

Lawrence Berkeley National Laboratory

LBL Publications

Title

Nonlinear Matrix Approximation with Radial Basis Function Components

Permalink

<https://escholarship.org/uc/item/0wt4g7hb>

Authors

Rebrova, Elizaveta
Tang, Yu-Hang

Publication Date

2021-06-03

Peer reviewed

NONLINEAR MATRIX APPROXIMATION WITH RADIAL BASIS FUNCTION COMPONENTS

ELIZAVETA REBROVA AND YU-HANG TANG
 COMPUTATIONAL RESEARCH DIVISION
 LAWRENCE BERKELEY NATIONAL LABORATORY
 BERKELEY, CA 94598, USA
 ELIZAVETA.REBROVA@GMAIL.COM, TANG.MAXIN@GMAIL.COM

ABSTRACT. We introduce and investigate matrix approximation by decomposition into a sum of radial basis function (RBF) components. An RBF component is a generalization of the outer product between a pair of vectors, where an RBF function replaces the scalar multiplication between individual vector elements. Even though the RBF functions are positive definite, the summation across components is not restricted to convex combinations and allows us to compute the decomposition for *any* real matrix that is not necessarily symmetric or positive definite. We formulate the problem of seeking such a decomposition as an optimization problem with a nonlinear and non-convex loss function. Several modern versions of the gradient descent method, including their scalable stochastic counterparts, are used to solve this problem. We provide extensive empirical evidence of the effectiveness of the RBF decomposition and that of the gradient-based fitting algorithm. While being conceptually motivated by singular value decomposition (SVD), our proposed nonlinear counterpart outperforms SVD by drastically reducing the memory required to approximate a data matrix with the same L_2 -error for a wide range of matrix types. For example, it leads to 2 to 6 times memory save for Gaussian noise, graph adjacency matrices, and kernel matrices. Moreover, this proximity-based decomposition can offer additional interpretability in applications that involve, e.g., capturing the inner low-dimensional structure of the data, retaining graph connectivity structure, and preserving the acutance of images.

1. INTRODUCTION

There is an extensive body of work that aims to approximate large multi-dimensional data via simple low-dimensional components. Such approximations save the storage memory and expose simpler underlying structures in the data. The singular value decomposition (SVD) for matrices is arguably the standard example of such a low-rank approximation, while the canonical polyadic (CP) decomposition [1, 2], where every component is represented by the outer product of two or more vectors, is the generalization for higher-order tensors. Much less is known for approximations that are low-complexity but not necessarily low-rank. In this work, we propose one such approximation by generalizing the outer product between vectors. By *generalizing*, we mean to replace the scalar product between vector elements by a nonlinear function, e.g. the one-dimensional radial basis function (RBF) in this case. Just like a component formed by an outer product, each nonlinear component is indexed by two vectors with potentially one additional scaling factor. However, the nonlinear relationship introduced between the vectors leads to more expressive components as manifested by, for example, their higher rank. This in turn yields better overall approximations with less explicit components comparing to the SVD in many situations.

The generalization of the outer product using RBF functions stems from their use in kernel machines [3]. In that context, a kernel is a positive definite function that evaluates the inner product between data points after implicitly mapping them to a higher-dimensional vector space. Specifically, the RBF kernel between two feature vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ is a real-valued function of the form

$$(1) \quad k_{\text{RBF}}(\mathbf{u}, \mathbf{v}) = \exp\left(-\frac{\|\mathbf{u} - \mathbf{v}\|^2}{2h^2}\right),$$

where h is a scalar parameter. The mapping, which is nonlinear, introduces flexibility for modeling functions and decision boundaries that would otherwise be difficult to learn in the original feature space. Moreover, the RBF kernel can be conveniently interpreted as a similarity measure, *i.e.* the value of the kernel is small if input vectors are far away from each other.

In our work, an *RBF component* created from two vectors $\mathbf{u} = (u_1, \dots, u_n) \in \mathbb{R}^n$, $\mathbf{v} = (v_1, \dots, v_m) \in \mathbb{R}^m$ is a kernel matrix $K = (k_{ij}) \in \mathbb{R}^{n \times m}$ with the elements defined as

$$(2) \quad k_{ij} = \exp\left(-|u_i - v_j|^2\right).$$

An important difference here, as opposed to the case in (1), is that the vectors \mathbf{u}, \mathbf{v} are not known inputs but rather *learnable* parameters for expressing a given target data matrix. In particular, there is no need to specify the length-scale parameter h as in the standard RBF definition (1), since optimizing \mathbf{u} and \mathbf{v} can achieve an equivalent effect. Similarly to the RBF function (1), an RBF component (2) exposes the underlying geometric structures in the data matrices based on the proximity between the elements of the learned vectors. However, we emphasize that the approximations are not specific to any special geometry-related features in the data. Linear combinations of such kernel matrices can act as universal approximators of general matrices while competing successfully in efficiency with SVD decomposition.

