

WEARING A MASK: COMPRESSED REPRESENTATIONS OF VARIABLE-LENGTH SEQUENCES USING RECURRENT NEURAL TANGENT KERNELS

Sina Alemohammad, Hossein Babaei, Randall Balestriero, Matt Y. Cheung, Ahmed Imtiaz Humayun
Daniel LeJeune, Naiming Liu, Lorenzo Luzi, Jasper Tan, Zichao Wang, Richard G. Baraniuk

Rice University

ABSTRACT

High dimensionality poses many challenges to the use of data, from visualization and interpretation, to prediction and storage for historical preservation. Techniques abound to reduce the dimensionality of fixed-length sequences, yet these methods rarely generalize to variable-length sequences. To address this gap, we extend existing methods that rely on the use of kernels to variable-length sequences via use of the Recurrent Neural Tangent Kernel (RNTK). Since a deep neural network with ReLU activation is a Max-Affine Spline Operator (MASO), we dub our approach Max-Affine Spline Kernel (MASK). We demonstrate how MASK can be used to extend principal components analysis (PCA) and t-distributed stochastic neighbor embedding (t-SNE) and apply these new algorithms to separate synthetic time series data sampled from second-order differential equations.

Index Terms— Dimensionality reduction, Variable-length, Neural tangent kernel, t-SNE, Recurrent neural network

1. INTRODUCTION

Many signal processing and machine learning problems involve high-dimensional data. In order to analyze and exploit such data effectively, we must find meaningful low-dimensional representations. On the one hand, it is often extremely valuable to visualize data as a part of the exploration process, yet we are limited as humans to visualizing in only a few dimensions, typically two or three. On the other hand, the data may be nonlinearly embedded in a high-dimensional space, yet the meaningful attributes for a prediction task of interest may lie in a low-dimensional space where linear predictors achieve high accuracy.

Consequently, myriad methods for dimensionality reduction have been developed. These include linear parametric transformations, such as Principal Components Analysis (PCA) [1], sparse PCA [2], and Independent Components Analysis (ICA) [3]; nonlinear parametric methods, such as bottleneck features [4]; and non-parametric methods, such as multi-dimensional scaling (MDS) [5], Isomap [6] and t-distributed stochastic neighbor embedding (t-SNE), all of

which transform data points from the input space to a low-dimensional space. Many of these methods find embeddings for collections of data points in new spaces that preserve similarities or distances. For example, in classical MDS, data points are embedded into a space of lower dimensions to preserve pairwise Euclidean distance. Such a distance, while simple and efficient, yields a linear embedding that is incapable of preserving nonlinear structures in data, corroborating research for more representative distance measures that replace Euclidean distance. Isomap preserves the geodesic distance of the points [6], and t-SNE preserves distances based on fitted probability distributions [7]. Some works take these a step further and use pairwise distances based on kernel functions, a technique commonly used in machine learning. Such kernel functions have successfully been applied to PCA [8], Isomap [9, 10] and t-SNE [11].

Unfortunately, a challenge to all these methods is that the pairwise distance function requires the two data points under comparison to have the same number of dimensions, restricting these methods to fixed-length data. This is particularly troubling for time-series data such as natural language or physics continuity equations [12, 13, 14], wherein one may need to deal with sequences of different lengths. One way to solve this is to use a pairwise distance function that can take inputs of different lengths.

Recent advances in deep learning theory have revealed that the training dynamics of infinite-width neural networks (NNs) learned via first-order gradient descent with Gaussian parameter initialization are captured by a so-called *neural tangent kernel* (NTK) [15]. In particular, [16, 17] showed that a kernel can be derived from an infinitely wide recurrent neural networks called the recurrent neural tangent kernel (RNTK). The RNTK inherits the universal representation power of infinitely wide RNNs [18] while also representing a practical kernel that has been shown to outperform both kernels induced from non-recurrent neural architectures and classical kernels on a number of time series tasks [16]. RNTK also gracefully handles data-points of varying lengths. The above results and properties of RNTK make it a promising yet under-explored candidate for time-series dimensionality reduction applications.

