CAN A FRUIT FLY LEARN WORD EMBEDDINGS?

Anonymous authors
Paper under double-blind review

ABSTRACT

The mushroom body of the fruit fly brain is one of the best studied systems in neuroscience. At its core it consists of a population of Kenyon cells, which receive inputs from multiple sensory modalities. These cells are inhibited by the anterior paired lateral neuron, thus creating a sparse high dimensional representation of the inputs. In this work we study a mathematical formalization of this network motif and apply it to learning the correlational structure between words and their context in a corpus of unstructured text, a common natural language processing (NLP) task. We show that this network can learn semantic representations of words and can generate both static and context-dependent word embeddings. Unlike conventional methods (e.g., BERT, GloVe) that use dense representations for word embedding, our algorithm encodes semantic meaning of words and their context in the form of sparse binary hash codes. The quality of the learned representations is evaluated on word similarity analysis, word-sense disambiguation, and document classification. It is shown that not only can the fruit fly network motif achieve performance comparable to existing methods in NLP, but, additionally, it uses only a fraction of the computational resources (shorter training time and smaller memory footprint).

1 INTRODUCTION

Deep learning has made tremendous advances in computer vision, natural language processing and many other areas. While taking high-level inspiration from biology, the current generation of deep learning methods are not necessarily biologically realistic. This raises the question whether biological systems can further inform the development of new network architectures and learning algorithms that can lead to competitive performance on machine learning tasks or offer additional insights into intelligent behavior. Our work is inspired by this motivation. We study a well-established neurobiological network motif from the fruit fly brain and investigate the possibility of reusing it for solving common machine learning tasks in NLP. We consider this exercise as a toy model example illustrating the possibility of "reprogramming" of naturally occurring algorithms and behaviors (clustering combinations of input stimuli from olfaction, vision, and thermo-hydro sensory system) into a target algorithm of interest (learning word embeddings from raw text) that the original biological organism does not naturally engage in.

The mushroom body (MB) is a major area of the brain responsible for processing of sensory information in fruit flies. It receives inputs from a set of projection neurons (PN) conveying information from several sensory modalities. The major modality is olfaction (Bates et al., 2020), but there are also inputs from the PN responsible for sensing temperature and humidity (Marin et al., 2020), as well as visual inputs (Vogt et al., 2016; Caron et al., 2020). These sensory inputs are forwarded to a population of approximately 2000 Kenyon cells (KCs) through a set of synaptic weights (Li et al., 2020). KCs are reciprocally connected through an anterior paired lateral (APL) neuron, which sends a strong inhibitory signal back to KCs. This recurrent network effectively implements winner-takes-all competition between KCs, and silences all but a small fraction of top activated neurons (Dasgupta et al., 2017). This is the network motif that we study in this paper; its schematic is shown in Fig. 1. KCs also send their outputs to mushroom body output neurons (MBONs), but this part of the MB network is not included into our mathematical model.

Behaviorally, it is important for a fruit fly to distinguish sensory stimuli, e.g., different odors. If a fruit fly senses a smell associated with danger, it's best to avoid it; if it smells food, the fruit fly might want to approach it. The network motif shown in Fig. 1 is believed to be responsible for

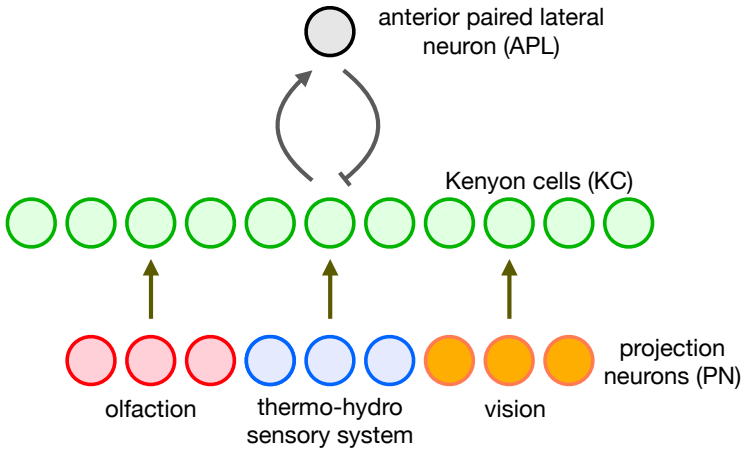

Figure 1: Network architecture. Several groups of PNs corresponding to different modalities send their activities to the layer of KCs, which are inhibited through the reciprocal connections to the APL neuron.

clustering sensory stimuli so that similar stimuli elicit similar patterns of neural responses at the level of KCs to allow generalization, while distinct stimuli result in different neural responses, to allow discrimination. Importantly, this biological network has evolved to accomplish this task in a very efficient way.

In computational linguistics there is a long tradition (Harris, 1954) of using distributional properties of linguistic units for quantifying semantic similarities between them, as summarized in the famous quote by JR Firth: "a word is characterized by the company it keeps" (Firth, 1957). This idea has led to powerful tools such as Latent Semantic Analysis (Deerwester et al., 1990), topic modelling (Blei et al., 2003), and language models like word2vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), and, more recently, BERT (Devlin et al., 2018) which relies on the Transformer model (Vaswani et al., 2017). Specifically word2vec models are trained to maximize the likelihood of a word given its context, GloVe models utilize global word-word co-occurence statistics, and BERT uses a deep neural network with attention to predict masked words (and the next sentence). As such, all these methods utilize the correlations between individual words and their context in order to learn useful word embeddings.

