embeddings. To avoid learning *from scratch* the entire sparse matrix A , an autoencoding scheme was set up, so that the target embedding E was mapped onto intermediate layers, then quantised, then mapped back to itself (via a trainable dictionary D).

The layers of the implemented autoencoder are described in Table 1, where concrete sizes are given, assuming a vocabulary size of $2^{17} = 131,072$ words, 300-dimensional underlying embedding, and a sparse embedding of size $k = 1024$.

Table 1. Sparse Auto-encoding model structure

Parameter set	Intra-layer operations	Output shape
Word embedding (all tokens)	\mathbf{E}	$(2^{17}, 300) \in \mathbb{R}$
Hidden layer	$\max(\mathbf{W}^H \mathbf{x} + \mathbf{b}^H, 0)$	$(2^{17}, 300 \times 8) \in \mathbb{R}^+$
Pre-binary linear layer	$\mathbf{W}^L \mathbf{x} + \mathbf{b}^L$	$(2^{17}, 1024) \in \mathbb{R}$
Batch normalization layer	$\text{batchnormalize}(\mathbf{x})$	$(2^{17}, 1024) \in \mathbb{R}$
Rectification	$\max(\mathbf{x}, 0)$	$(2^{17}, 1024) \in \mathbb{R}^+$
Gaussian noise	$\mathbf{x} + N(0, \beta_t)$	$(2^{17}, 1024) \in \mathbb{R}$
Top- α sparsification	$\text{drop-all-but-}\alpha(\mathbf{x}, \alpha_t)$	$(2^{17}, 1024) \in \mathbb{R}^+$
Sparse encoded version	$\mathbf{A} = \mathbf{x}$	$(2^{17}, 1024) \in \mathbb{R}$
Output ($=\mathbf{E}^*$)	$\mathbf{W}^F \mathbf{x} + \mathbf{b}^F$	$(2^{17}, 300) \in \mathbb{R}$

The quantity being optimised via gradient descent was the l_2 error between the output embedding E^* , and the input E , and this scheme was accelerated using ADAM [7].

Note that no regularisation terms were used. There is regularisation being implicitly applied due to the top- α sparsity constraint, coupled with the batch normalization that occurs in the layer before the sparse encoding. These combine to constrain the sparse encoding A within reasonable bounds.

In order to improve convergence, a scheme was used whereby a variable σ (initially zero) was incremented (by 0.01) after each epoch *only if* the l_2 error over the epoch averaged less than 0.01. The variables α_t and β_t were then dynamically adjusted as follows:

$$\begin{aligned}\alpha_t &= 50\% \times (e^{10\sigma}) + \alpha \times (1 - e^{10\sigma}) \\ \beta_t &= 0.2 \times e^{0.01\sigma}\end{aligned}$$

Having this path-dependent speed regulator enabled learning pressure to be applied incrementally through training in a way that was sensitive to current training progress. In particular, the sparsity factor α_0 was started out at ‘easy’ values, so that the initial weights could migrate to areas where they were at least tackling the underlying problem - and α_t then moved asymptotically closer to α over the course of a thousand epochs. Similarly, the Gaussian noise layer seemed to improve search in later periods, and this was decayed more slowly.

After training, the noise-layer was switched off, and the network-learned A and E^* were used for the performance analyses below.

Representation of Sparse Encodings. Suppose a non-negative sparse encoding over a k -dimensional space is required, and a specific ‘bit budget’ of n bits is given. Since the majority of encoding values will be zero, the representation need only store the addresses of the non-zero elements, and their values.

Moreover, because the ordering of the list of values is an additional degree of freedom, one can increase the information content by storing (a) the length of the list, (b) the locations of each of the values in declining numerical order, (c) the highest value, and (d) the percentage ratio between each successive pair (typically in the range [70 %...100 %]).

Supposing that 3-bits is sufficient for each ratio (d)⁴, then the required sparsity ratio α can be determined:

$$\alpha = \frac{n}{k(\log_2(k) + 3)}$$

In the experiments that follow, the α values are chosen to be within this bound.

In addition to the $n \times V$ bits of storage required, there is the additional overhead of storing the dictionary D - which consists of $k \times d$ elements, which can be a significant factor if k is large.