The task of finding and interpreting an optimal RBF approximation results in unique challenges distinct from SVD-based approximation algorithms. This includes the absence of known direct algorithms to find the best decomposition, the lack of orthogonality between the individual components, and the lack of a clear ordering of the components based on their norm. However, we show that trading-in the linearity of the components has great advantages in memory efficiency and additional expressivity of the decomposition due to the Euclidean distance in the structure of the RBF kernels. We also demonstrate that iterative gradient-based optimization methods make the problem of finding a good approximation tractable.

This paper is structured as follows. In Section 2, we present a synthetic case where an RBF approximation is naturally optimal because the matrix to be approximated is created from the difference between two RBF components. This sets the scene for a formal description of our approximation method, including a discussion of the success and sophistication of the learning process, as presented in Section 3. In Section 4, we provide extensive empirical evidence of the general applicability of the method to a variety of natural matrix models without *a priori* underlying structures that make them particularly amenable to the RBF approximation. In Section 5, we discuss how additional knowledge about the structure in data can be obtained using the learned RBF components in several situations, including a manifold learning example, a graph structure learning example, and an image compression example. In Section 6, we compare the proposed approach with related works. In Section 7, we conclude the paper with a discussion of multiple future directions related to the analysis and applications of the proposed RBF approximation.

2. A MOTIVATING EXAMPLE

Let us consider a special case when an RBF component defined in (2) is square and symmetric, namely, $\mathbf{u} = \mathbf{v}$ and

$$(3) \quad \tilde{k}_{ij} := \exp[-(u_i - u_j)^2], \quad \text{for each entry } \tilde{k}_{ij} \text{ in } \tilde{K} \in \mathbb{R}^{n \times n}.$$

Depending on the vector \mathbf{u} , a symmetric RBF component can have vastly different structures from a rank one constant matrix to a full-rank near identity matrix. In the case where \mathbf{u} is randomly sampled from a regular enough distribution, the spectrum of the normalized matrix $n^{-1}K$ is concentrated around the spectrum of its expectation $n^{-1}\mathbb{E}K$ [4]. For example, when the vector elements are *i.i.d.* standard Gaussians, with high probability, the RBF component has 9 eigenvalues above 10^{-5} .

A fact that we take as an essential advantage for approximating general matrices is that a linear combination of two or more RBF components is not necessarily semi-positive definite. Figure 1 shows one such example of a two-component matrix $K_{exact2} = 5K^{(1)} - 4K^{(2)}$, where $K^{(1)}$ and $K^{(2)}$ are the RBF components defined per (3). Without loss of generality, we take $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ to be standard normal vectors smoothed by one-dimensional Gaussian filters with standard deviations of 3 and 6, respectively. Even though the matrix can be perfectly reconstructed from only $2n$ parameters that encode $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$, the best rank-4 SVD approximation has a mean square error above 0.01 and delivers a visibly imperfect reconstruction using $4n$ parameters.

In this paper, we are motivated by the following *inverse* questions: Would this be possible to recover vectors $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ if we were just given a matrix K_{exact2} , similar to how the SVD reconstruction in Figure 1 was found? More generally, given a matrix K that is known to be equal to, or well-approximated by a linear combination of r general RBF components as per (2), can we learn the vectors $\mathbf{u}^{(k)}, \mathbf{v}^{(k)}$ and scalar coefficients that define this linear combination? Further, can we find a linear combination of the RBF components to approximate a generic




FIGURE 1. The best rank 1, 2 and 4 approximations of K_{exact2} matrix are visibly not ideal, even though the latter one requires twice more parameters than the exact functional expression of $K_{exact2} = 5K^{(1)} - 4K^{(2)}$ where the matrices $K^{(i)}$ are given by the form (3).