1.1. Contributions

We propose *MASK* (Max-Affine Spline Kernel), a kernel-based time series dimensionality reduction method based on RNTK with ReLU activation (and thus can be regarded as a max-affine spline operator (MASO); see [19, 20] for detailed explanations). We show that by simply replacing the commonly used kernels (such as Gaussian and polynomial) with the RNTK in classic dimensionality reduction algorithms such as PCA and t-SNE, we can enable these algorithms (which are traditionally not well suited for time series data) to work well with time series data. We experimentally validate our proposed MASK with PCA (RNTK PCA) and t-SNE (RNTK t-SNE) on various synthetic time series data generated by linear dynamical systems and demonstrate MASK's superior performance in discovering clear patterns in each dataset in the reduced dimensional space compared to PCA and t-SNE with classic kernels.

2. RNTK AS A SIMILARITY MEASURE

In this section, we introduce RNTK and briefly show how to compute it given a set of time series sequences of potentially varying lengths. Note that this kernel matrix can be treated as a similarity measure which then leads to its application to dimensionality reduction that we explain in more detail in the next section.

Consider a Deep Neural Network (DNN) mapping $f_\theta(\mathbf{x})$ with parameters θ and a pair of inputs \mathbf{x} and \mathbf{x}' , the *Neural Tangent Kernel* (NTK) [15] is defined as:

$$\Theta(\mathbf{x}, \mathbf{x}') := \langle \nabla_\theta f_\theta(\mathbf{x}), \nabla_\theta f_\theta(\mathbf{x}') \rangle. \quad (1)$$

where θ is Gaussian distributed. In the regime of infinite width of the layers employed by f_θ , the above converges to an analytical form [21, 22, 23, 24, 25] which can be used to compute input similarities or gain insights into DNNs training and representation power [26, 27, 28, 29]. When the consider DNN as a recurrent architecture, namely a Recurrent Neural Network, the limiting kernel is denoted as the RNTK [16, 17]. The key property that we will exploit in this paper is the fact that RNTK provides a novel way to obtain similarity measure of different length inputs thanks to the underlying recurrent modeling of the inputs.

We overview the derivation of the RNTK from a single-layer RNN with average pooling [30]. Given an input data sequence $\mathbf{x} = \{\mathbf{x}(t)\}_{t=1}^T$ of fixed length T with $\mathbf{x}^{(t)} \in \mathbb{R}^m$, an RNN with n units in each hidden layer performs the following recursive computation at each layer and each time step

$$\mathbf{g}^{(t)}(\mathbf{x}) = \frac{\sigma_w}{\sqrt{n}} \mathbf{W} \mathbf{h}^{(t-1)}(\mathbf{x}) + \frac{\sigma_u}{\sqrt{n}} \mathbf{U} \mathbf{x}(t) + \sigma_b \mathbf{b} \quad (2)$$

$$\mathbf{h}^{(t)}(\mathbf{x}) = \phi(\mathbf{g}^{(t)}(\mathbf{x})) \quad (3)$$

where $\mathbf{W} \in \mathbb{R}^{n \times n}$, $\mathbf{b} \in \mathbb{R}^n$, $\mathbf{U} \in \mathbb{R}^{n \times m}$, and $\phi(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ is the activation function that acts entry-wise on a vector.

$\mathbf{h}^{(0)}(\mathbf{x})$ is the initial hidden state, which is set to zero. The output of an RNN at time t is computed via

$$f_\theta^{(t)}(\mathbf{x}) = \frac{1}{\sqrt{n}} \mathbf{V} \mathbf{h}^{(t)}(\mathbf{x}) \quad (4)$$

and the final output of the RNN is computed via

$$f_\theta(\mathbf{x}) = \sum_{t=1}^T f_\theta^{(t)}(\mathbf{x}) = \sum_{t=1}^T \frac{1}{\sqrt{n}} \mathbf{V} \mathbf{h}^{(t)}(\mathbf{x}) \in \mathbb{R}^d. \quad (5)$$

The learnable parameters $\theta := \text{vect}[\mathbf{W}, \mathbf{U}, \mathbf{b}, \mathbf{V}]$ are initialized to Gaussian $\mathcal{N}(0, 1)$ random variables. The set $\mathcal{S} = \{\sigma_w, \sigma_u, \sigma_b\}$ represents the initialization hyperparameters.

