



The Spectral Underpinning of word2vec

Ariel Jaffe^{1*}, Yuval Kluger^{1,2,3}, Ofir Lindenbaum¹, Jonathan Patsenker¹, Erez Peterfreund⁴ and Stefan Steinerberger⁵

¹Program in Applied Mathematics, Yale University, New Haven, CT, United States, ²Department of Pathology, Yale School of Medicine, New Haven, CT, United States, ³Interdepartmental Program in Computational Biology and Bioinformatics, New Haven, CT, United States, ⁴School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel, ⁵Department of Mathematics, University of Washington, Seattle, WA, United States

Word2vec introduced by Mikolov et al. is a word embedding method that is widely used in natural language processing. Despite its success and frequent use, a strong theoretical justification is still lacking. The main contribution of our paper is to propose a rigorous analysis of the highly nonlinear functional of word2vec. Our results suggest that word2vec may be primarily driven by an underlying spectral method. This insight may open the door to obtaining provable guarantees for word2vec. We support these findings by numerical simulations. One fascinating open question is whether the nonlinear properties of word2vec that are not captured by the spectral method are beneficial and, if so, by what mechanism.

OPEN ACCESS

Edited by:

Dabao Zhang, Purdue University, United States

Reviewed by:

Yuguang Wang, Max-Planck-Gesellschaft (MPG), Germany
Halim Zeghdoudi, University of Annaba, Algeria

*Correspondence:

Ariel Jaffe
a.jaffe@yale.edu

Specialty section:

This article was submitted to Mathematics of Computation and Data Science, a section of the journal Frontiers in Applied Mathematics and Statistics

Received: 10 August 2020

Accepted: 21 October 2020

Published: 03 December 2020

Citation:

Jaffe A, Kluger Y, Lindenbaum O, Patsenker J, Peterfreund E and Steinerberger S (2020) The Spectral Underpinning of word2vec. *Front. Appl. Math. Stat.* 6:593406. doi: 10.3389/fams.2020.593406

Keywords: dimensionality reduction, word embedding, spectral method, word2vec, skip-gram model, nonlinear functional

1 INTRODUCTION

Word2vec was introduced by Mikolov et al. [1] as an unsupervised scheme for embedding words based on text corpora. We will try to introduce the idea in the simplest possible terms and refer to [1–3] for the way it is usually presented. Let $\{x_1, x_2, \dots, x_n\}$ be a set of elements for which we aim to compute a numerical representation. These may be words, documents, or nodes in a graph. Our input consists of an $n \times n$ matrix P with non-negative elements P_{ij} , which encode, by a numerical value, the relationship between the set $\{x_i\}_{i=1}^n$ and a set of context elements $\{c_j\}_{j=1}^n$. The meaning of contexts is determined by the specific application, where in most cases the set of contexts is equal to the set of elements (i.e. $c_i = x_i$ for any $i \in \{1, \dots, n\}$) [2].

The larger the value of P_{ij} , the larger the connection between x_i and c_j . For example, such a connection can be quantified by the probability that a word appears in the same sentence as another word. Based on P , Mikolov defined an energy function which depends on two sets of vector representations $\{w_1, \dots, w_n\}$ and $\{v_1, \dots, v_n\}$. Maximizing the functional with respect to these sets yields $\{w'_1, \dots, w'_n\}$ and $\{v'_1, \dots, v'_n\}$ which can serve as a low dimensional representations for the words and contexts respectively. Ideally, this embedding should encapsulate the relations captured by the matrix P .

Assuming a uniform prior over the n elements, the energy function $L : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$, introduced by Mikolov et al. [1] can be written as

$$L(w, v) = \langle w, Pv \rangle - \sum_{i=1}^n \log \left(\sum_{j=1}^n \exp(w_i v_j) \right). \quad (1)$$




FIGURE 1 | Illustration of point set drawn from two distinct Gaussian distributions. The result of maximizing over the word2vec functional (black) is closely tracked (up to scale) by the optimizer of the spectral method (blue) and the eigenvector (red). In **Figure 7**, we present a scatter plot comparing the values of \hat{w} and $\sqrt{\lambda n} \cdot u$.

The exact relation between **1** and the formulation in [4] appears in **Supplementary Material**. Word2vec is based on maximizing this expression over all $(w, v) \in \mathbb{R}^n \times \mathbb{R}^n$

$$(w^*, v^*) = \operatorname{argmax}_{(w,v)} L(w, v).$$

There is no reason to assume that the maximum is unique. It has been observed that if x_i and x_j are similar elements in the data set (namely, words that frequently appear in the same sentence), then v_i^*, v_j^* or w_i^*, w_j^* tend to have similar numerical values. Thus, the values $\{w_1^*, \dots, w_n^*\}$ are useful for embedding $\{x_1, \dots, x_n\}$. One could also try to maximize the symmetric loss that arises from enforcing $w = v$ and is given by $L : \mathbb{R}^n \rightarrow \mathbb{R}$

$$L(w) = \langle w, Pw \rangle - \sum_{i=1}^n \log \left(\sum_{j=1}^n \exp(w_i w_j) \right). \quad (2)$$

In **Section 5** we show that the symmetric functional yields a meaningful embedding for various datasets. Here, the interpretation of the functional is straight-forward: we wish to pick $w \in \mathbb{R}^n$ in a way that makes $\langle w, Pw \rangle$ large. Assuming P is diagonalizable, this is achieved for w that is a linear combination of the leading eigenvectors. At the same time, the exponential function places a penalty over large entries in w .

Our paper initiates a rigorous study of the energy functional $L(w)$, however, we emphasize that a complete description of the energy landscape $L(w)$ remains an interesting open problem. We also emphasize that our analysis has direct implications for computational aspects as well: for instance, if one were interested in maximizing the nonlinear functional, the maximum of its linear approximation (which is easy to compute) is a natural starting point. A simple example is shown in **Figure 1**: the underlying dataset contains 200 points in \mathbb{R}^{10} where the first 100 points are drawn from a Gaussian distribution, and the second 100 points are drawn from a second Gaussian distribution. The matrix P is the row-

stochastic matrix induced by a Gaussian kernel $K_{ij} = \exp(-\|x_i - x_j\|^2/\alpha)$ where α is a scaling parameter discussed in **Section 5**. We observe that, up to scaling, the maximizer of the energy functional (black) is well approximated by the spectral methods introduced below.