3 Experiments

3.1 Corpus

The text corpus used here was the concatenation of (a) the ‘One Billion Word Benchmark’ [8], and (b) a cleaned version of English Wikipedia (August 2013 dump, using only pages with greater than 20 pageviews), which was pre-processed by removing non-textual elements, sentence splitting, and tokenization. After preprocessing the corpus contained 1.2 billion tokens (680 million and 524 million from each respective source)⁵.

⁴ While (c) might be stored with higher fidelity, the remaining ratios are less exacting.

⁵ Importantly, these resources have been made freely available without restrictive licenses, and in the same spirit, the code for this paper is being released under a permissive license.

Models were derived using windows of 5 tokens to each side of the focus word. Rather than focus on a word occurrence limit (as is common), a vocabulary-size of $2^{17} = 131,072$ words was chosen, since (a) it made batching more convenient, and (b) the limit is arbitrary either way.

3.2 Test Datasets

We evaluated each word representation on seven datasets covering similarity and analogy tasks, using the test framework of Levy et al. [9], which has code and data available at: <https://bitbucket.org/omerlevy/hyperwords>.

Word Similarity. Five datasets were used to evaluate word similarity: the popular WordSim353 [10] partitioned into two datasets, WordSim Similarity and WordSim Relatedness [11, 12]; Bruni et al.’s MEN dataset [13]; Radinsky et al.’s Mechanical Turk dataset [14]; and Luong et al.’s Rare Words dataset [15]. All these datasets contain word pairs together with human-assigned similarity scores. The word vectors are evaluated by ranking the pairs according to their cosine similarities, and measuring the correlation (Spearman’s ρ) with the human ratings.

Word Analogy. The two analogy datasets present questions of the form “ a is to a^* as b is to b^* ”, where b^* is hidden, and must be guessed from the entire vocabulary. MSR’s analogy dataset [16] contains 8000 morpho-syntactic analogy questions, such as “good is to best as smart is to smartest”. Google’s analogy dataset [17] contains 19544 questions, about half of the same kind as in MSR (syntactic analogies), and another half of a more semantic nature, such as capital cities (“Paris is to France as Tokyo is to Japan”). After filtering questions involving out-of-vocabulary words, i.e. words that did not appear in the pruned Corpus, we remain with 7118 instances in MSR and 19296 instances in Google.

As in [18], the analogy questions are answered using both 3COSADD as well as 3COSMUL.

3.3 Results

Embeddings with representations compressed by the Lloyd algorithm and the non-negative sparse encoding methods outlined above were run, with results shown in Tables 2 and 3.

In addition to the reconstructed embeddings E^* for each method, the *raw results* of the similarity and analogy tests were run on the intermediate sparse embeddings A themselves.

As a point of comparison, a random vector approach (binarised locality sensitive hashing: “LSH” [19]) was also tested, with the same 900-bit constraint.

Table 2. Similarity results

Method	WordSim Similarity	WordSim Relatedness	Bruni et al.	Radinsky et al.	Luong et al.
			MEN	M. Turk	Rare words
GloVe baseline	66.2 %	52.1 %	69.1 %	63.2 %	22.8 %
Lloyd-8	65.9 %	51.5 %	68.6 %	62.3 %	22.7 %
$E^*(k = 4096, \alpha = 1.50 \%)$	65.0 %	50.2 %	68.5 %	62.7 %	22.2 %
$E^*(k = 1024, \alpha = 6.75 \%)$	64.7 %	51.0 %	67.7 %	62.7 %	21.9 %
$A(k = 4096, \alpha = 1.50 \%)$	63.7 %	45.2 %	62.4 %	49.5 %	18.0 %
$A(k = 1024, \alpha = 6.75 \%)$	69.2 %	51.3 %	69.5 %	62.1 %	17.4 %
LSH-900	62.0 %	45.8 %	65.9 %	59.2 %	22.0 %