In our work we ask the following question: can the correlations between words and their contexts be extracted from raw text by the biological network of KCs, shown in Fig. 1? Further, how do the word representations learned by KCs differ from those obtained by existing NLP methods? Although this network has evolved to process sensory stimuli from olfaction and other modalities and not to "understand" language it uses a general purpose algorithm to embed inputs (from different modalities) into a high dimensional space with several desirable properties, which we discuss below.

Our approach relies on a recent proposal that the recurrent network of mutually inhibited KCs can be used as a "biological" model for generating sparse binary hash codes for the input data presented at the projection neuron layer (Dasgupta et al., 2017). It was argued that a matrix of random weights projecting from PN layer into the KCs layer leads to the highly desirable property of making the generated hash codes locality sensitive, i.e., placing similar inputs close to each other in the embedding space and pushing distinct stimuli far apart. A subsequent study (Ryali et al., 2020) has demonstrated that the locality sensitivity of the hash codes can be significantly increased, compared to the random case, if the matrix of weights from PN to KCs is learned from data. The learning is unsupervised.

Biologically, there is an ongoing debate in the neuroscience community regarding whether these projections are random. For instance, (Caron et al., 2013) argues for the random model, while (Zheng et al., 2020) presents evidence of the non-random structure of this network, which is related to the frequency of presented odors. Since the goal of our work is to build a useful AI system and not mimic every detail of the biological system, we adopt the data-driven synaptic weight strategy even if fruit flies may use random projections. As is clearly demonstrated in (Ryali et al., 2020), learned synapses lead to better performance.

Our main contributions in this work are the following:

1. Inspired by the fruit fly network, we propose an algorithm that makes it possible to generate binary (as opposed to continuous) word embeddings for words and their context. We systematically evaluate the performance of this algorithm on word similarity task, word-sense disambiguation, and document classification.
2. We demonstrate that our binary embeddings result in tighter and better separated clusters of concepts compared to GloVe embeddings.
3. We show that training the fruit fly network requires an order of magnitude smaller compute time than training the classical NLP architectures, like BERT, at the expense of relatively small decrease in classification accuracy.

Related Work Our work builds on several ideas previously discussed in the literature. The first idea is that fruit fly olfactory network can generate high quality hash codes for the input data in random (Dasgupta et al., 2017) and data-driven (Ryali et al., 2020) cases. There are two algorithmic differences of our approach compared to these previous studies. First, our network uses representational contraction, rather than expansion when we go from the PN layer to the KCs layer. Second, (Dasgupta et al., 2017; Ryali et al., 2020) construct hash codes for data coming from a single modality (e.g., images, or word vectors), while the goal of the present paper is to learn correlations between two different "modalities": target word and its context. The second idea pertains to the training algorithm for learning the PN→KCs synapses. We use a biologically plausible algorithm of (Krotov & Hopfield, 2019) to do this, with modifications that take into account the wide range of frequencies of different words in the training corpus (we discuss these differences in section 2.1). Also, similarly to (Dasgupta et al., 2017; Ryali et al., 2020) the algorithm of (Krotov & Hopfield, 2019) is used for learning the representations of the data, and not correlations between two types of data (context and target) as we do in this paper.

2 Learning Algorithm

Consider a training corpus. Each sentence can be decomposed into a collection of w-grams of consecutive words. If the word tokens come from a predefined vocabulary of size N_{voc}, the input to the algorithm is a vector of size $2N_{\text{voc}}$. This vector consists of two blocks: the context (the first N_{voc} elements), and the target (the remaining N_{voc} elements); see Fig. 2. In this work w is assumed to be an odd integer, and the target word is assumed to be the center of the w-gram. The target word

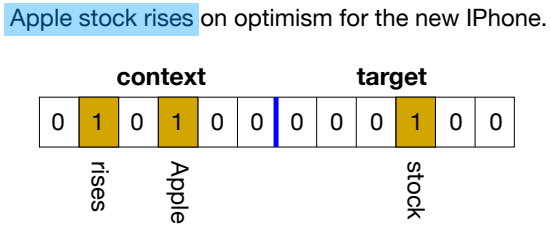

Figure 2: The encoding method. The input vector consists of two blocks separated by the (thick) blue line. Assuming $w = 3$, a center word "stock" is the target word and the two flanking words, highlighted in light blue, form a context.

is one-hot encoded in the target block, and the context words are binary encoded as a bag of words in the context block. The window w slides along the text corpus, and for each position generates a training vector $\mathbf{v^A} = \{v_i^A\}_{i=1}^{2N_{\text{voc}}}$, where the index A enumerates different w-grams, and index i enumerates positions in the context-target vector. These training vectors are passed to the learning algorithm. The goal of the algorithm is to learn correlations between the context and the target blocks.

2.1 Mathematical Formulation

Mathematically, the objective of the training algorithm is to distribute a set of context-target pairs among K buckets, so that similar pairs end up in similar buckets. In order to achieve this, the learning algorithm takes two inputs: a set of training vectors $\mathbf{v^A} \in \{0, 1\}^{2N_{\text{voc}}}$, and a vector of

occurrence probabilities $\mathbf{p} = \{p_i = f_{(i \mod N_{\text{voc}})}\}_{i=1}^{2N_{\text{voc}}} \in \mathbb{R}^{2N_{\text{voc}}}$, where f_j is the probability of observing word j in the training corpus. The learning can be formalized as a minimization of the energy function defined by

$$E = -\sum_{A \in \text{data}} \frac{\left\langle \mathbf{W}_{\hat{\mu}}, \mathbf{v}^{\mathbf{A}}/\mathbf{p} \right\rangle}{\left\langle \mathbf{W}_{\hat{\mu}}, \mathbf{W}_{\hat{\mu}} \right\rangle^{1/2}}, \qquad \text{where} \quad \hat{\mu} = \arg\max_{\mu} \left\langle \mathbf{W}_{\mu}, \mathbf{v}^{\mathbf{A}} \right\rangle \tag{1}$$