2 MOTIVATION AND RELATED WORKS

Optimizing over energy functions such as **1** to obtain vector embeddings is done for various applications, such as words [4], documents [5] and graphs [6]. Surprisingly, very few works addressed the analytic aspects of optimizing over the word2vec functional. Hashimoto et. al. [7] derived a relation between word2vec and stochastic neighbor embedding [8]. Cotterell et al. [9] showed that when P is sampled according to a multinomial distribution, optimizing over **1** is equivalent to exponential family PCA [10]. If the number of elements is large, optimizing over **1** becomes impractical. As an efficient alternative, Mikolov et al. [4] suggested a variation based on negative sampling. Levy and Goldberg [11] showed that if the embedding dimension is sufficiently high, then optimizing over the negative sampling functional suggested in [4] is equivalent to factorizing the shifted Pointwise Mutual Information matrix. This work was extended in [12], where similar results were derived for additional embedding algorithms such as [3, 13, 14]. Decomposition of the PMI matrix was also justified by Arora et al. [15], based on a generative random walk model. A different approach was introduced by Landgraf [16], that related the negative sampling loss function to logistic PCA.

In this work, we focus on approximating the highly nonlinear word2vec functional by Taylor expansion. We show that in the regime of embedding vectors with small magnitude, the functional can be approximated by the spectral decomposition of the matrix P . This draws a natural connection between word2vec and classical, spectral embedding methods such as [17, 18]. By rephrasing word2vec as a spectral method in the “small vector limit,” one gains access to a large number of tools that allow one to rigorously establish a framework under which word2vec can enjoy provable guarantees, such as in [19, 20].

3 RESULTS

We now state our main results. In **Section 3.1** we establish that the energy functional $L(w, v)$ has a nice asymptotic expansion around $(v, w) = (0, 0) \in \mathbb{R}^n \times \mathbb{R}^n$ and corresponds naturally to a spectral method in that regime. Naturally, such an asymptotic expansion is only feasible if one has some control over the size of the entries of the extremizer. We establish in **Section 3.2** that the vectors maximizing the functional are not too large. The results in **Section 3.2** are closely matched by numerical results: in particular, we observe that $\|w\| \sim \sqrt{n}$ in practice, a logarithmic factor smaller than our upper bound. The proofs are given in **Section 4** and explicit numerical examples are shown

in Section 5. In Section 6 we show empirically that the relation between word2vec and the spectral approach holds also for embedding in more than one dimension.

3.1 First Order Approximation for Small Data

The main idea is simple: we make an ansatz assuming that the optimal vectors are roughly of size $\|w\|, \|v\| \sim 1$. If we assume that the vectors w, v are fairly “typical” vectors of size ~ 1 , then each entry is expected to scale approximately as $\sim n^{-1/2}$. Our main observation is that this regime is governed by a regularized spectral method. Before stating our theorem, let \lesssim denote the inequality up to universal multiplicative constants.

Theorem 3.1 (Spectral Expansion). If $\|v\|_\infty, \|w\|_\infty \lesssim n^{-1/2}$, then

$$L(w, v) = \langle w, Pv \rangle - \frac{1}{n} \left(\sum_{i=1}^n w_i \right) \left(\sum_{j=1}^n v_j \right) - \frac{1}{n} \sum_{i,j=1}^n \frac{w_i^2 v_j^2}{2} - n \log n + \mathcal{O}(n^{-1}).$$

Naturally, since we are interested in maximizing this quantity, the constant factor $n \log n$ plays no role. The leading terms can be rewritten as

$$\langle w, Pv \rangle - \frac{1}{n} \left(\sum_{i=1}^n w_i \right) \left(\sum_{j=1}^n v_j \right) = \left\langle w, \left(P - \frac{1}{n} \mathbf{1} \right) v \right\rangle,$$

where $\mathbf{1}$ is the matrix all of whose entries are 1. This suggests that the optimal v, w maximizing the quantity should simply be the singular vectors associated to the matrix $P - \frac{1}{n} \mathbf{1}$. The full expansion has a quadratic term that serves as an additional regularizer. The symmetric case (with ansatz $v = w$) is particularly simple, since we have

$$L(w) = \langle w, Pw \rangle - \frac{1}{n} \left(\sum_{i=1}^n w_i \right)^2 - \frac{\|w\|^4}{2n} - n \log n + \mathcal{O}(n^{-1}).$$

Assuming P is similar to a symmetric matrix, the optimal w should be well described by the leading eigenvector of $\left(P - \frac{1}{n} \mathbf{1} \right)$ with an additional regularization term ensuring that $\|w\|$ is not too large. We consider this simple insight to be the main contribution of this paper, since it explains succinctly why an algorithm like word2vec has a chance to be successful. We also give a large number of numerical examples showing that in many cases the result obtained by word2vec is extremely similar to what we obtain from the associated spectral method.

3.2 Optimal Vectors Are Not Too Large

Another basic question is as follows: how large is the norm of the vector(s) maximizing the energy function? This is of obvious importance in practice, however, as seen in Theorem 3.1, it also has some theoretical relevance: if w has large entries, then clearly one cannot hope to capture the exponential nonlinearity with a

polynomial expansion. Assuming $\|P\| \leq 1$, the global maximizer w^* of the second-order approximation

$$L_2(w) = \langle w, Pw \rangle - \frac{1}{n} \left(\sum_{i=1}^n w_i \right)^2 - \frac{\|w\|^4}{2n} - n \log n, \tag{3}$$

satisfies

$$\|w\|^* \leq \sqrt{2n}.$$

This can be seen as follows: if $\|P\| \leq 1$, then $\langle w, Pw \rangle \leq \|w\|^2$. Plugging in $w = 0$ shows that the maximal energy is at least size $-n \log n$. For any vector exceeding $\sqrt{2n}$ in size, we see that the energy is less than that establishing the bound. We obtain similar boundedness properties for the fully nonlinear problem for a fairly general class of matrices.

Theorem 3.2 (Generic Boundedness.). Let $P \in \mathbb{R}^{n \times n}$ satisfy $\|P\| < 1$. Then

$$w = \operatorname{argmax}_w \langle w, Pw \rangle - \sum_{i=1}^n \log \left(\sum_{j=1}^n \exp(w_i w_j) \right),$$

satisfies

$$\|w\|^2 \leq \frac{n \log n}{1 - \|P\|}.$$

While we do not claim that this bound is sharp, however it does nicely illustrate that the solutions of the optimization problem must be bounded. Moreover, if they are bounded, then so are their entries; more precisely, $\|w\|^2 \lesssim n$ implies that, for ‘flat’ vectors, the typical entry is of size $\lesssim 1$ and thus firmly within the approximations that can be reached by a Taylor expansion. It is clear that a condition such as $\|P\| < 1$ is required for boundedness of the solutions. This can be observed by considering the row-stochastic matrix

$$P = \begin{pmatrix} 1 - \varepsilon & \varepsilon \\ \varepsilon & 1 - \varepsilon \end{pmatrix}.$$

Writing $w = (w_1, w_2)$, we observe that the arising functional is quite nonlinear even in this simple case. However, it is fairly easy to understand the behavior of the gradient ascent method on the w_1 -axis since

$$\frac{\partial}{\partial w_1} L(w_1, w_2) |_{w_2=0} = 2w_1 \left(1 - \varepsilon - \frac{e^{w_1^2}}{1 + e^{w_1^2}} \right),$$

is monotonically increasing until $w_1 \sim \pm \sqrt{\log \varepsilon^{-1}}$. Therefore it is, a priori, unbounded since ε can be arbitrarily close to 0.