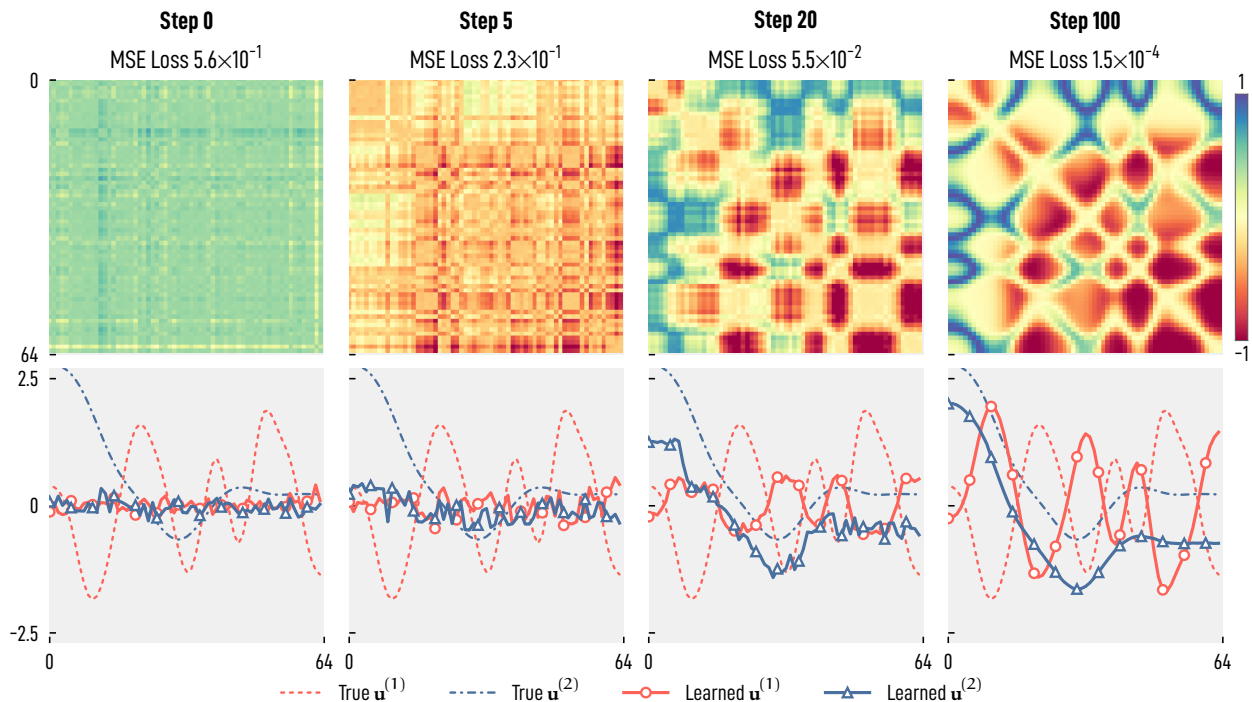


FIGURE 2. Learning process via Adam method of the two RBF component approximation of K_{exact2} matrix: initial matrix is nearly constant (first row), since randomly initialized component vectors $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ are chosen to have small coordinates all about $1e-2$ (second row), and, as the optimization proceeds, after 5, 20 and 100 steps, the component vectors gradually learn the shape of the true component vectors, and the learned matrix approximates K_{exact2} .

data matrix K with no initial structure assumed? This is parallel to finding an SVD approximation of a given rank for a target matrix but poses a much more challenging optimization problem due to the nonlinearity of the RBF components. In this paper, we give promising answers to all three key questions above. In Figure 2, we demonstrate that for a matrix K_{exact2} , the component vectors $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ can be found by an efficient iterative optimization process. In the next section, we dive into the specifics of this training process that allows us to learn the RBF component vectors and coefficients for a given data matrix.

3. THE OPTIMIZATION PROBLEM

For a matrix $K = (k_{ij}) \in \mathbb{R}^{n \times m}$, constructing an r -component RBF approximation is a problem of finding a sequence of column vectors $\{\mathbf{u}^{(k)} \in \mathbb{R}^n, k = 1, \dots, r\}$ and row vectors $\{\mathbf{v}^{(k)} \in \mathbb{R}^m, k = \{1, \dots, r\}\}$, a vector of coefficients $\mathbf{a} \in \mathbb{R}^r$ and a scalar offset b such that

$$(4) \quad k_{ij} \approx k_{ij}^0 := b + \sum_{k=1}^r a_k \exp \left[- \left(u_i^{(k)} - v_j^{(k)} \right)^2 \right]$$

This leads to the following optimization problem

$$(5) \quad \operatorname{argmin}_{\{\mathbf{u}^{(k)}, \{\mathbf{v}^{(k)}, \mathbf{a}, b\}} \mathcal{L}(K, K^0), \quad \text{where } \mathcal{L}(K, K^0) = \frac{1}{nm} \sum_{i,j=1,1}^{n,m} \left(k_{ij} - k_{ij}^0 \right)^2$$

is the mean squared error loss function. Although other losses could be considered depending on an application, we choose an L_2 -based loss to establish direct comparison with SVD-based approximations, which is the best low-rank approximation for the loss as guaranteed by the Eckart–Young–Mirsky theorem [5].

Global minimum structure. The problem (5) is nonlinear, non-convex, and potentially of very high dimension $\sim \mathcal{O}((n+m)r)$. In particular, the global minimum may not be unique as multiple equivalent solutions exist due to the permutation of the RBF components. It is only obvious in the one-component case, where the entries of K correspond to one-dimensional proximities between the vector elements as discussed above, that the set of global minima consists of the ground truth vector \mathbf{u} uniquely determined up to reflection and a constant shift. Empirically, we can see in Figure 2 how two learned component vectors differ from true component vectors only by isometries.

Adaptive iterative gradient-based optimization: Adam. We turn to *gradient-based optimization techniques* for solving the problem (5). It is known that the stochastic gradient descent method used with adaptive step size algorithms can successfully handle complex non-convex objectives [6, 7, 8]. Their practical use in the deep network community goes way beyond their established theoretical guarantees. As visualized in Figure 3, the learning dynamic of the two-component matrix K , introduced in Section 2, suggests that Adam [9] can generally outperform other popular gradient descent algorithms. Therefore, we use it for all subsequent experiments presented in the paper.