In this equation $\mathbf{W} \in \mathbb{R}^{K \times 2N_{\text{voc}}}$ is a matrix of synaptic connections, given as $\mathbf{W} = \{\mathbf{W}_\mu\} = \{W_{\mu i}\}$, projecting from PN layer (individual neurons in the layer are denoted by the index i) to the KC layer (individual neurons in the KC layer are denoted by the index μ). There are $2N_{\text{voc}}$ neurons in the PN layer and K neurons in the KC layer. The inner product $\langle \mathbf{X}, \mathbf{Y} \rangle = \sum_{i=1}^{2N_{\text{voc}}} X_i Y_i$ is defined as a contraction over index i of PN cells. In the numerator of the energy function the binary encoded w-gram is divided by the probabilities of occurrences of individual words element-wise, so that the numerator can be written as

$$\left\langle \mathbf{W}_{\hat{\mu}}, \mathbf{v}^{\mathbf{A}}/\mathbf{p} \right\rangle = \sum_{i=1}^{2N_{\text{voc}}} W_{\hat{\mu} i} \frac{v_i^A}{p_i}$$

Probabilities $\mathbf{p}$ are calculated based on the frequencies of words in the training corpus. The vocabulary contains N_{voc} most frequent words in the corpus, thus all the elements of p_i are non-zero and the element-wise division is well defined.

Intuitively, the goal of the training algorithm is to adjust the weights of the neural network so that they are aligned with w-grams that are frequently present in the training corpus. We rely on the assumption that semantically related w-grams share several "core" words, while a few individual words might be substituted by synonyms/antonyms. The minimization of the energy function (1) is accomplished by the iterative update of the weights satisfying the following learning rule (Krotov & Hopfield, 2019; Ryali et al., 2020)

$$\tau \frac{dW_{\mu i}}{dt} = g\Big[\sum_j W_{\mu j} v_j^A\Big]\Big[\frac{v_i^A}{p_i} - \Big(\sum_j W_{\mu j} \frac{v_j^A}{p_j}\Big) W_{\mu i}\Big] \tag{2}$$

In this differential equation the activation function is equal to one for a maximally driven hidden unit (Kenyon cell), and is equal to zero otherwise

$$g\Big[x_\mu\Big] = \delta_{\mu,\hat{\mu}}, \qquad \text{where} \quad \hat{\mu} = \arg\max_{\mu} \Big[x_\mu\Big] \tag{3}$$

The time constant τ is inversely proportional to the learning rate, and $\delta_{\mu,\hat{\mu}}$ is a Kronecker delta symbol.

2.2 Bio-Hashing

After learning is complete the hash codes for the inputs can be generated in the following way. Given the binary encoded w-gram $\mathbf{v}^{\mathbf{A}}$,

$$H_\mu = \begin{cases} 1, & \text{if } \langle \mathbf{W}_\mu, \mathbf{v}^{\mathbf{A}} \rangle \text{ in the top } k \text{ of all KCs activations} \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

This is a crude mathematical approximation of the biological computation performed by the PN–KC–APL neural network (Dasgupta et al., 2017; Ryali et al., 2020). An input $\mathbf{v}^{\mathbf{A}}$ generates an input current $\langle \mathbf{W}_\mu, \mathbf{v}^{\mathbf{A}} \rangle$ into the KC neurons using feedforward weights $W_{\mu i}$. The recurrent network of KCs and the APL neuron silences all but a small fraction of KCs. Those cells that remain active are assigned state 1, while the rest of the KCs are assigned the inactive state 0.

Notice, that equation (4) makes it possible to generate the hash codes for both individual words (static word embeddings like word2vec and GloVe) and phrases (similar to Transformer models). In the static case, the input $\mathbf{v}^{\mathbf{A}}$ has all zeros in the context block and a one-hot encoded word in the target block. In the context-dependent case, both blocks have binary encoded input words. Importantly, both context-dependent and static embeddings are mapped into the same space of sparse binary hash codes (a vector of K elements, with k ones in it). We show below that these hash codes capture semantic meaning of the target word and the context in which it is used. For the rest of the paper we refer to the parameter k in equation (4) as the *hash length*.

3 EMPIRICAL EVALUATION

The KC network shown in Fig. 1 was trained on the OpenWebText Corpus (Gokaslan & Cohen, 2019), which is a 32GB corpus of unstructured text containing approximately 6B tokens. The details of the training protocols and the hyperparameters are reported in section 6 in the supplement.

3.1 STATIC WORD EMBEDDINGS EVALUATION

Our aim here is to demonstrate that the sparse embeddings obtained by the fruit fly network motif are competitive with existing state-of-the-art GloVe embeddings (Pennington et al., 2014) (trained on Wiki2014 and Gigaword 5). We show this by evaluating the semantic similarity of static word embeddings. Several common benchmark datasets are used: WS353 (Finkelstein et al., 2002), MEN (Bruni et al., 2014), RW (Luong et al., 2013), SimLex (Hill et al., 2015), RG-65 (Rubenstein & Goodenough, 1965), Mturk (Halawi et al., 2012). These datasets contain pairs of words with human-annotated similarity scores between them. Following previous work (Tissier et al., 2019; Sokal, 1958), cosine similarity is used for calculating the similarity between embeddings produced by each of the two words in the pair, and Spearman's correlation coefficient is calculated between this similarity and the human annotated score. The results are reported in Table 1.