In practice, one often uses word2vec for matrices whose spectral norm is $\|P\| = 1$ and which have the additional property of being row-stochastic. We also observe empirically that the global optimizer w^* has a mean value close to 0 (the expansion in Theorem 3.1 suggests why this would be the case). We achieve a similar boundedness theorem in which the only relevant operator norm is that of the operator restricted to the subspace of vectors having mean 0.

Theorem 3.3 (Boundedness for row-stochastic matrices). Let $P \in \mathbb{R}^{n \times n}$ be a row-stochastic matrix and let

$$P_S : \{w \in \mathbb{R}^n : w_1 + \dots + w_n = 0\} \rightarrow \mathbb{R}^n,$$

denote the restriction of P to that subspace and suppose that $\|P_S\| < 1$. Let

$$w = \arg \max_w \langle w, Pw \rangle - \sum_{i=1}^n \log \left(\sum_{j=1}^n \exp(w_i w_j) \right).$$

If w has a mean value sufficiently close to 0,

$$\left| \left\langle w, \frac{\mathbf{1}}{\sqrt{n}} \right\rangle \right| \leq \frac{1 - \|P_S\|}{3} \|w\|,$$

where $\mathbf{1} = (1, 1, 1, \dots, 1)$, then

$$\|w\|^2 \leq \frac{2n \log n}{1 - \|P_S\|}.$$

The 2×2 matrix given above, illustrates that some restrictions are necessary, in order to obtain a nicely bounded gradient ascent. There is some freedom in the choice of the constants in Theorem 3.3. Numerical experiments show that the results are not merely theoretical: extremizing vectors tend to have a mean value sufficiently close to 0 for the theorem to be applicable.

3.3 Outlook

Summarizing, our main arguments are as follows:

- (1) The energy landscape of the word2vec functional is well approximated by a spectral method (or regularized spectral method) as long as the entries of the vector are uniformly bounded. In any compact interval around 0, the behavior of the exponential function can be appropriately approximated by a Taylor expansion of sufficiently high degree.
- (2) There are bounds that suggests that the energy of the embedding vector scale as $\sqrt{n \log n}$; this means that, for “flat” vectors, the individual entries grow at most like $\sqrt{\log n}$. Presumably this is an artifact of the proof.
- (3) Finally, we present examples in **Section 4** showing that in many cases the embedding obtained by maximizing the word2vec functional are indeed accurately predicted by the second order approximation.

This suggests various interesting lines of research: it would be nice to have refined versions of Theorems 3.2 and 3.3 (an immediate goal being the removal of the logarithmic dependence and perhaps even pointwise bounds on the entries of w). Numerical experiments indicate that Theorems 3.2 and 3.3 are at most a logarithmic factor away from being optimal. A second natural avenue of research proposed by our paper is to differentiate the behavior of word2vec and that of the associated spectral method: are the results of word2vec (being intrinsically nonlinear) truly different from the behavior of the spectral method (arising as its linearization)? Or, put differently, is the nonlinear aspect of word2vec that is *not* being captured by the spectral method helpful for embedding?

4 PROOFS

Proof of Theorem 3.1. We recall our assumption of $\|w\|_\infty \lesssim n^{-1/2}$ and $\|v\|_\infty \lesssim n^{-1/2}$ (where the implicit constant affects all subsequent constants). We remark that the subsequent arguments could also be carried out for any $\|w\|_\infty, \|v\|_\infty \lesssim n^{-\epsilon}$ at the cost of different error terms; the arguments fail to be rigorous as soon as $\|w\|_\infty \sim 1$, since then, a priori, all terms in the Taylor expansion of e^x could be of roughly the same size. We start with the Taylor expansion

$$\begin{aligned} \sum_{j=1}^n e^{w_i v_j} &= \sum_{j=1}^n \left(1 + w_i v_j + \frac{w_i^2 v_j^2}{2} + \mathcal{O}(n^{-3}) \right) \\ &= n + \sum_{j=1}^n \left(w_i v_j + \frac{w_i^2 v_j^2}{2} \right) + \mathcal{O}(n^{-2}). \end{aligned}$$

In particular, we note that

$$\left| \sum_{j=1}^n \left(w_i v_j + \frac{w_i^2 v_j^2}{2} \right) \right| \lesssim 1.$$

We use the series expansion

$$\log(n + x) = \log n + \frac{x}{n} - \frac{x^2}{2n^2} + \mathcal{O}\left(\frac{|x|^3}{n^3}\right)$$

to obtain

$$\begin{aligned} \log \left(\sum_{j=1}^n e^{w_i v_j} \right) &= \log n + \frac{1}{n} \sum_{j=1}^n \left(w_i v_j + \frac{w_i^2 v_j^2}{2} \right) \\ &\quad - \frac{1}{2n^2} \left(\sum_{j=1}^n w_i v_j + \frac{w_i^2 v_j^2}{2} \right)^2 + \mathcal{O}(n^{-3}). \end{aligned}$$

Here, the second sum can be somewhat simplified since

$$\begin{aligned} \frac{1}{2n^2} \left(\sum_{j=1}^n w_i v_j + \frac{w_i^2 v_j^2}{2} \right)^2 &= \frac{1}{2n^2} \left(\sum_{j=1}^n (w_i v_j + \mathcal{O}(n^{-2})) \right)^2 \\ &= \frac{1}{2n^2} \left(\mathcal{O}(n^{-1}) + \sum_{j=1}^n w_i v_j \right)^2 \\ &= \frac{1}{2n^2} \left(\sum_{j=1}^n w_i v_j \right)^2 + \mathcal{O}(n^{-3}) \\ &= \frac{w_i^2}{2n^2} \left(\sum_{j=1}^n v_j \right)^2 + \mathcal{O}(n^{-3}) \end{aligned}$$

Altogether, we obtain that

$$\begin{aligned} \sum_{i=1}^n \log \left(\sum_{j=1}^n e^{w_i v_j} \right) &= \sum_{i=1}^n \left(\log n + \frac{1}{n} \sum_{j=1}^n \left(w_i v_j + \frac{w_i^2 v_j^2}{2} \right) \right. \\ &\quad \left. - \frac{w_i^2}{2n^2} \left(\sum_{j=1}^n v_j \right)^2 + \mathcal{O}(n^{-3}) \right) \end{aligned}$$

$$= n \log n + \frac{1}{n} \sum_{i,j=1}^n w_i v_j + \frac{1}{n} \sum_{i,j=1}^n \frac{w_i^2 v_j^2}{2} - \frac{1}{n^2} \left(\sum_{i=1}^n \frac{w_i^2}{2} \right) \left(\sum_{j=1}^n v_j \right)^2 + \mathcal{O}(n^{-2}).$$

Since $\|w\|_\infty, \|v\|_\infty \leq n^{-1/2}$, we have

$$\frac{1}{n^2} \left(\sum_{i=1}^n \frac{w_i^2}{2} \right) \left(\sum_{j=1}^n v_j \right)^2 \leq n^{-1}$$

and have justified the desired expansion.