Vanishing gradients and small magnitude initialization. The problem of vanishing gradient may halt the training process when entries of the learned vectors have large mutual distances. To illustrate this issue, let us consider the problem of learning the vector \mathbf{u} from a matrix K of the simplest one-component symmetric form (3). The corresponding loss function is

$$(6) \quad \mathcal{L} = \frac{1}{n^2} \sum_{ij} \left[\exp(-(u_i - u_j)^2) - k_{ij} \right]^2$$

and its (i, j) -th gradient is given by $-2/n^2 (u_i - u_j) \exp(-(u_i - u_j)^2)$, which approaches zero both for very large and especially small values of $|u_i - u_j|$. In particular, the training stalls (or stops) if we start with the points that are too far from each other. We partially alleviate this issue by making sure to initialize the parameters in component vectors \mathbf{u} with small enough values. The experimental evidence confirms that this makes the training process more consistent: the histograms in Figure 3 show the distributions of the losses achieved on the batch of 100 runs, all started randomly with the independent standard normal entries, multiplied by a factor of 1 and 0.1 respectively. We can see that smaller initialization values lead to a larger fraction of independent training runs resulting in a global minimum. We can also see in Figure 2 how the magnitude of the parameters gradually grows during the learning process.

Batch training: only the best-run matters. However, the histograms in Figure 3 show a non-negligible fraction of runs that do not converge to a global minimum, for both initialization options. The crucial advantage of our problem comparing to the majority of typical machine learning problems (when a gradient-based approach is used for training the model) is that we always know the ground truth matrix K , and we do not need the learned parameters to generalize to the unseen data. Thus, we can track the error during the training process. The following process can be employed to get a consistent convergence result: after training a batch of gradient descent processes with various random starts, we keep only the best one (or several), and restart all the others. In this paper, we employ a simpler approach: we simply run a *batch of independent learning processes* concurrently and remember that it is enough to get just one good run to learn an efficient RBF approximation. Essentially, our task




FIGURE 3. Learning two (exactly-parametrized regime) and four (over-parametrized regime) RBF component approximations of K_{exact2} : (A) and (B) show that standard adaptive gradient-based learning methods, such as Adam, AdamW and Adagrad, can find approximation K_{exact2} with the MSE loss below $1e-5$ in about 1000 iterations of 100 independent runs, and that convergence tends to be quicker in the over-parametrized regime. (C) and (D) show the distribution of the resulting MSE approximation loss over 100 independent randomly started runs after 10K iteration steps of the Adam algorithm, depending on the starting random distribution of the component vector coefficients (standard Gaussian, or 10 times damped). The initialization with the smaller random values, and especially the over-parametrization, lead to additional robustness of the learning process.

requires an algorithm design that will result in a non-negligible (and not necessarily high) probability of successful training.

Random initialization and over-parametrization. It is known that the gradient-based methods are very sensitive to the initialization, and it is always advantageous to have a warm start for the training. Although the magnitudes of the entries of K give some information about the distances between the coordinates of the component vectors, such effect is hard to quantify especially for fitting an approximation with multiple RBF components. Here, we do not tune the starting distribution for the specific data matrix, but rather initialize randomly. The promise of random initialization was proposed in several recent papers, including [10], where the authors demonstrate that random initialization works for a tensor CP-fitting problem in the over-parametrized regime. In our RBF fitting problem, over-parametrization also leads to the increased speed and robustness of the training process, as we can see on the right two displays in Figure 3.

Stochastic counterparts. *Stochastic* gradient descent is often used in training deep neural networks to reduce the computational cost while also providing the ability to overcome local minima as a side effect of the introduced randomization. The key here is to estimate the gradient of the loss function with respect to the learnable parameters using a small subset of the training examples chosen randomly during each gradient descent iteration. In the context of solving (5) for the RBF decomposition, this translates into evaluating the loss function over a subset of the elements of the target matrix. In other words, we evaluate the loss function as

$$(7) \quad \tilde{\mathcal{L}} = \frac{1}{|I|} \sum_{i,j \in I} (k_{ij} - k_{ij}^0)^2,$$

where I is an index set that contains a random subset of the elements of the matrix coordinates and $|I|$ denotes its cardinality.

Let us conclude this section with a disclaimer that it is still an important direction for future work to analyze the loss landscape specific to the RBF fitting problem and to develop tailored initialization and optimization processes for this particular shape of the loss function.