Leveraging the equivalence of infinite-width NNs with Gaussian processes (GPs) [21, 22, 23, 24, 25], as $n \rightarrow \infty$, each coordinate of the RNN output at each time step converges to a Gaussian process for two input sequence of possible different length $\mathbf{x} = \{\mathbf{x}(t)\}_{t=1}^T$ and $\mathbf{x}' = \{\mathbf{x}'(t)\}_{t=1}^{T'}$ with kernel

$$\mathcal{K}^{(t,t')}(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\theta \sim \mathcal{N}} [[f_\theta^{(t)}(\mathbf{x})]_i \cdot [f_\theta^{(t')}(\mathbf{x}')]_i], \forall i \in [d] \quad (6)$$

known as the NN-GP kernel. The pre-activation $\mathbf{g}^{(t)}(\mathbf{x})$ and gradient vectors $\boldsymbol{\delta}^{(t)}(\mathbf{x}) := \sqrt{n}(\nabla_{\mathbf{g}^{(t)}(\mathbf{x})} f_\theta(\mathbf{x}))$ also converge to zero mean GPs with kernels

$$\Sigma^{(t,t')}(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\theta \sim \mathcal{N}} [[\mathbf{g}^{(t)}(\mathbf{x})]_i \cdot [\mathbf{g}^{(t')}(\mathbf{x}')]_i] \quad \forall i \in [n] \quad (7)$$

$$\Pi^{(t,t')}(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\theta \sim \mathcal{N}} [[\boldsymbol{\delta}^{(t)}(\mathbf{x})]_i \cdot [\boldsymbol{\delta}^{(t')}(\mathbf{x}')]_i] \quad \forall i \in [n]. \quad (8)$$

The RNTK for two inputs \mathbf{x} and \mathbf{x}' with length T and T' can thus be obtained as

$$\Theta(\mathbf{x}, \mathbf{x}') = \left(\sum_{t=1}^T \sum_{t'=1}^{T'} (\Pi^{(t,t')}(\mathbf{x}, \mathbf{x}') \cdot \Sigma^{(t,t')}(\mathbf{x}, \mathbf{x}')) + \mathcal{K}^{(t,t')}(\mathbf{x}, \mathbf{x}') \right) \otimes \mathbf{I}_d. \quad (9)$$

where \mathbf{I}_d is the identity matrix of size d . For $\phi(z) = \text{ReLU}(z) = \max(0, z)$, an analytical formula exists for the RNTK, enabling fast point-wise evaluation of RNTK on the input data. The full and detailed derivation of the RNTK is given in [17]. Given a means to compute the RNTK, we can use the RNTK to compute similarities between data points of varying length and extend classical dimensionality reduction techniques.

3. MASK: RNTK-BASED DIMENSIONALITY REDUCTION

We now present MASK, the incorporation of the RNTK into standard dimensionality reduction methods. The key idea is to

replace the pairwise distance (or similarity) measure used in these methods with a new distance computed by RNTK. We introduce two specific instances of MASK: RNTK PCA and RNTK t-SNE. While we specifically use PCA and t-SNE as examples, we emphasize that MASK is generally applicable and that any dimensionality reduction technique that operates off of kernels is compatible with MASK.

3.1. RNTK PCA

PCA constructs new dimensions by finding the linear combination of the original dimensions that are uncorrelated with each additional dimension containing the most variance possible from the data [31]. It can be shown that these dimensions are the eigenvectors of the data's covariance matrix and are thus linear combinations of the original data dimensions, yielding a linear embedding. Many times, the data is not linearly separable in its original dimensions, and it may be beneficial to first map the data points into a higher-dimensional space before finding its covariance matrix's eigenvectors to allow nonlinear embeddings. Using the kernel trick, one need not compute the higher-dimensional data points; one simply needs to compute the inner products between the higher-dimensional data points (i.e. the output of the kernel function) in the form of a matrix K with entries $K_{i,j} = \langle \Phi(\mathbf{x}_i), \Phi(\mathbf{x}_j) \rangle$ for some given mapping function Φ and following the method in [32]. RNTK PCA is obtained by using the RNTK outputs as the kernel: $K_{i,j} = \Theta(\mathbf{x}_i, \mathbf{x}_j)$ from Eq. 1.