Proof of Theorem 3.2. Setting $w = 0$ results in the energy

$$L(w) = -n \log n.$$

Now, let w be a global maximizer. We obtain

$$\begin{aligned} -n \log n &\leq \langle w, Pw \rangle - \sum_{i=1}^n \log \left(\sum_{j=1}^n e^{w_i v_j} \right) \\ &\leq \|P\| \|w\|^2 - \sum_{i=1}^n \log(e^{w_i^2}) \leq (\|P\| - 1) \|w\|^2 \end{aligned}$$

which is the desired result.

Proof of Theorem 3.3. We expand the vector w into a multiple of the constant vector of norm 1, the vector

$$\frac{1}{\sqrt{n}} = \left(\frac{1}{\sqrt{n}}, \frac{1}{\sqrt{n}}, \dots, \frac{1}{\sqrt{n}} \right),$$

and the orthogonal complement via

$$w = \left\langle w, \frac{1}{\sqrt{n}} \right\rangle \frac{1}{\sqrt{n}} + \left(w - \left\langle w, \frac{1}{\sqrt{n}} \right\rangle \frac{1}{\sqrt{n}} \right),$$

which we abbreviate as $w = \tilde{w} + (w - \tilde{w})$. We expand,

$$\begin{aligned} \langle w, Pw \rangle &= \langle \tilde{w}, P\tilde{w} \rangle + \langle \tilde{w}, P(w - \tilde{w}) \rangle + \langle w - \tilde{w}, P\tilde{w} \rangle + \langle w - \tilde{w}, P(w - \tilde{w}) \rangle. \end{aligned}$$

Since P is row-stochastic, we have $P\tilde{w} = \tilde{w}$ and thus $\langle \tilde{w}, P\tilde{w} \rangle = \|\tilde{w}\|^2$. Moreover, we have

$$\langle w - \tilde{w}, P\tilde{w} \rangle = \langle w - \tilde{w}, \tilde{w} \rangle = 0$$

since $w - \tilde{w}$ has mean value 0. We also observe, again because $w - \tilde{w}$ has mean value 0, that

$$\langle \tilde{w}, P(w - \tilde{w}) \rangle = \langle \tilde{w}, P_S(w - \tilde{w}) \rangle.$$

Collecting all these estimates, we obtain

$$\frac{\langle w, Pw \rangle}{\|w\|^2} \leq \frac{\|\tilde{w}\|^2}{\|w\|^2} + \frac{\|\tilde{w}\|}{\|w\|} \frac{\|w - \tilde{w}\|}{\|w\|} \|P_S\| + \frac{\|w - \tilde{w}\|^2}{\|w\|^2} \|P_S\|.$$

We also recall the Pythagorean theorem,

$$\|\tilde{w}\|^2 + \|w - \tilde{w}\|^2 = \|w\|^2.$$

Abbreviating $x = \|\tilde{w}\|/w$, we can abbreviate our upper bound as

$$\frac{\langle w, Pw \rangle}{\|w\|^2} \leq x^2 + x\sqrt{1-x^2} \|P_S\| + (1-x^2) \|P_S\|.$$

The function,

$$x \rightarrow x\sqrt{1-x^2} + (1-x^2)$$

is monotonically increasing on $[0, 1/3]$. Thus, assuming that

$$x = \frac{\|\tilde{w}\|}{\|w\|} \leq \frac{1 - \|P_S\|}{3},$$

we get, after some elementary computation,

$$\begin{aligned} \frac{\langle w, Pw \rangle}{\|w\|^2} &\leq \left(\frac{1 - \|P_S\|}{9} \right)^2 + \frac{1 - \|P_S\|}{9} \sqrt{1 - \left(\frac{1 - \|P_S\|}{9} \right)^2} \|P_S\| \\ &+ \left(1 - \left(\frac{1 - \|P_S\|}{9} \right)^2 \right) \|P_S\| \leq 0.2 + 0.8 \|P_S\|. \end{aligned}$$

However, we also recall from the proof of Theorem 3.2 that

$$\sum_{i=1}^n -\log \left(\sum_{j=1}^n e^{w_i v_j} \right) \leq -\|w\|^2.$$

Altogether, since the energy in the maximum has to exceed the energy in the origin, we have

$$\begin{aligned} -n \log n &\leq \langle w, Pw \rangle - \sum_{i=1}^n \log \left(\sum_{j=1}^n e^{w_i v_j} \right) \leq (0.2 + 0.8 \|P_S\|) \|w\|^2 \\ &- \|w\|^2 \end{aligned}$$