Dataset	Ours	GloVe
MEN	57.0	**69.5**
WS353	**64.6**	64.0
SIMLEX	24.2	**31.5**
RW	41.3	**46.8**
RG	69.0	**74.2**
Mturk	**60.8**	57.5

Table 1: Evaluation on word similarity datasets. For each dataset Spearman's rank correlation coefficient is reported; the best score is in **bold**. GloVe 300d embeddings are used. Hyperparameter settings for our model: $K = 400$, $w = 11$.

We observe that our binary word embeddings demonstrate competitive performance compared to GloVe. Specifically, the fruit fly network demonstrates slightly higher correlation scores compared to GloVe on WS353 and Mturk, and slightly lower scores on the remaining four datasets. This indicates that the fruit fly network can learn meaningful binary semantic representations directly from raw text. An added advantage of our embeddings is that they require only a fraction (approx. 3%) of the memory footprint required for continuous word embeddings (assuming they have the same legth), since a real value requires 32-bits per vector element, whereas a boolean value requires only 1-bit of memory. We have also compared the performace of our word embeddings with the performance of binarized GloVe using a variety of common binarization techniques, see section 7 of the supplement for details.

3.2 WORD CLUSTERING

A nice aspect of the proposed fruit fly embeddings is that they result in tighter and better separated clusters than the conventional GloVe embeddings. To evaluate this property we started with hash codes for individual words and performed agglomerative clustering via complete link, using the cosine distance as the metric. The clustering algorithm was terminated at 200 clusters (we experimented with possible choices of this parameter, such as $200, 500, 1000, 2000, 3000, 5000$, and arrived at similar conclusions). The results for the inter-cluster similarity vs. intra-cluster similarity are shown in Fig. 3 (panel A). It is clear from this scatter plot that the average distance between the points within a cluster is smaller (higher similarity) for fruit fly embeddings compared to GloVe embeddings. At the same time, the distance between the closest clusters is larger or equal (smaller similarity). This property makes fruit fly embeddings appealing from the clustering perspective. We have also inspected the semantic structure of the clusters obtained this way. An example of the hierarchical clustering diagram (lower part of the tree containing 42 leaves) is shown in Fig. 3 (panel B). We clearly observe semantically coherent clusters resulting from the fruit fly word embeddings.

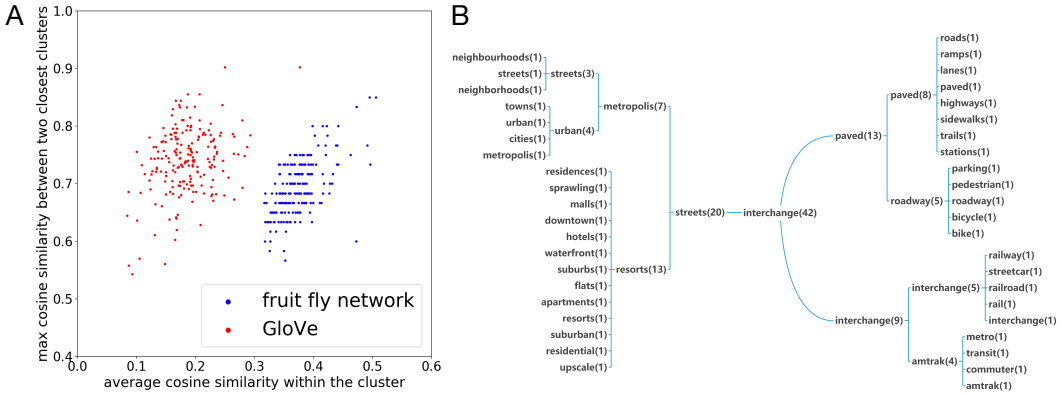

Figure 3: Panel A: maximum cosine similarity (minimal distance) to a point from the closest cluster vs. average cosine similarity between the points within the cluster. Panel B: an example of a cluster generated by the agglomerative clustering, the integer number associated with each node corresponds to the number of daughter leaves in that cluster. The root node corresponds to "interchange (42)".

3.3 Context-Dependent Word Embeddings

Here, we evaluate the effectiveness of our fruit fly inspired approach for contextual word embeddings, as opposed to static (or context-independent) embeddings from above. We use the WiC (Pilehvar & Camacho-Collados, 2018) and SCWS (Huang et al., 2012) benchmarks for the evaluation of context-sensitive word embeddings for word sense disambiguation. Both the datasets comprise pairs of sentences that contain a target word, and the task is to determine whether the two target words share a similar semantic meaning in the corresponding contexts. The WiC dataset is modeled as a binary prediction task, with 1 denoting that the target words have the same sense, and 0 indicating that they mean different things. The SCWS dataset is modeled as a rank prediction task, since for each pair of sentences and target words, it reports the average human similarity scores (from 10 Amazon Mechanical Turkers per pair).

word in context	10 nearest neighbor words in the hash code space
design of apple latest iphone	design, web, features, graphics, radeon, android, apps, ios, apple, nvidia
filling sweet apple pie recipe	chocolate, sweet, crispy, noodles, syrup, coconut, cheese, sauce, butter, cinnamon
money in bank checking account	money, credit, loans, account, cash, funds, savings, paying, banks, pension
boat on the bank of the river	river, lake, near, island, creek, canyon, valley, mountains, area, shore

Figure 4: For every word (highlighted in green) in context (left), 10 nearest neighbor words in the binary hashing space are shown (right). Context allows the algorithm to disambiguate the target word's meaning.

Before presenting quantitative results, we qualitatively examine how the fruit fly network performs on context sentence pairs for target words "apple" and "bank" in Fig. 4. We show the top $q = 10$ nearest neighbor words for the context dependent target word. These examples clearly indicate that the "correct" sense of the word has been found ("apple" the device manufacturer has different nearest neighbors from the fruit, and "bank" the financial institution from the natural feature).