4. EFFICIENCY OF THE RBF APPROXIMATION

As we discussed in Section 2, an RBF approximation is more efficient than any low-rank approximation for at least *some* matrices. Further, observing that *every* matrix is exactly a linear combination of several rank-one matrices, we claim that a natural analog holds: every matrix can be approximated by a linear combination of the RBF components given by (2). One can notice that by constructing $u_{i'} = v_{j'}$ for a fixed pair of indices (i', j') and setting all other elements very large so that $|u_i - v_j|, i \neq i' \text{ or } j \neq j'$ is greater than multiple length scales of the RBF function, we obtain a matrix \tilde{K} very close to being having exactly one non-zero entry. This gives a trivial, albeit memory-heavy, universal approximation for any matrix in $\mathbb{R}^{n \times m}$ using at most nm components of type (2). Of course, approximating only one matrix element using an RBF component is completely impractical, since storing a truncated SVD decomposition or even storing the elements of the original matrix in verbatim would take less memory. This motivates the next key question we address in this paper: *For what natural models is the RBF approximation of a general data matrix K more successful than one based on SVD?*

In this section, we compare the efficiency of the approximations by RBF and SVD components in terms of reconstruction quality as measured by the MSE loss (5) over a variety of the data matrix models. Note that in the symmetric case, right and left singular vectors coincide, and so the SVD-based low-rank approximation depends on twice fewer parameters. Parallel to this, we use symmetric RBF components of form (3) to approximate symmetric matrices. Finally, it is worth noting that there are "adversarial" settings where SVD approximations will show better results than the RBF ones. For example, a rank 1 matrix can be approximated exactly by one SVD component, but generally not by one RBF component. However, as the rank r of the target matrix grows, we can see that the rank $r' \lesssim r$ approximations by the RBF components quickly become more efficient. See also the discussion about the correct functional form below in Section 7.

4.1. Synthetic Data. First, we consider the following artificial random matrix models: 1) a square matrix of random Gaussian noise, 2) a rectangular matrix of random Gaussian noise, 3) the adjacency matrix of the Erdős-Rényi graph, and 4) the adjacency matrix of the Barabási-Albert graph [11]. We show that RBF component approximations are more efficient than truncated SVD approximations even for these highly unstructured models. This suggests that the RBF approximation is not merely more efficient due to the exploitation of special structures in the data, *e.g.* as exemplified in Figure 1. The RBF components learn *general* data matrices better while being defined by the same number of parameters as the standard outer products because they are more informative and flexible as manifested by the higher number of ranks that each component possesses.

Figure 4 (A)/(B) shows the decay of the mean square loss as we increase the number r of RBF components. For the random matrices with independent standard Gaussian entries, approximately twice the number of SVD ranks are required to achieve the same mean square error with the RBF approximations.

Figure 4(C)/(D) show the decay rate of reconstruction loss for the adjacency matrices of a Erdős-Rényi random graph and a Barabási-Albert random graph. The Erdős-Rényi graph has 40 vertices and a probability of 0.5 for each independent undirected edge. The Barabási-Albert random graph also has 40 vertices and is formed by the preferential attachment model so that each new vertex is attached with 3 undirected edges. See, *e.g.*, [11] for further details on the construction. For very small r , the MSE loss of the RBF and SVD models is similar for both graph models while the RBF approximations have a small advantage. However, the RBF model exhibits much a faster reduction of loss as the number of components grow. In particular, for the Erdős-Rényi model, exact reconstruction with an MSE error below 10^{-5} requires 18 RBF components and 39 SVD components, respectively. For the preferential attachment model, 8 RBF and 37 SVD components are needed for the exact reconstruction.

4.2. Molecular Distance Matrix. In the next example, we try to approximate a distance matrix between a set of molecules in the reproducing kernel Hilbert space (RKHS) induced by a marginalized graph kernel [12, 13, 14]. Here, each molecule is represented by a graph where nodes represent atoms and edges represent inter-atomic interactions. The inner product between two molecules under the graph kernel is the expectation of the path similarities between a simultaneous random walk on the two corresponding molecular graphs. Given a graph kernel $\langle \cdot, \cdot \rangle$ and a pair of molecular graphs G and G' , the distance is computed using the geometric definition

$$(8) \quad d_{G,G'}^2 = \langle G - G', G - G' \rangle = \langle G, G \rangle + \langle G', G' \rangle - 2\langle G, G' \rangle,$$

where the subtractions in the intermediate step are to be understood as happening in the feature space.

The matrix in Figure 5(A) is computed from a random subset of 256 molecules from the QM7 dataset [15]. Figure 5(B)/(C) show the MSE losses in log scale of the RBF approximations and the SVD approximation with the




FIGURE 4. A comparison of MSE error achieved by RBF and SVD approximations depending on the number of components/ranks. Various full-rank random models considered: (A) rectangular and (B) square standard asymmetric Gaussian matrices and (C) and (D) adjacency matrices of random graphs. Consistently over these models, several times less parameters are needed to achieve the same MSE accuracy with the RBF approximation.




FIGURE 5. A comparison of RBF and SVD approximations, as well as different solution algorithms for RBF approximations, of a molecular kernel distance matrix in terms of accuracy and time consumption. (A) the molecular distance matrix; (B) construction loss of RBF and SVD approximations; (C) zoom-in comparison of the performance of full-batch gradient descent (GD) and mini-batch stochastic gradient descent (SGD) algorithms; (D) comparison of GD and SGD in terms of time-to-accuracy.

same number of ranks. To further demonstrate the advantage of the RBF approximation, in Table 1, we collect the number of components required to achieve certain approximation accuracy, in particular, there is more than seven times memory save to reach a 0.0005 MSE loss.