and therefore,

$$\|w\|^2 \leq \frac{2n \log n}{1 - \|P_S\|}.$$

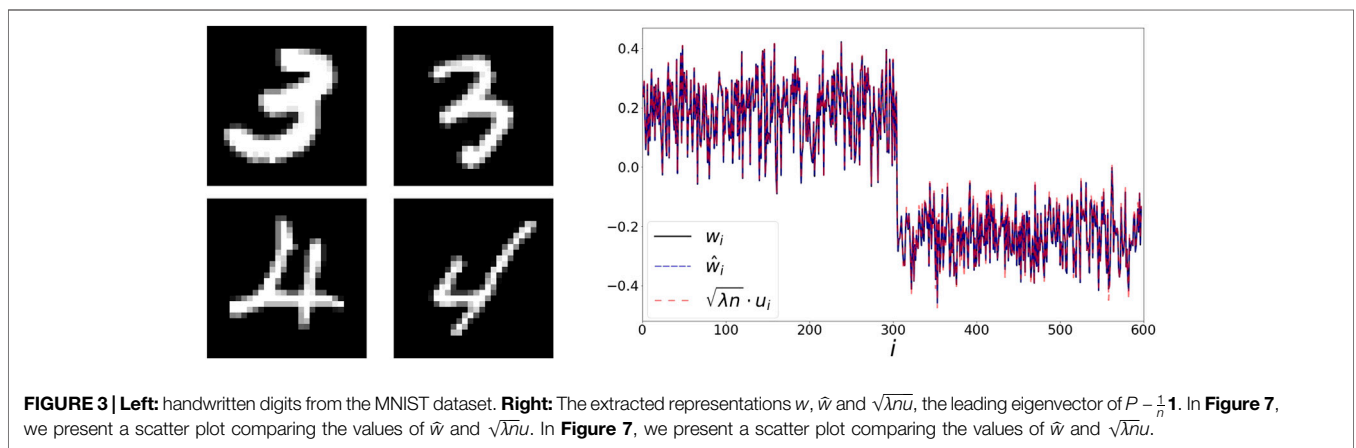
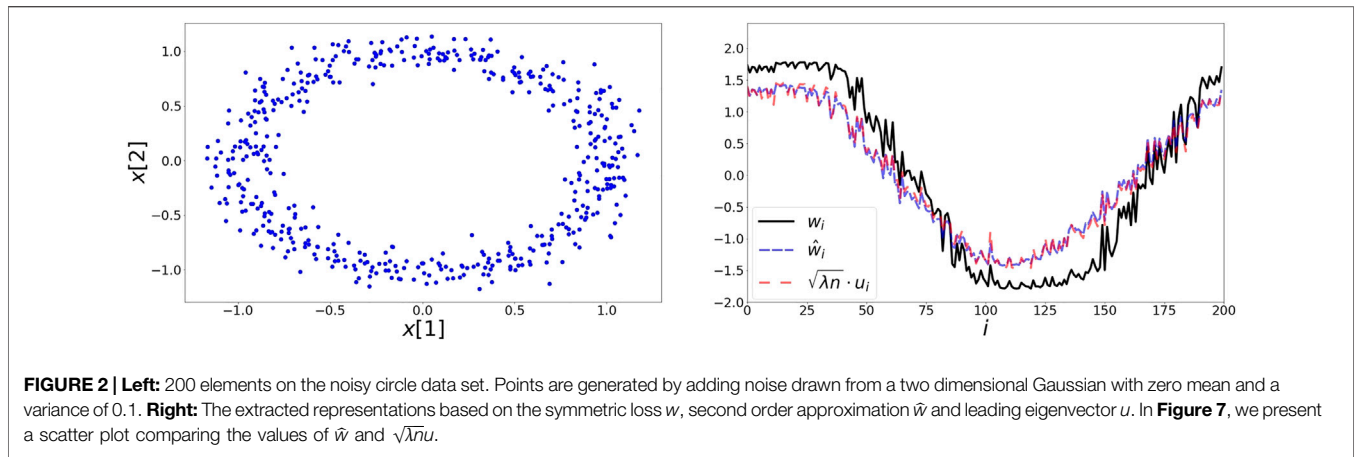
5 EXAMPLES

We validate our theoretical findings by comparing, for various datasets, the representation obtained by the following methods: i) optimizing over the symmetric functional in **1**, ii) optimizing over the spectral method suggested by Theorem 3.1 and iii) computing the leading eigenvector of $P - \frac{1}{n} \mathbf{1}$. We denote by w , \tilde{w} and u be the three vectors obtained by (i)–(iii), respectively. The comparison is performed for two artificial datasets, two sets of images, a seismic dataset and a text corpus. For the artificial, image and seismic data, the matrix P is obtained by the following steps: we compute a pairwise kernel matrix

$$K(x_i, x_j) = \exp \left(- \frac{\|x_i - x_j\|^2}{\alpha} \right),$$

where α is a scale parameter set as in [21] using a max-min scale. The max-min scale is set to

$$\alpha = \max_j \left[\min_{i,i \neq j} (\|x_i - x_j\|^2) \right], i, j = 1, \dots, n. \tag{4}$$



This global scale guarantees that each point is connected to at least one other point. Alternatively, adaptive scales could be used as suggested in [22]. We then compute P via

$$P_{ij} = K_{ij} / \sum_{l=1}^N K_{il}$$

The matrix P can be interpreted as a random walk over the data points, (see for example [18]). Theorem 3.1 holds for any matrix P that is similar to a symmetric matrix, here we use the common construction from [18], but our results hold for other variations as well. To support our approximation in Theorem 3.1, **Figure 7** shows a scatter plot of the scaled vector u vs. \hat{w} . In addition, we compute the correlation coefficient between u and \hat{w} by

$$\frac{\rho(u, \hat{w}) = (u - \mu)^T (\hat{w} - \hat{\mu})}{\|u - \mu\| \|\hat{w} - \hat{\mu}\|},$$

where μ and $\hat{\mu}$ are the means of u and \hat{w} respectively. A similar measure is done for w and u . In addition, we illustrate that the norm $\|w\|$ is comparable to \sqrt{n} , which supports the upper bound in Theorem 3.3.