For the quantitative comparison, we contrast our method against contextual embeddings from BERT (Devlin et al., 2018) and GloVe (Pennington et al., 2014). For BERT we use the 768-dimensional embeddings from the uncased-large model, and for GloVe we use the 300-dimensional embeddings, both from Spacy (`www.spacy.io`). Since BERT outputs contextual embeddings for each word in a sentence, we simply compute the cosine similarity between the embedding vectors for the target words for each pair of instances. For GloVe, we use a context window of size w centered at each of the target words and compute the average embedding for each window and compute the cosine similarity between the two window vectors. For the fruit fly network, given the effectiveness of the top-q nearest neighbor words (as seen in Fig. 4), we devise a two component scoring function. The first component is the dot-product between the context-dependent hash codes for the

two target words plus w length context blocks, denoted J_{dot}. The second is the number of common contextual nearest neighbors of the two target words among the top-q neighbors of each (scaled to be between 0 and 1), denoted J_{nn}. The final score is given as $J = \alpha \cdot J_{dot} + (1 - \alpha) \cdot J_{nn}$, where $\alpha \in [0, 1]$ is a hyperparameter. For all the methods, we predict a WiC pair to be positive if the score is above a threshold value θ. For SCWS, the ranking is proportional to the scores above θ, with the rest scored as zero. The hyperparameter θ is tuned for all the methods independently. Finally, for a fair comparison, all methods use the same 20k vocabulary.

We report the performance of our context-dependent word embeddings for both SCWS and WiC in Table 2 and Table 3, respectively. For both benchmarks we report the results from a 5-fold cross-validation study, where each fold (in turn) is used as a development set, and the remaining four folds as the test set. We select the optimal hyperparameters (including θ, α, q, k, w) for all the methods using only the first fold; no training is done since we evaluate only the pretrained embeddings. The tables report the Spearman rank correlation on SCWS, and the accuracy on WiC.

Method	mean	std
BERT	**54.6**	0.54
GloVe ($w = 3$)	40.9	1.3
GloVe ($w = 0$)	54.4	0.10
Ours ($w = 0$)	49.1	0.36

Table 2: SCWS dataset: mean and std for Spearman rank correlation. The best window value is also shown.

Method	mean	std
BERT	**61.2**	0.22
GloVe ($w = 7$)	54.9	0.26
GloVe ($w = 0$)	50.1	0.25
Ours ($w = 21$)	57.7	0.27

Table 3: WiC dataset: mean and std for accuracy. The best window value is also shown.

On SWCS (Table 2), while BERT performs the best, we see that the context-independent embeddings (using $w = 0$) are better for both GloVe and our method. The reason is that about 86.5% of the word pairs in SCWS are different words, and can be distinguished without looking at the context. Unlike SCWS, the WiC benchmark uses the same target word (with only minor variations in some cases) in both contexts, and therefore a context-independent approach is not expected to perform well. Indeed, on WiC (Table 3), we clearly observe that context-independent GloVe vectors ($w = 0$) are not very good, and our method, that uses the joint scoring function J combining both the hash code and nearest neighbor scores, is better than context-dependent GloVe ($w = 7$). Finally, for both SCWS and WiC our results are also relatively close to BERT.

Model	Accuracy
GloVe(100d)	80.0
GloVe(300d)	80.5
Ours ($k = 30$)	81.5
Ours ($k = 50$)	**81.8**

Table 4: Classification accuracy for word embeddings on the Stanford Sentiment Treebank data set.

Model	Accuracy
GloVe (100d)	80.3
GloVe (300d)	**81.4**
Ours ($k = 30$)	80.4
Ours ($k = 50$)	80.2

Table 5: Classification accuracy for different word embeddings on the 20 Newsgroup dataset.

3.4 Document Classification

We also compare our binary embeddings with GloVe (Pennington et al., 2014) on a document classification task. The benchmarks we use are 20 Newsgroups (20NewsGroups) and the Stanford Sentiment Treebank (Socher et al., 2013). The 20 Newsgroups data set contains around 18,000 documents, partitioned evenly into 20 different groups, and the Stanford Sentiment Treebank data set contains movie reviews reflecting their sentiment as either positive or negative. We use the TextCNN (Kim, 2014) classifier that uses the pre-trained fruit-fly or GloVe embeddings. For fair comparison, we use the same model parameters (e.g., kernel size, filter dimension) while testing different embeddings. The results in Table 4 and Table 5 show that our sparse binary encoding has similar performance to GloVe.

4 Computational Complexity

The computational complexity of our method can be evaluated by analyzing equations (2,3) for the weight updates. In these equations $\mathbf{v}^{\mathbf{A}}$ is a sparse vector, which has only w non-zero elements in it.

device	K	batch-size	GPU mem	time
V100 $\times$ 3	400	2000$\times$ 3	122MB	17m
V100 $\times$ 3	400	10000$\times$ 3	150MB	8m
V100 $\times$ 3	600	2000$\times$ 3	232MB	24m
V100 $\times$ 3	600	10000*3	267MB	11.5m
CPU 44cores	400	2000	-	76m
CPU 44cores	400	10000	-	25m

Table 6: Training time (per epoch) and memory footprint of our method on GPUs and CPUs. For the GPU implementation, three V100 GPUs interconnected with 100GB/s (bi-directional) NVLink® were used. For the CPU implementation, the computation was done on two 22-core CPUs. CPU memory is 137GB. The results are reported for window $w = 11$.

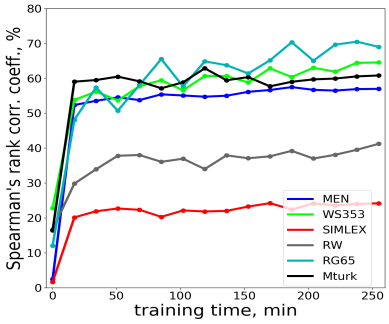

Figure 5: Spearman's correlation on word similarity datasets (see Section 3.1) vs. training time. Each point is one epoch.