3.2. RNTK t-SNE

In t-SNE, the pairwise similarity of two data points \mathbf{x}_i and \mathbf{x}_j is defined as a conditional probability:

$$p_{j|i} = \frac{\exp\left(\frac{-\|\mathbf{x}_i, \mathbf{x}_j\|^2}{2\sigma_i^2}\right)}{\sum_{k \neq i} \exp\left(\frac{-\|\mathbf{x}_i, \mathbf{x}_k\|^2}{2\sigma_i^2}\right)}, \quad (10)$$

where σ_i is a parameter [7]. Next, a probability distribution is defined based on the distances of the low-dimensional embeddings:

$$q_{j|i} = \frac{\exp(-\|\mathbf{y}_i - \mathbf{y}_j\|^2)}{\sum_{k \neq i} \exp(-\|\mathbf{y}_i - \mathbf{y}_k\|^2)}. \quad (11)$$

The low-dimensional embeddings \mathbf{y}_i for each data point \mathbf{x}_i is then obtained by minimizing the Kullback-Leibler divergence $C = \sum_i \sum_j p_{j|i} \log \frac{p_{j|i}}{q_{j|i}}$ between the two induced probability distributions.

To apply the RNTK function into t-SNE, we first define the RNTK distance based on Eq. 1 and the Pearson dissimilarity [33]:

$$d(\mathbf{x}_i, \mathbf{x}_j) = \sqrt{1 - \frac{\Theta(\mathbf{x}_i, \mathbf{x}_j)}{\sqrt{\Theta(\mathbf{x}_i, \mathbf{x}_i)\Theta(\mathbf{x}_j, \mathbf{x}_j)}}}. \quad (12)$$

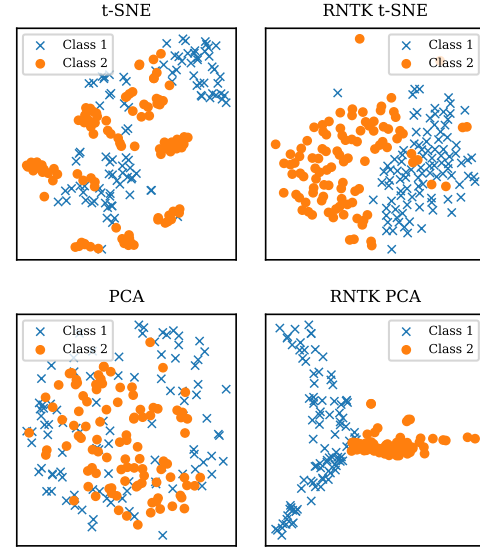


Fig. 1. Demonstration of MASK dimensionality reduction to separate variable-length time series data. *Top left:* t-SNE on original sequences padded with zeros fails to separate the data. *Top right:* RNTK t-SNE successfully separates the data into distinct clusters corresponding to the different classes of time series. *Bottom left:* PCA on original sequences padded with zeros fails to separate the data. *Bottom right:* RNTK PCA successfully separates the data into distinct clusters corresponding to the different classes of time series.

The RNTK t-SNE pairwise conditional probabilities for two data points \mathbf{x}_i and \mathbf{x}_j can then be defined as:

$$p_{j|i} = \frac{\exp\left(\frac{-d(\mathbf{x}_i, \mathbf{x}_j)^2}{2\sigma_i^2}\right)}{\sum_{k \neq i} \exp\left(\frac{-d(\mathbf{x}_i, \mathbf{x}_k)^2}{2\sigma_i^2}\right)}, \quad (13)$$

which differs from classical t-SNE simply by replacing its Euclidean distance with the novel RNTK-based distance. Afterwards, the standard t-SNE method of finding the Kullback-Leibler divergence between p and q proceeds as usual to produce the low-dimensional MASK embeddings.

4. EXPERIMENTS

We now experimentally demonstrate that two specific examples of MASK, namely RNTK PCA and RNTK t-SNE, produce more distinct patterns on several synthetic datasets compared to PCA and t-SNE with classic kernels.¹

4.1. Dataset

Many real-world time series data can be modeled by linear dynamic systems. Therefore, for our experiments, we use

¹Our code can be found at <https://github.com/dlej/MASK>.

synthetic time series data generated from the following family of differential equations:

$$t^2 \frac{\partial^2 x(t)}{\partial t^2} - a_1 \frac{\partial x(t)}{\partial t} - a_0 t x(t) = 0. \quad (14)$$

To sample time series data of different classes, we draw $a_0 \sim \mathcal{N}(0, 0.0025)$ and $a_1 \sim \mathcal{N}(-4, 4)$ for class 1 and $a_0 \sim \mathcal{N}(0, 0.01)$ and $a_1 \sim \mathcal{N}(-20, 1)$ and for class 2. We sample 100 sequences of length 30 for each class and then randomly drop elements with probability 0.2 in each sequence to shorten the sequences, resulting in sequences with varying lengths.