5.1 Noisy Circle

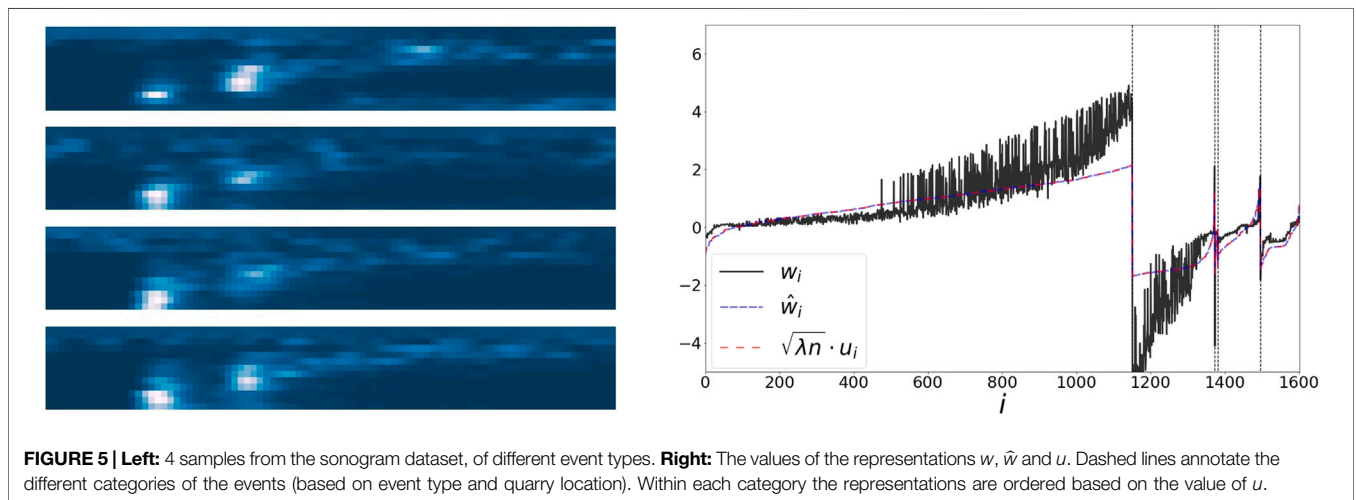
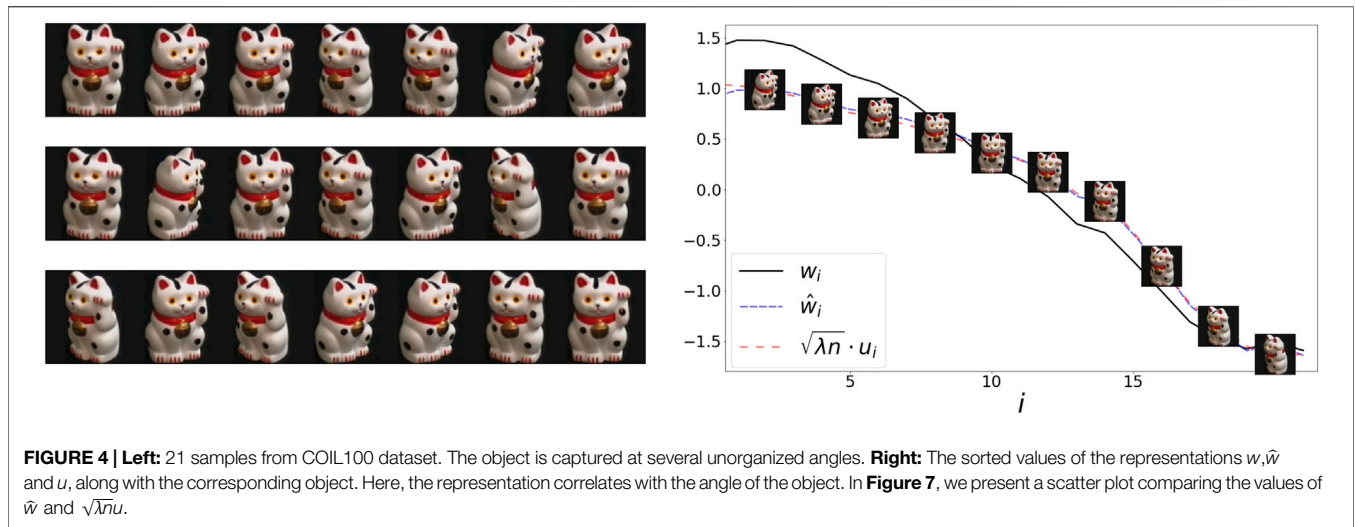
Here, the elements $\{x_1, \dots, x_{200} \in \mathbb{R}^2\}$ are generated by adding Gaussian noise with mean 0 and $\sigma^2 = 0.1$ to a unit circle (see the left panel of **Figure 2**). The right panel shows the extracted representations w , \hat{w} along with the leading eigenvector u scaled by $\sqrt{\lambda n}$ where λ is the corresponding eigenvalue. The correlation coefficients $\rho(w, u)$ and $\rho(\hat{w}, u)$ are equal to 0.98, 0.99 respectively.

5.2 Binary MNIST

Next, we use a set of 300 images of the digits 3 and 4 from the MNIST dataset [23]. Two examples from each category are presented in the left panel of **Figure 3**. Here, the extracted representations w and \hat{w} match the values of the scaled eigenvector u (see right panel of **Figure 3**). The correlation coefficients $\rho(w, u)$ and $\rho(\hat{w}, u)$ are both higher than 0.999.

5.3 COIL100

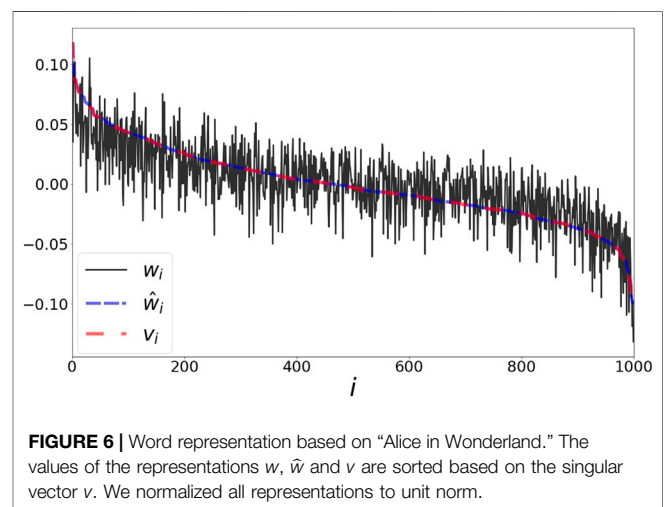
In this example, we use images from Columbia Object Image Library (COIL100) [24]. Our dataset contains 21 images of a cat captured at several pose intervals of 5 degrees (see left panel of **Figure 4**). We extract the embedding w and \hat{w} and reorder them based on the true angle of the cat at every image. In the right



panel, we present the values of the reordered representations w, \hat{w} and u overlaid with the corresponding objects. The values of all representations are strongly correlated with the angle of the object. Moreover, the correlation coefficients $\rho(w, u)$ and $\rho(\hat{w}, u)$, are 0.97 and 0.99 respectively.

5.4 Seismic Data

Seismic recordings could be useful for identifying properties of geophysical events. We use a dataset collected in Israel and Jordan, described in [25]. The data consists of 1632 seismic recordings of earthquakes and explosions from quarries. Each recording is described by a sonogram with 13 frequency bins, and 89 time bins [26] (see the left panel of **Figure 5**). Events could be categorized into 5 groups using manual annotations of their origin. We flatten each sonogram into a vector, and extract embeddings w, \hat{w} , and u . In the right panel of this figure, we show the extracted representations of all events. We use dashed



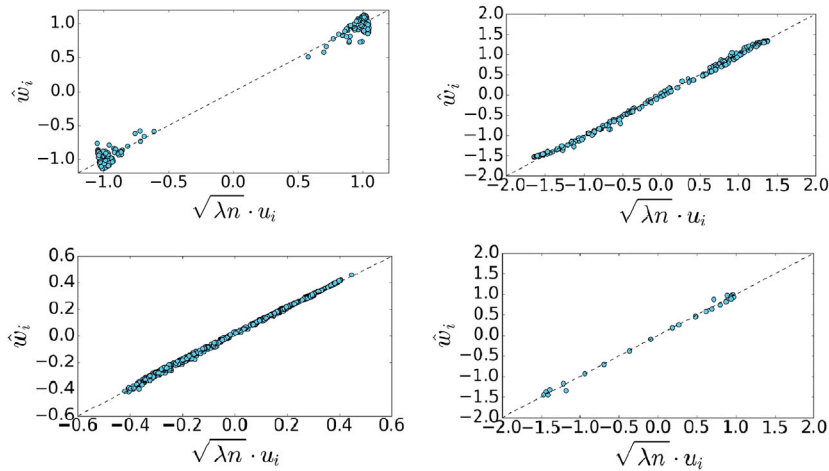


FIGURE 7 | Scatter plots of the scaled eigenvector of $P - \frac{1}{n}\mathbf{1}$, denoted by u , and the minimizer of the approximated functional (in Eq. 3), denoted by \hat{w} . From top left to bottom right, results based on data from: two Gaussian clusters, a circle, binary MNIST images and COIL100.

lines to annotate the different categories and sort the values within each category based on u . The coefficient $\rho(w, v)$ is equal to 0.89, and $\rho(\hat{w}, v) = 1$.