Thus, for a minibatch of size $|BS|$, the computational complexity of evaluating the dot product with weights is $K \cdot w \cdot |BS|$. Additionally, the argmax operation requires $K \cdot |BS|$ operations. We will assume that the largest parameters in our model are the size of the corpus $|A| \approx 10^{10}$, and the size of the vocabulary $N_{\text{voc}} \approx 10^4 - 10^5$. Additionally we use large minibatches $|BS| \approx N_{\text{voc}}$. Calculation of the second term in (2) requires $O(K \cdot N_{\text{voc}})$ operations in addition to $K \cdot w \cdot |BS|$ operations for calculating the dot-product for each data point. Since the algorithm has to go over the entire corpus, this computation needs to be repeated $|A|/|BS|$ times per epoch. Thus, the overall computational complexity of our method is $O\Big(K \cdot |A|(w + N_{\text{voc}}/|BS|)\Big) \approx K \cdot |A| \cdot w$ per epoch. Thus, in the leading order it does not grow with the size of the vocabulary, which is a nice computational feature.

From the practical perspective, typical wall-clock training time and memory requirements per epoch are shown in Table 6. As is shown in Fig. 5, accurate solutions are obtained after about $2 - 3$ epochs; improvements beyond that are relatively small. Thus, our algorithm is capable of producing competitive models in a couple of hours. Contrast this with approximately 24 hours training time for GloVe (Pennington et al., 2014); 4 days of training on 16 TPUs for $\text{BERT}_{\text{BASE}}$; and 4 days on 64 TPUs for $\text{BERT}_{\text{LARGE}}$ (Devlin et al., 2018) (the last two numbers assume training corpus of size 250B tokens vs. 6B tokens considered in this paper). The record breaking training time of 47 minutes for BERT requires the use of 1472 NVIDIA® V100 GPUs each with 32GB of memory and a specialized DGX™ server architecture (Narasimhan, 2019). These are substantially larger computational resources than those required for training the fruit fly network.

5 Discussion and Conclusions

In this work we asked the intriguing question whether the core computational algorithm of one of the best studied networks in neuroscience – the network of KCs in the fruit fly brain – can be repurposed for solving a well defined machine learning task, namely, learning word embeddings from text. We have shown that, surprisingly, this network can indeed learn the correlations between the words and their context, and produce high quality word embeddings. On the semantic similarity task the fruit fly word embeddings stand in line with SOTA GloVe word embeddings (slightly outperforming GloVe on some of the benchmarks, and slightly loosing to GloVe on others, see Table 1). On the word-in-context task the fruit fly network outperforms GloVe by almost 3%, but looses to BERT by 3.5%, see Table 3. The small gap in classification accuracy compared with BERT, however, is outweighted by the benefit of requiring significantly smaller computational resources to obtain these fruit fly embeddings, as we have explained in Section 4, see Table 6. We view this result as an example of a general statement that biologically inspired algorithms might be more compute efficient compared with their classical (non-biological) counterparts, even if they slightly lose in terms of accuracy. We have also shown that the fruit fly embeddings result in much tighter and better separated clusters of semantic concepts, compared to those produced by GloVe embeddings, see Fig. 3.

REFERENCES

20NewsGroups. 20 newsgroups dataset. URL http://people.csail.mit.edu/jrennie/20Newsgroups/.

Alexander Shakeel Bates, Philipp Schlegel, Ruairí JV Roberts, Nikolas Drummond, Imaan FM Tamimi, Robert Gillies Turnbull, Xincheng Zhao, Elizabeth C Marin, Patricia Demetria Popovici, Serene Dhawan, et al. Complete connectomic reconstruction of olfactory projection neurons in the fly brain. *BioRxiv*, 2020.

David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.

Elia Bruni, Nam Khanh Tran, and Marco Baroni. Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49(1):1–47, January 2014. ISSN 1076-9757.

Sophie JC Caron, Vanessa Ruta, LF Abbott, and Richard Axel. Random convergence of olfactory inputs in the drosophila mushroom body. *Nature*, 497(7447):113–117, 2013.

Sophie Jeanne Cecile Caron, Jinzhi Li, Brennan Dale Mahoney, and Miles Solomon Jacob. Two parallel pathways convey distinct visual information to the drosophila mushroom body. *bioRxiv*, 2020.

Moses S. Charikar. Similarity Estimation Techniques from Rounding Algorithms. In *Annual ACM Symposium on Theory of Computing*, pp. 380–388, 2002. ISBN 978-1-58113-495-7. doi: 10.1145/509907.509965.

Sanjoy Dasgupta, Charles F. Stevens, and Saket Navlakha. A neural algorithm for a fundamental computing problem. *Science*, 358(6364):793–796, 2017.

Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. Indexing by latent semantic analysis. *Journal of the American society for information science*, 41 (6):391–407, 1990.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. Placing search in context: The concept revisited. *ACM Trans. Inf. Syst.*, 20(1), 2002. ISSN 1046-8188.

John R Firth. A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis*, 1957.

Aaron Gokaslan and Vanya Cohen. OpenWebText Corpus. http://Skylion007.github.io/OpenWebTextCorpus, 2019.

Yunchao Gong and Svetlana Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *CVPR 2011*, pp. 817–824. IEEE, June 2011. ISBN 978-1-4577-0394-2. doi: 10.1109/CVPR.2011.5995432.

Guy Halawi, Gideon Dror, Evgeniy Gabrilovich, and Yehuda Koren. Large-scale learning of word relatedness with constraints. *KDD*, pp. 1406–1414, 2012.