Table 3. Analogy results

Method	Google	MSR
	Add/Mul	Add/Mul
GloVe baseline	67.1 %/68.5 %	53.4 %/56.6 %
Lloyd-8	65.9 %/67.4 %	51.9 %/54.5 %
$E^*(k = 4096, \alpha = 1.50 \%)$	62.5 %/66.4 %	51.8 %/54.1 %
$E^*(k = 1024, \alpha = 6.75 \%)$	62.4 %/62.9 %	49.0 %/50.1 %
$A(k = 4096, \alpha = 1.50 \%)$	37.6 %/40.8 %	27.3 %/29.9 %
$A(k = 1024, \alpha = 6.75 \%)$	52.5 %/55.1 %	40.5 %/43.8 %
LSH-900	53.0 %/53.1 %	41.2 %/42.2 %

4 Discussion

4.1 Level Quantisation vs. More Sophisticated Methods

The level-quantisation approaches to compression work extremely well, and are relatively simple to implement. Assuming values would otherwise be stored as 32-bit floats, the Lloyd method achieves very similar scores, with only 3-bits per value (the overhead of storing the quantisation levels is only the equivalent of storing 8 of the original word-vectors). However, they do not accomplish the goal of learning about the underlying embedding through an efficient compression algorithm.

4.2 Performance of Sparse Embeddings

As can be seen from Table 2, the reconstruction (E^*) results for both sparse embedding methods are only marginally below those of the original embeddings, and generally better than the LSH re-representation. This is satisfying, because it shows that the GPU-friendly method outlined here actually reconstructs the embedding without a significant loss in performance, within the same ‘bit budget’ as the quantisation methods.

The fact that the sparse encodings A can also perform as embeddings on their own (without involving the learned dictionary D) is encouraging, since it implies that there is more information about the underlying language that can be obtained from existing word embeddings ‘for free’.

Interestingly, their performance on similarity tasks is far higher than on the analogy ones (particularly in the case of $k = 4096$ which has an α of only 1.50%). This can be understood by considering the algebra of the sparse non-negative vectors being used. For similarity purposes, non-negative vectors can be scored using the same a cosine measure that works for more general dense, real-valued vectors. However, for the analogy tasks, there are implicit subtractions being done, which result in direction vectors that are not part of the same algebra. This deserves further work, to see whether other operators would be able to make use of the sparse vector spaces’ geometry more fully (potentially including ideas from [20]).

4.3 Interpretability of the Sparse “ A ”

Table 4 (which lists the highest weighted words in each dimension that ‘motorbike’ is also most highly weighted) clearly demonstrates that the A sparse representation has learned something about the structure of the English language ‘for free’, using only data obtained from an embedding trained on unlabelled data itself.

Table 4. Top ‘motorbike’ dimensions

Model	Top words in each of first 7 dimensions
GloVe baseline	<ul style="list-style-type: none"> · lb., four-bladed, propeller, propellers, two-bladed, ... · passerine, 1975-79, rennae, fyrstenberg, edw, coots, ... · banchoston, oshiomhole, 30-sept, holmer, smithee, recon, ... · http://www.nytimes.com, (888), receival, jamiat, shyi, ... · subjunctive, purley, 11-july, broaddus, muharram, ebit, ... · proximus, pattani, 31-feb, wgc, 30-nov, crossgen, 2,631, ... · officership, tvcolumn, integrable, salticidae, o-157, ...
$A_{(\alpha=6.75\%)}^{(k=1024)}$	<ul style="list-style-type: none"> · vehicles, vehicle, cars, scrappage, car, 4x4, armored, ... · prix, races, race, laps, vettel, rikknen, sprint, ... · ski, coal, gas, taxicab, nuclear, wine, cellphone, ... · kool, electrons, pulpit, ehta, gallen, gasol, birdman, ... · eric, anglo, tornadoes, rt, asteroids, dera, rim, ... · wear, trousers, dresses, jeans, wearing, worn, pants, ... · stabbed, kercher, 16-year-old, 15-year-old, 18-year-old, ...

5 Conclusion

The joint goals of good compression rates and cognitive plausibility are realistic and achievable.

Using the GPU-friendly sparsity Winner-Take-All Autoencoder scheme described, sparse, non-negative encodings have been demonstrated that combine high compression with interpretability ‘for free’.

Further work will include the investigation of operators that respect the geometry of these sparse vectors, so that analogy tests might perform in-line with the word similarity scores (which are purely direction-based).