MSE	# SVD ranks	# RBF components
< 0.01	3	1
0.001	30	7
0.0005	77	10

TABLE 1.

Green plots in Figure 5 refer to the RBF approximation obtained by the stochastic optimization approach (discussed above at the end of Section 3) and black plots refer to the learning process based on the full gradient. In Figure 5 (C), we can see that a non-stochastic training achieves slightly smaller MSE loss for the same parameter set-up, and Figure 5 (D) shows that it costs us a 2.5-3 times longer training time. Here, the batch size is 100,

the learning rate is 0.1 and the number of Adam iteration steps is chosen to be $10000k$ for k component approximation since this is approximately the number of iterations required for the best run of a non-stochastic training process to stabilize at some loss level. Stochastic version mini-batch size is set to be $5kn$, where k is the number of components learned and $n = 256$ is the size of the matrix. Moreover, a stochastic training process can be administered significantly lighter than described, allowing to search for the RBF approximation under the limited memory resources. So, in Figure 5 (B), we demonstrate that stochastic Adam can find an RBF approximation with k being up to 128 that significantly outperforms corresponding SVD-based approximation. For this experiment, the batch size is taken to be 50, there is a hard limit to 50000 iterations and mini-batch size is set to be $8n$ for every number of components k . We note that 64 RBF components are observed to be enough for an almost perfect RBF approximation (with the MSE loss below $1e-4$).

5. INTERPRETATION OF THE RESULTING DECOMPOSITION

RBF components of form (2) and (3) can discover one-dimensional distance-related relationships because the underlying radial basis function is a distance-based positive definite kernel. As we have discussed in the previous section, the RBF approximation can reach the same L_2 -error with the SVD approximation using significantly fewer components and parameters across a variety of natural matrix models. However, the advantages of the RBF approximation go beyond memory savings. Just like any approximation of a large data object using a few simpler and lower-dimensional components, the RBF approximation provides a means for capturing simpler underlying structures in the data. In this section, we illustrate this on a sequence of diverse applications, *e.g.*, learning one-dimensional geometry in a three-dimensional manifold, uncovering structures in graphs, and compressing image data.

Before we proceed with the details, let us note that the applications that we consider are among the most popular and widely studied applied math problems. Consequently, each of them is associated with multiple specialized and deeply studied solution approaches. Here, our aim is *not* to compare the RBF-based approach with the state-of-the-art solutions specialized to any of these applications, but rather to demonstrate the universality of the RBF approach over several diverse tasks. We hope that this can motivate the competitive promise of the RBF approximations and set a stage for further exploration.

One-dimensional manifold learning. Consider an S-shaped surface

$$\left\{ \left(\sin t, 2y, \text{sign}(t)(\cos t - 1) \right) \mid t \in \left[-\frac{3}{2}\pi, \frac{3}{2}\pi\right], y \in [0, 2] \right\}$$

in a 3D space, where t represents the “time” coordinate along the curve while y represents the coordinate along the “width” dimension of the letter S. A noisy dataset $X = (\mathbf{x}_1, \dots, \mathbf{x}_{1000})$ is drawn from this surface by sampling t and y uniformly 1000 times and then perturbing the points by an additive Gaussian noise $Z \sim \mathcal{N}(0, \delta \mathbf{I})$.

The approximation target is the soft distance matrix K for the dataset, namely, its (i, j) -th entry is defined as $k_{ij} = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|^2/2)$ for $i, j, = 1, \dots, 1000$. One symmetric RBF component as per definition (3) is fitted to approximate K , resulting in a vector $\mathbf{u} \in \mathbb{R}^{1000}$ that approximately recovers the time coordinate t . See the relationship between the ground truth order along the curve ($\mathbf{t} = (t_1, \dots, t_{1000})$) and learned order (\mathbf{u}) for four noise levels δ from 0 (no noise) to 0.6 in Figure 6.

Community detection. The stochastic block model (SBM) is a random graph model with planted clusters, widely used for community detection and clustering tasks. See, *e.g.*, [16] for a survey of modern approaches to the SBM-based clustering, as well as their limitations. One of the most standard approaches to find k clusters in an SBM matrix is spectral clustering using the entries in the k top leading eigenvectors. In contrast, only one RBF component can suffice to successfully detect multiple clusters. As demonstrated in Figure 7, we consider an 80-vertex graph with 5 communities each containing 8, 12, 16, 20 and 24 vertices, respectively. All clusters have an intra-cluster probability of 0.8 and an inter-cluster probability of 0.2. Like in the previous example, we fit one symmetric RBF component (3) to the adjacency matrix of the graph. A one-dimensional clustering is carried out on the learned coordinates of the component vector u to assign the respective graph vertices into 5 clusters.