Zellig S Harris. Distributional structure. *Word*, 10(2-3):146–162, 1954.

Felix Hill, Roi Reichart, and Anna Korhonen. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695, December 2015.

Eric H. Huang, Richard Socher, Christopher D. Manning, and Andrew Y. Ng. Improving Word Representations via Global Context and Multiple Word Prototypes. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2012.

Yoon Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.

Dmitry Krotov and John J. Hopfield. Unsupervised learning by competing hidden units. *Proceedings of the National Academy of Sciences*, 116(16):7723–7731, 2019.

Feng Li, Jack Lindsey, Elizabeth C Marin, Nils Otto, Marisa Dreher, Georgia Dempsey, Ildiko Stark, Alexander Shakeel Bates, Markus William Pleijzier, Philipp Schlegel, et al. The connectome of the adult drosophila mushroom body: implications for function. *bioRxiv*, 2020.

Thang Luong, Richard Socher, and Christopher Manning. Better word representations with recursive neural networks for morphology. In *Conference on Computational Natural Language Learning*, pp. 104–113, 2013.

Elizabeth C Marin, Ruairí JV Roberts, Laurin Büld, Maria Theiss, Markus W Pleijzier, Tatevik Sarkissian, Willem J Laursen, Robert Gillies Turnbull, Philipp Schlegel, Alexander Shakeel Bates, et al. Connectomics analysis reveals first, second, and third order thermosensory and hygrosensory neurons in the adult drosophila brain. *BioRxiv*, 2020.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

Shar Narasimhan. Nvidia clocks world's fastest bert training time and largest transformer based model, paving path for advanced conversational ai. `https://devblogs.nvidia.com/training-bert-with-gpus/`, 2019.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.

Mohammad Taher Pilehvar and Jose Camacho-Collados. Wic: the word-in-context dataset for evaluating context-sensitive meaning representations. *arXiv preprint arXiv:1808.09121*, 2018.

Herbert Rubenstein and John B. Goodenough. Contextual correlates of synonymy. *Communications of the ACM*, 8(10):627–633, 1965.

Chaitanya K. Ryali, John J. Hopfield, Leopold Grinberg, and Dmitry Krotov. Bio-Inspired Hashing for Unsupervised Similarity Search. *arXiv preprint arXiv:2001.04907*, 2020.

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

RR Sokal. A statistical method for evaluating systematic relationships. *Univ. Kansas, Sci. Bull.*, 38: 1409–1438, 1958.

Julien Tissier, Christophe Gravier, and Amaury Habrard. Near-lossless binarization of word embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 7104–7111, 2019.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Katrin Vogt, Yoshinori Aso, Toshihide Hige, Stephan Knapek, Toshiharu Ichinose, Anja B Friedrich, Glenn C Turner, Gerald M Rubin, and Hiromu Tanimoto. Direct neural pathways convey distinct visual information to drosophila mushroom bodies. *Elife*, 5:e14009, 2016.

Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral Hashing. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou (eds.), *Advances in Neural Information Processing Systems 21*, pp. 1753–1760. 2009.

Zhihao Zheng, Feng Li, Corey Fisher, Iqbal J Ali, Nadiya Sharifi, Steven Calle-Schuler, Joseph Hsu, Najla Masoodpanah, Lucia Kmecova, Tom Kazimiers, et al. Structured sampling of olfactory input by the fly mushroom body. *bioRxiv*, 2020.

6 APPENDIX A. TRAINING PROTOCOLS AND HYPERPARAMETER CHOICES.

The fruit fly network was trained on the OpenWebText Corpus (Gokaslan & Cohen, 2019), which is a 32GB corpus of unstructured text containing approximately 6B tokens. Individual documents were concatenated and split into sentences. A collection of w-grams were extracted from each sentence by sliding a window of size w along each sentence from the beginning to the end. Sentences shorter than w were removed. The vocabulary was composed of $N_{\text{voc}} = 20000$ most frequent tokens in the corpus.

Training was done for N_{epoch}. At each epoch all the w-grams were shuffled, organized in mini-batches, and presented to the learning algorithm. The learning rate was linearly annealed starting from the maximal value ε_0 at the first epoch to nearly zero at the last epoch.

The training algorithm has the following hyperparameters: size of the KC layer K, window w, overall number of training epochs N_{epoch}, initial learning rate ε_0, minibatch size, and hash length k. All models presented in this paper were trained for $N_{\text{epoch}} = 15$. The optimal ranges of the hyperparameters are: learning rate is $\varepsilon_0 \approx 10^{-4} - 5 \cdot 10^{-4}$; $K \approx 200 - 600$; $w \approx 9 - 15$; minibatch size $\approx 2000 - 15000$; hash length k is reported for each individual experiment.