We emphasize that we choose the system in Eq. 14 only for demonstrative purpose and that MASK works generally on other types of time series data.

4.2. Experiment Setup

We compute RNTK using the default set of parameters: $\sigma_w = 2$, $\sigma_u = 0.316$, and $\sigma_b = 10^{-5}$. To compare against fixed-length methods, i.e., PCA and t-SNE with classic kernels, we pad the end of the sequences with zeros until they are of the same length. We used the Scikit-learn [34] implementations of PCA and t-SNE with default parameters, although the results were similar even when tuning para

4.3. Results

Fig. 1 visualizes the dimensionality reduction results comparing PCA and t-SNE with their RNTK replaced counterparts. The color and shapes of points represent different classes. We clearly observe obvious separation of classes in both RNTK PCA and RNTK t-SNE whereas classic PCA and t-SNE fails to separate the data. These results validate the proposed MASK's capability in extracting useful patterns from the synthetic data, demonstrating its potential as a practical kernel dimensionality reduction technique for time series data of varying lengths.

5. CONCLUSIONS

We have demonstrated that, through MASK, we can extend existing, well-proven dimensionality reduction techniques to variable-length data using the RNTK. However, dimensionality reduction is not the only method that can be extended in this manner; any kernel- or distance-based method can be extended using MASK, for example, clustering. Our promising preliminary experimental results on synthetic data validate the utility of MASK. Future avenues of research include evaluating MASK on other modalities of time series data, such as text, audio, speech, and natural language and extending the current technique to other kernel-based methods beyond PCA and t-SNE. MASK will revolutionize the way we think

about dimensionality reduction of different length data sequences and affords us with tremendous future directions for real-world data.

6. REFERENCES

- [1] Karl Pearson, "On lines and planes of closest fit to systems of points in space," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.
- [2] Hui Zou, Trevor Hastie, and Robert Tibshirani, "Sparse principal component analysis," *J. Comput. Graphical Statist.*, vol. 15, no. 2, pp. 265–286, 2006.
- [3] Pierre Comon, "Independent component analysis, a new concept?," *Signal Process.*, vol. 36, no. 3, pp. 287–314, 1994.
- [4] Frantisek Grézl, Martin Karafiát, Stanislav Kontár, and Jan Cernocky, "Probabilistic and bottle-neck features for lvcsr of meetings," in *IEEE Int. Conf. Acoust., Speech, Signal Process.* IEEE, 2007, vol. 4, pp. IV–757.
- [5] Michael AA Cox and Trevor F Cox, "Multidimensional scaling," in *Handbook of Data Visualization*, pp. 315–347. Springer, 2008.
- [6] Joshua B Tenenbaum, Vin De Silva, and John C Langford, "A global geometric framework for nonlinear dimensionality reduction," *science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [7] Laurens van der Maaten and Geoffrey Hinton, "Visualizing data using t-sne," *J. Mach. Learn. Res.*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [8] Sebastian Mika, Bernhard Schölkopf, Alex J Smola, Klaus-Robert Müller, Matthias Scholz, and Gunnar Rätsch, "Kernel pca and de-noising in feature spaces," in *Advances Neural Inf. Process. Syst.*, 1999, pp. 536–542.
- [9] Heeyoul Choi and Seungjin Choi, "Kernel isomap," *Electronics letters*, vol. 40, no. 25, pp. 1612–1613, 2004.
- [10] Heeyoul Choi and Seungjin Choi, "Robust kernel isomap," *Pattern Recognition*, vol. 40, no. 3, pp. 853–862, 2007.
- [11] Andrej Gisbrecht, Alexander Schulz, and Barbara Hammer, "Parametric nonlinear dimensionality reduction using kernel t-sne," *Neurocomputing*, vol. 147, pp. 71–82, 2015.