Approximation of graph adjacency matrices. We return to the example of the Erdős-Rényi (40, 0.5) random graph considered in Section 4 with the observation that the L_2 -based MSE loss is probably not the most important measure of a meaningful graph compression. Instead, what one might want to preserve in the approximation is the connectivity of the graph. In Figure 8, we can visually compare the difference in how well the RBF and SVD approximation models can capture the structural patterns in the adjacency matrix. To quantify this observation,



FIGURE 6. First row: Visualization of the S-curve dataset with various added noise levels. Second row: Linear relation between Ground truth (\mathbf{t}) “time” coordinate parametrizing given S dataset and Learned embedding (\mathbf{u}) demonstrates that the increasing or decreasing order of the coordinates of \mathbf{u} can be used to uncover main one-dimensional structure of the dataset.




FIGURE 7. Left: adjacency matrix of an SBM random graph with 5 communities with 8, 12, 16, 20 and 24 vertices, sampled from the model with intra-cluster probability of 0.8 and an inter-cluster probability of 0.2. Middle: a one component RBF approximation visually captures all 5 communities whereas the rank one SVD-based approximation fails. Right: One-dimensional clustering of the coordinates of the RBF approximation component vector allows to recover community attribution for all the vertices of the given graph.




FIGURE 8. Left: adjacency matrix of an Erdős-Rényi random graph on 40 vertices with edge probability 0.5. Middle: a 2 component RBF approximation captures visibly more local structure comparing to the rank 2 SVD approximation. Right: RBF-based edge prediction algorithm dominates its SVD-based counterpart for 2-component and 7-component approximations. Seven component RBF-based approximation allows for an almost exact edge reconstruction.

we formalize the following *edge prediction algorithm*: we threshold the real-valued entries of the approximation matrix by some cutoff value $\varepsilon \in [0, 1]$ and predict an edge if the element of the approximation matrix is greater than ε (thus, it is close enough to its value in the binary adjacency matrix). We confirm the advantage of the RBF approximation in retaining the true edge structure by the difference in the ROC curve and AUC score ([17]) for the SVD-based and RBF-based edge prediction algorithms with 2 and 7 components (see Figure 8, right).

Image data compression. It is common to consider natural image data as inherently low-rank, this assumption makes it both theoretically and practically amenable for various denoising and compressive sensing tasks. Our preliminary experiments suggest that, unlike exactly low-rank matrices, real images can be approximated with the RBF components at least just as well as with the SVD-based components. Moreover, RBF-based approximation is even better at preserving visual acutance using the same number of parameters. Here, we compare the quality of RBF and SVD approximations on two images: 1) a baboon picture that is widely used in the image processing literature, and 2) a poster from a 2009 documentary “Space Tourists”. For RBF approximations, we fit the images using multiple components of the general asymmetric form (2).

For the “baboon” image, MSE-losses are slightly better in the RBF approximation as shown in Figure 9, while the visual quality of the compressed image is notably higher. For the “Space Tourists” poster, MSE-losses are very similar in approximations with the same number of components, however, we can also observe the advantages of the RBF compression in the visual quality both in the text and the portrait parts of the poster. It is an important follow-up direction of the current work to define the quantitative measure that captures the differences in the sharpness of the image compression better than the generic L_2 -based MSE loss.

6. RELATED WORK

The literature on the low-rank approximation of matrices is truly abundant.

While SVD is known to give the best L_2 norm approximation of a given rank, numerous treatments develop efficient algorithms for computing and approximating SVD, making the computations scalable and robust, and use them to do dimension reduction, denoising, and learning from data. For a survey, see, e.g., [18, 19]. Let us briefly note here that slightly more sophisticated low-rank approximation tasks, such as low-rank tensor approximations, or non-negative factorizations, rapidly increase in complexity of training, as well as existence and uniqueness questions. It is common to turn to the iterative optimization algorithms to solve them (e.g., [20]).

Beyond the low-rank approaches, special properties of the underlying data can be used to learn its approximation by some simple components. One classical example is approximation of kernel matrices. Many approaches were introduced to make kernel trick scalable, among the most prominent ones is the approximation of a kernel matrix by an scalar product of higher-dimensional vectors, namely $K(\mathbf{x}^i, \mathbf{x}^j) \approx \sum_k \phi(\mathbf{x}^i)_k \phi(\mathbf{x}^j)_k$. The function ϕ here can be constructed using randomly samples from the Fourier transformation of the kernel, versions of the Nyström method, or using polynomial approximation [21, 22, 23]. All these methods use positive semi-definiteness