5.5 Text Data

As a final evaluation we use a corpus of words from the book “Alice in Wonderland” as processed in [27]. To define a co-occurrence matrix, we scan the sentences using a window size covering 5 neighbors before and after each word. We subsample the top 1000 words in terms of occurrences in the book. The matrix P is then defined by normalizing the co-occurrence matrix. In Figure 6 we present centered and normalized versions of the representations w , \hat{w} and the leading left singular vector v of $P - \frac{1}{n}\mathbf{1}$. The coefficient $\rho(w, v)$ is equal to 0.77, and $\rho(\hat{w}, v) = 1$.

governed by a regularized spectral method. For simplicity, we have restricted our analysis to a one dimensional symmetric representations, i.e. $w = v \in R^N$. Here, we demonstrate empirically that this result holds even when embedding n elements x_1, \dots, x_n in higher dimensions.

Let $w_i \in R^d$ be the embedding vector associated with x_i , where $d \ll n$ is the embedding dimension. The symmetric word2vec functional is given by

$$L(W) = \text{Tr}(W^T P W) - \sum_{i=1}^n \log \left(\sum_{j=1}^n \exp(w_i^T w_j) \right), \quad (5)$$

where $W = [w_1, \dots, w_n]^T \in R^{n \times d}$. A similar derivation to the one presented in Theorem 3.1 (the one dimensional case) yields the following approximation of 5,

$$L_2(W) = \text{Tr} \left(W^T \left(P - \frac{1}{n} \mathbf{1} \right) W \right) - \frac{\|W^T W\|_F^2}{2n} - n \log n. \quad (6)$$

6 MULTI-DIMENSIONAL EMBEDDING

In Section 3 we have demonstrated that under certain assumptions, the maximizer of the energy functional in 2 is

Note that both the symmetric functional in 5 and its approximation in 6 are invariant to multiplying W with an


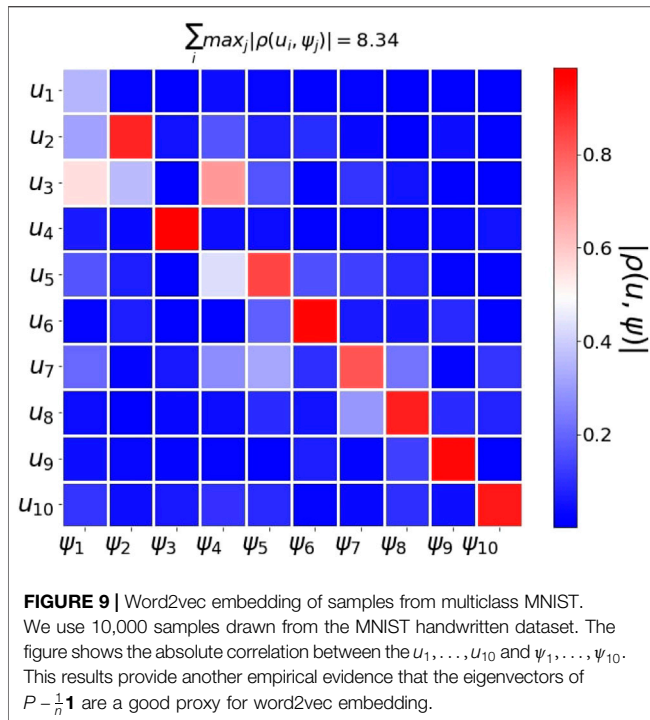


FIGURE 8 | Word2vec embedding of a collection of Gaussians. We use 2,500 points generated according to five distinct Gaussians $\mathcal{N}(r \cdot i \cdot \mathbf{1}, 2 \cdot I)$, where $\mathbf{1}$ is a 10–dimensional all ones vector, and r is scalar that controls the separation between the Gaussian centers. The figure shows the absolute correlation between u_1, \dots, u_4 and ψ_1, \dots, ψ_4 for $r = 8, 9$, and 10. For each value of r , we compute the sum of diagonal elements in the correlation matrix. This results demonstrate the strong agreement between word2vec embedding and the spectral representation based on $P - \frac{1}{n}\mathbf{1}$.



orthogonal matrix. That is, W and WR produce the same value in both functionals, where $R \in O(d)$.

To understand how the maximizer of 5 is governed by a spectral approach, we perform the following comparison. i) We obtain the optimizer W of 5 via gradient descent, and compute its left singular vectors, denoted u_1, \dots, u_d . ii) We compute the right singular vectors of $P - \frac{1}{n}\mathbf{1}$, denoted by ψ_1, \dots, ψ_d . iii) Compute the pairwise absolute correlation values $\rho(u_i, \psi_j)$.

We experiment on two datasets: 1) A collection of 5 Gaussians, and 2) images of hand written digits from MNIST. The transition probability matrix P was constructed as described in Section 5.

6.1 Data from Distinct Gaussians

In this experiment we generate a total of 2,500 samples, consisting of five sets of size 500. The samples in the i -th set are drawn independently according to a Gaussian distribution $\mathcal{N}(r \cdot i \cdot \mathbf{1}, 2 \cdot I)$, where $\mathbf{1}$ is a 10-dimensional all ones vector, and r is scalar that controls the separation between the Gaussian centers.

Figure 8 shows the absolute correlation value of the pairwise correlation between u_1, \dots, u_4 and ψ_1, \dots, ψ_4 for $r = 8, 9$, and 10. The correlation between the result obtained via the word2vec functional 5 and the right singular vectors of $P - \frac{1}{n}\mathbf{1}$ increase when the separation between the Gaussians is high.

REFERENCES

- Mikolov T, Chen K, Corrado G, Dean J. Efficient estimation of word representations in vector space. In: Bengio Y and LeCun Y editors. 1st International conference on learning representations, ICLR 2013; 2013 May 2–4; Scottsdale, Arizona, USA. Workshop Track Proceedings (2013).