7 APPENDIX B. COMPARISON WITH BINARIZED GLOVE.

Method	Hash Length (k)						Hash Length (k)					
	4	8	16	32	64	128	4	8	16	32	64	128
	MEN (69.5/68.1)						**WS353** (64.0/47.7)					
Ours	**34.0**	**49.9**	**55.9**	56.7	55.3	51.3	**43.2**	**52.1**	**55.3**	**57.4**	**60.3**	51.7
LSH	16.9	23.7	35.6	42.6	53.6	63.4	8.2	20.7	30.0	34.7	43.9	50.3
FlyHash	27.5	37.7	46.6	**57.6**	**67.3**	**71.6**	20.9	32.9	41.9	48.4	57.6	61.7
ITQ	0.1	7.7	10.5	16.5	30.4	50.5	-6.6	-6.1	-2.4	-4.4	6.1	24.8
SH	9.4	17.0	22.9	37.6	52.9	65.4	15.4	14.1	19.5	32.3	43.1	58.4
PCAH	12.5	21.8	27.6	39.6	53.4	68.1	6.4	6.3	20.6	33.9	49.8	**62.6**
NLB	-	-	-	-	46.1	63.3	-	-	-	-	30.1	44.9
	SIMLEX (31.5/29.8)						**RW** (46.8/31.4)					
Ours	**13.4**	16.5	**22.8**	22.1	21.1	17.0	11.0	**22.6**	25.8	36.9	38.6	35.2
LSH	6.8	11.9	17.0	21.2	26.8	30.9	10.8	16.3	21.8	27.8	36.3	45.0
FlyHash	10.4	**17.2**	**22.8**	**28.5**	**32.4**	**35.2**	19.9	21.3	**30.9**	**40.5**	**47.6**	**53.3**
ITQ	7.0	1.6	4.3	5.5	11.8	18.2	13.7	5.3	6.6	6.9	12.5	26.5
SH	9.3	15.6	15.9	17.0	23.1	31.2	**22.6**	21.5	24.3	28.8	36.1	45.8
PCAH	4.4	10.3	11.0	17.3	24.1	31.6	12.4	16.7	21.5	30.3	36.9	44.4
NLB	-	-	-	-	20.5	31.4	-	-	-	-	25.1	34.3
	RG (74.2/67.6)						**Mturk** (57.5/61.9)					
Ours	24.0	40.4	**51.3**	62.3	63.2	55.8	**44.0**	**49.0**	**52.2**	**60.1**	57.7	55.2
LSH	21.2	35.4	44.6	55.1	63.1	70.1	16.0	23.1	33.2	35.6	42.7	55.5
FlyHash	**36.6**	**49.0**	49.5	**66.1**	**69.6**	70.9	29.3	35.8	41.4	50.4	**59.0**	**61.6**
ITQ	-17.5	-8.9	26.3	41.7	50.5	66.2	9.9	7.8	10.1	17.7	32.8	47.3
SH	4.5	5.8	20.3	42.9	61.3	**72.6**	18.9	17.6	27.5	35.45	48.1	57.9
PCAH	1.9	9.6	19.8	40.9	53.3	68.2	15.5	15.1	27.1	41.7	46.5	56.2

Table 7: Evaluation on word similarity datasets. For each dataset and hash length, the best (second best) score is in **bold** (underlined). The performance for GloVe embeddings is reported next to the name of each dataset in the format 300d/100d. Spearman's rank correlation coefficient is reported for common baselines that binarize GloVe (300d) embeddings together with our results. Hyperparameter settings for our algorithm: $K = 400$, $w = 11$.

Our aim here is to demonstrate that the fruit fly word embeddings are competitive with existing state-of-the-art binarization methods applied to pretrained GloVe embeddings. We show this by evaluating the semantic similarity of static word embeddings, using several common benchmark datasets: WS353 (Finkelstein et al., 2002), MEN (Bruni et al., 2014), RW (Luong et al., 2013), SimLex (Hill et al., 2015), RG-65 (Rubenstein & Goodenough, 1965), and Mturk (Halawi et al., 2012). These datasets contain pairs of words with human-annotated similarity scores between them. We also include a comparison with GloVe (Pennington et al., 2014) word embeddings (trained on

Wiki2014 and Gigaword 5). Since our representations are binary (in contrast to GloVe), we binarize GloVe embeddings and report their performance using a number of common hashing methods, LSH/SimHash (Charikar, 2002) (random *contractive* projections followed by binarization based on sign), FlyHash (Dasgupta et al., 2017) (random *expansive* projections followed by k-winner take all binarization), ITQ (Gong & Lazebnik, 2011) (iterative quantization), SH (spectral hashing) (Weiss et al., 2009), PCAH (Gong & Lazebnik, 2011) (PCA followed by binarization based on sign). Where available, we include evaluation from NLB (Tissier et al., 2019) (autoencoder-based binarization).

Following previous work (Tissier et al., 2019; Sokal, 1958), model similarity score for binary representations is evaluated as $sim(v_1, v_2) = (n_{11} + n_{00})/n$, where n_{11} (n_{00}) is the number of bits in v_1 and v_2 that are both 1 (0), and n is the length of $v_{1,2}$. Cosine similarity is used for real-valued representations. The results are reported in Table 7. For each dataset, we report performance across a range of hash lengths $\{4, 8, 16, 32, 64, 128\}$. For methods that incorporate randomness (LSH, FlyHash, ITQ), we report the average across 5 runs. As is clear from Table 7, our binary word embeddings demonstrate competitive performance compared to published methods for GloVe binarization, and our algorithm can learn meaningful binary semantic representations directly from raw text. Importantly, our algorithm does not require training GloVe embeddings first before binarizing them.

8 Appendix C. Details of Technical Implementation.

From the practical perspective, efficient implementation of the learning algorithm for the fruit fly network requires the use of sparse algebra, atomic updates, and block-sparse data access. Our algorithm is implemented in CUDA as a back-end, while python is used as an interface with the main functions.

The typical memory footprint of our approach is very small. About $100 - 270$MB GPU memory is allocated for the operators $W_{\mu i}$, $\mathbf{v}^{\mathbf{A}}$ and temporary fields; while approximately 140GB CPU memory is needed to store the input data, array of random numbers for shuffle operations and shuffled indices. For GPU implementation, the model data is stored in the GPU's memory, while the input data is stored in the CPU memory. The parallelization strategy in our implementation is based on two aspects. First, each minibatch of data is divided into smaller sub-minibatches which are processed on different GPUs. Second, all the operations (dense-sparse matrix multiplications, $\arg\max$ operation, and weight updates) are executed in parallel using multiple threads.