- [12] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Ng, and C. Potts, "Recursive deep models for semantic compositionality over a sentiment treebank," in *Proc. Conf. Empirical Methods Natural Language Process. (EMNLP)*, Oct. 2013, pp. 1631–1642.
- [13] A. Graves, A. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech and Signal Process. (ICASSP)*, May 2013, pp. 6645–6649.
- [14] E. Cakir, S. Adavanne, G. Parascandolo, K. Drossos, and T. Virtanen, "Convolutional recurrent neural networks for bird audio detection," in *2017 25th European Signal Processing Conference (EUSIPCO)*, 2017, pp. 1744–1748.
- [15] Arthur Jacot, Franck Gabriel, and Clément Hongler, "Neural tangent kernel: Convergence and generalization in neural networks," in *Advances in Neural Information Processing Systems*, 2018, pp. 8571–8580.
- [16] Sina Alemohammad, Zichao Wang, Randall Balestriero, and Richard Baraniuk, "The recurrent neural tangent kernel," *arXiv preprint arXiv:2006.10246*, 2020.
- [17] Greg Yang, "Tensor programs ii: Neural tangent kernel for any architecture," *arXiv preprint arXiv:2006.14548*, 2020.
- [18] Anton Maximilian Schäfer and Hans Georg Zimmermann, "Recurrent neural networks are universal approximators," in *Proc. 16th Int. Conf. Artif. Neural Netw. - Volume Part I*, 2006, p. 632–640.
- [19] Zichao Wang, Randall Balestriero, and Richard Baraniuk, "A max-affine spline perspective of recurrent neural networks," in *Int. Conf. Learn. Representations*, 2018.
- [20] Randall Balestriero and richard baraniuk, "A spline theory of deep learning," 10–15 Jul 2018, vol. 80 of *Proc. Mach. Learn. Res.*, pp. 374–383.
- [21] Jaehoon Lee, Jascha Sohl-dickstein, Jeffrey Pennington, Roman Novak, Sam Schoenholz, and Yasaman Bahri, "Deep neural networks as gaussian processes," in *International Conference on Learning Representations*, 2018.
- [22] David Duvenaud, Oren Rippel, Ryan Adams, and Zoubin Ghahramani, "Avoiding pathologies in very deep networks," in *Artificial Intelligence and Statistics*, 2014, pp. 202–210.
- [23] Roman Novak, Lechao Xiao, Yasaman Bahri, Jaehoon Lee, Greg Yang, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-dickstein, "Bayesian deep convolutional networks with many channels are gaussian processes," in *Int. Conf. Learn. Representations*, 2019.
- [24] Adrià Garriga-Alonso, Carl Edward Rasmussen, and Laurence Aitchison, "Deep convolutional networks as shallow gaussian processes," in *Int. Conf. Learn. Representations*, 2019.
- [25] Greg Yang, "Tensor programs i: Wide feedforward or recurrent neural networks of any architecture are gaussian processes," *arXiv preprint arXiv:1910.12478*, 2019.
- [26] Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai, "Gradient descent finds global minima of deep neural networks," in *Int. Conf. Mach. Learn.*, 2019, pp. 1675–1685.
- [27] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song, "On the convergence rate of training recurrent neural networks," in *Advances in Neural Information Processing Systems*, 2019, pp. 6676–6688.
- [28] Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang, "Learning and generalization in overparameterized neural networks, going beyond two layers," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2019, pp. 6155–6166.
- [29] Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu, "Stochastic gradient descent optimizes overparameterized deep relu networks," *arXiv preprint arXiv:1811.08888*, 2018.
- [30] Jeffrey L Elman, "Finding structure in time," *Cognitive science*, vol. 14, no. 2, pp. 179–211, 1990.
- [31] Hervé Abdi and Lynne J Williams, "Principal component analysis," *Wiley Interdisciplinary Reviews: Comput. Statist.*, vol. 2, no. 4, pp. 433–459, 2010.
- [32] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller, "Kernel principal component analysis," in *Int. Conf. Artificial Neural Netw.* Springer, 1997, pp. 583–588.
- [33] Victor Solo, "Pearson distance is not a distance," *arXiv preprint arXiv:1908.06029*, 2019.
- [34] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al., "Scikit-learn: Machine learning in python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.