6.2 Multi-Class MNIST

The data consists of 10,000 samples from the MNIST hand-written dataset with 1,000 images from each digit (0–9). We compute a 10-dimensional word2vec embedding W by optimizing (5).

Figure 9 shows the absolute correlation between the u_1, \dots, u_{10} and ψ_1, \dots, ψ_{10} . As evident from the correlation matrix, the results obtained by both methods span similar subspaces.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/Supplementary Material, further inquiries can be directed to the corresponding author.

AUTHOR CONTRIBUTIONS

AJ, EP, OL, SS, YK: designed research; EP, OL, AJ, SS, JP: performed research; SS, EP, AJ, OL contributed new reagents or analytic tools; EP, JP, OL: analyzed data; SS, EP, AJ, OL wrote the paper.

FUNDING

YK is supported in part by NIH grants UM1DA051410, R01GM131642, P50CA121974 and R61DA047037. SS was funded by NSF-DMS 1763179 and the Alfred P. Sloan Foundation. EP has been partially supported by the Blavatnik Interdisciplinary Research Center (ICRC), the Federmann Research Center (Hebrew University) Israeli Science Foundation research grant no. 1523/16, and by the DARPA PAI program (Agreement No. HR00111890032, Dr. T. Senator).

ACKNOWLEDGMENTS

The authors thank James Garritano and the anonymous reviewers for their helpful feedback.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fams.2020.593406/full#supplementary-material>

- Goldberg Y, Levy O. word2vec explained: deriving Mikolov et al.'s negative-sampling word-embedding method. arXiv preprint arXiv:1402.3722 (2014).
- Grover A Leskovec J. node2vec: scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining, San Francisco, CA, (2016) p. 855–64. doi:10.1145/2939672.2939754.

4. Mikolov T, Sutskever I, Chen K, Corrado GS, and Dean J. Distributed representations of words and phrases and their compositionality. *Adv Neural Inf Process Syst* (2013) 26:3111–9.
5. Le Q, Mikolov T. Distributed representations of sentences and documents. In: International conference on machine learning, Beijing, China, (2014): 1188–96.
6. Narayanan A, Chandramohan M, Venkatesan R, Chen L, Liu Y, and Jaiswal S. graph2vec: Learning distributed representations of graphs. arXiv preprint arXiv:1707.05005 (2017).
7. Hashimoto TB, Alvarez-Melis D, Jaakkola TS. Word embeddings as metric recovery in semantic spaces. *TACL* (2016) 4:273–86. doi:10.1162/tacl_a_00098.
8. Hinton GE Roweis ST. Stochastic neighbor embedding. *Adv Neural Inf Process Syst* (2003) 15:857–64.
9. Cotterell R, Poliak A, Van Durme B, and Eisner J. Explaining and generalizing skip-gram through exponential family principal component analysis. In: Proceedings of the 15th conference of the European chapter of the association for computational linguistics, (Valencia, Spain: Association for Computational Linguistics), Vol. 2 (2017). p. 175–81. Short Papers.
10. Collins M, Dasgupta S, and Schapire RE. A generalization of principal components analysis to the exponential family. *Adv Neural Inf Process Syst* (2002):617–24.
11. Levy O and Goldberg Y. Neural word embedding as implicit matrix factorization. *Adv Neural Inf Process Syst* (2014) 3:2177–85.
12. Qiu J, Dong Y, Ma H, Li J, Wang K, Tang J. “Network embedding as matrix factorization: unifying deepwalk, line, pte, and node2vec,” in Proceedings of the eleventh ACM international conference on web search and data mining, Marina Del Rey, CA (New York, NY: Association for Computing Machinery), (2018):459–67.
13. Perozzi B, Al-Rfou R, and Skiena S. “Deepwalk: online learning of social representations,” in Proceedings of the 20th ACM SIGKDD international conference on knowledge discovery and data mining, (New York, NY: Association for Computing Machinery), (2014):701–10.
14. Tang J, Qu M, Wang M, Zhang M, Yan J, and Mei Q. “Line: Large-scale information network embedding,” in Proceedings of the 24th international conference on world wide web, Florence, Italy (New York, NY: Association for Computing Machinery), (2015):1067–77.
15. Arora S, Li Y, Liang Y, Ma T, and Risteski A. Random walks on context spaces: towards an explanation of the mysteries of semantic word embeddings. arXiv abs/1502.03520 (2015).
16. Landgraf AJ and Bellay J. Word2vec skip-gram with negative sampling is a weighted logistic pca. arXiv preprint arXiv:1705.09755 (2017).
17. Belkin M. and Niyogi P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput* (2003). 15:1373–96. doi:10.1162/089976603321780317.
18. Coifman R. R. and Lafon S. Diffusion maps. *Appl Comput Harmon Anal* (2006) 21:5–30. doi:10.1016/j.acha.2006.04.006.
19. Singer A and Wu HT. Spectral convergence of the connection Laplacian from random samples. *Information Inference J IMA* (2017) 6:58–123. doi:10.1093/imaiai/iaw016.
20. Belkin M and Niyogi P. Convergence of Laplacian eigenmaps. *Adv Neural Inf Process Syst* (2007) 19:129–36.
21. Lafon S, Keller Y, and Coifman RR. Data fusion and multicue data matching by diffusion maps. *IEEE Trans Pattern Anal Mach Intell* (2006) 28:1784–97. doi:10.1109/tpami.2006.223.
22. Lindenbaum O, Salhov M, Yeredor A and Averbuch A. Gaussian bandwidth selection for manifold learning and classification. *Data Min Knowl Discov* (2020) 1–37. doi:10.1007/s10618-020-00692-x.
23. LeCun Y, Bottou L, Bengio Y, and Haffner P. Gradient-based learning applied to document recognition. *Proc IEEE* (1998) 86:2278–324. doi:10.1109/5.726791.
24. Nene SA, Nayar SK, and Murase H. Columbia object image library (coil-20). Technical Report CUCS006-96, Columbia University, Available at: https://www1.cs.columbia.edu/CAVE/publications/pdfs/Nene_TR96.pdf (1996).
25. Lindenbaum O, Bregman Y, Rabin N, and Averbuch A. Multiview kernels for low-dimensional modeling of seismic events. *IEEE Trans Geosci Rem Sens* (2018) 56:3300–10. doi:10.1109/tgrs.2018.2797537.
26. Joswig M. Pattern recognition for earthquake detection. *Bull Seismol Soc Am.* (1990). 80:170–86.
27. Johnsen P. A text version of “alice’s adventures in wonderland [Dataset]”. Availavle at: <https://gist.github.com/phillipj/4944029> (2019). Accessed August 8, 2020.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Jaffe, Kluger, Lindenbaum, Patsenker, Peterfreund and Steinerberger. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.