Acknowledgments. The author thanks DC Frontiers, the creators of the data-centric service ‘Handshakes’ (<http://www.handshakes.com.sg/>), for their willingness to support this on-going research. DC Frontiers is the recipient of a Technology Enterprise Commercialisation Scheme grant from SPRING Singapore, under which this work took place.

References

1. Murphy, B., Talukdar, P.P., Mitchell, T.: Learning effective and interpretable semantic models using non-negative sparse embedding. In: International Conference on Computational Linguistics (COLING 2012), Mumbai, India (2012). <http://aclweb.org/anthology/C/C12/C12-1118.pdf>
2. Griffiths, T.L., Steyvers, M., Tenenbaum, J.B.: Topics in semantic representation. *Psychol. Rev.* **114**(2), 211 (2007)
3. Vinson, D.P., Vigliocco, G.: Semantic feature production norms for a large set of objects and events. *Behav. Res. Methods* **40**(1), 183–190 (2008)
4. Pennington, J., Socher, R., Manning, C.D.: Glove: global vectors for word representation. *Proceedings of the Empirical Methods in Natural Language Processing (EMNLP 2014)* **12**, 1532–1543 (2014)
5. Lloyd, S.: Least squares quantization in PCM. *IEEE Trans. Inf. Theory* **28**(2), 129–137 (2006)
6. Makhzani, A., Frey, B.J.: A winner-take-all method for training sparse convolutional autoencoders (2014). CoRR, abs/1409.2752
7. Kingma, D., Ba, J.: Adam: a method for stochastic optimization (2014). arXiv preprint: [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
8. Chelba, C., Mikolov, T., Schuster, M., Ge, Q., Brants, T., Koehn, P.: One billion word benchmark for measuring progress in statistical language modeling (2013). CoRR, abs/1312.3005
9. Levy, O., Goldberg, Y., Dagan, I.: Improving distributional similarity with lessons learned from word embeddings. *Trans. Assoc. Comput. Linguist.* **3**, 211–225 (2015)
10. Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., Ruppin, E.: Placing search in context: the concept revisited. *ACM Trans. Inf. Syst.* **20**(1), 116–131 (2002)
11. Zesch, T., Müller, C., Gurevych, I.: Using wiktionary for computing semantic relatedness. In: *Proceedings of the 23rd National Conference on Artificial Intelligence, AAAI 2008*, vol. 2, pp. 861–866. AAAI Press (2008)

12. Agirre, E., Alfonseca, E., Hall, K., Kravalova, J., Paşca, M., Soroa, A.: A study on similarity and relatedness using distributional and wordnet-based approaches. In: Proceedings of Human Language Technologies: the 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, NAACL 2009, Stroudsburg, PA, USA, pp. 19–27. Association for Computational Linguistics (2009)
13. Bruni, E., Boleda, G., Baroni, M., Tran, N.-K.: Distributional semantics in technicolor. In: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers, ACL 2012, Stroudsburg, PA, USA, vol. 1, pp. 136–145. Association for Computational Linguistics (2012)
14. Radinsky, K., Agichtein, E., Gabrilovich, E., Markovitch, S.: A word at a time: computing word relatedness using temporal semantic analysis. In: Proceedings of the 20th International Conference on World Wide Web, WWW 2011, New York, NY, USA, pp. 337–346. ACM (2011)
15. Luong, M.-T., Socher, R., Manning, C.D.: Better word representations with recursive neural networks for morphology. In: CoNLL, Sofia, Bulgaria (2013)
16. Mikolov, T., tau Yih, W., Zweig, G.: Linguistic regularities in continuous space word representations. In: Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT-2013). Association for Computational Linguistics, May 2013
17. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space (2013). arXiv preprint: [arXiv:1301.3781](https://arxiv.org/abs/1301.3781)
18. Levy, O., Goldberg, Y., Ramat-Gan, I.: Linguistic regularities in sparse and explicit word representations. In: CoNLL, pp. 171–180 (2014)
19. Charikar, M.S.: Similarity estimation techniques from rounding algorithms. In: Proceedings of the Thirty-fourth Annual ACM Symposium on Theory of Computing, STOC 2002, New York, NY, USA, pp. 380–388. ACM (2002)
20. Mahadevan, S., Chandar, S.: Reasoning about linguistic regularities in word embeddings using matrix manifolds (2015). CoRR, abs/1507.07636