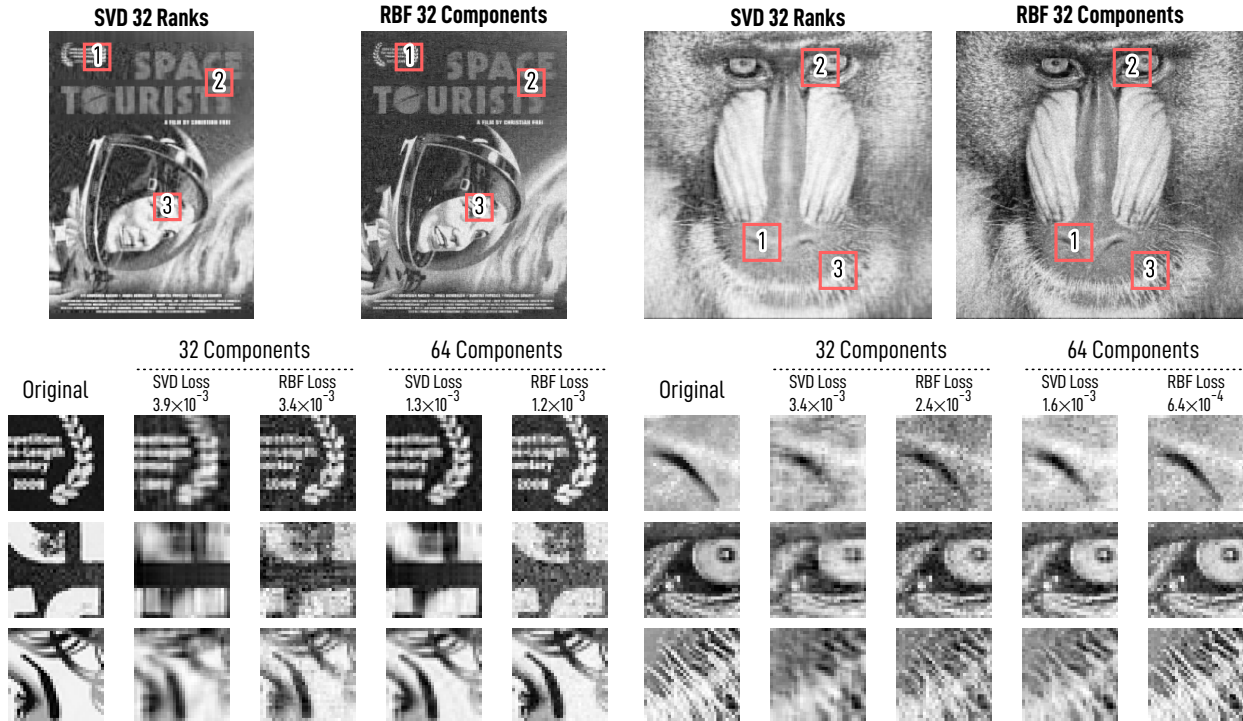


FIGURE 9. Visual comparison of the RBF and SVD based approximations for "Space tourists" poster and "Baboon" image with 32 and 64 components. Low rank approximations tend to be smoother and less noisy, while RBF components produce sharper images, better preserving local features of an image, such as textures and edges.

of the kernel matrix to be approximated, and sometimes its concrete functional form, and neither of which is imposed in our approach.

The idea to use kernel functions *for approximation* had several appearances in the recent literature, usually still limited to the approximation of positive semi-definite or otherwise very special matrices. For example, in [24], compositions of kernels used for nonparametric regression. In [25, 26], the authors consider approximations of functions defined on a known discrete dataset in \mathbb{R}^d by linear combinations of d -dimensional gaussian RBF functions, as well as by more sophisticated building blocks. Although our setting is crucially different from this one, where only the coefficients are learned, but the component matrices are known by design, which results in a linear system to solve instead of a highly non-convex optimization problem.

A method that is intimately related to our current approximation approach is the Radial-Basis-Function Network. An RBF network as defined in [27] is a three layer neural network with RBF nonlinearities in its hidden layer. Effectively, it learns, *i.e.* approximates, an unknown function as a linear combination of multivariate Gaussian kernel functions. RBF networks are known for the universal approximation properties [27]. The conceptual difference between our RBF approximation with neural networks trained by a supervised learning task is that we also learn the input data. From a neural network perspective, we treat the input training data as a set of learnable parameters in addition to the weights. Moreover, in the optimization process, we are not concerned with the generalizability of the learned model because a different set of parameters are learned for *each* concrete data matrix K . This makes our target different: we aim to make the model structure as light as possible. Clearly, an approximation exists for any data with enough overfitting, as illustrated in Section 4. Our primary interest is in the most compact one.

7. CONCLUSION AND FUTURE WORK

We present a novel framework for matrix data approximation based on nonlinear RBF components. We restrict the building blocks to the one-dimensional kernels, which gives us a simple way to control complexity¹. Unlike composite kernel learning, we consider arbitrary linear combinations of kernel matrix component and deliberately do not constrain the resulting matrix to be semi-positive definite. This relaxation enhances the flexibility of the approximation and improves efficiency. We demonstrate the advantage of the proposed approximation, both in terms of memory saving and interpretability, across various datasets. We employ gradient-based method for learning the nonlinear RBF components, discuss the particulars of the learning setup and demonstrate convergence to a low-component RBF approximation.

This is an initial exploratory treatment of a new method, which gives rise to an abundance of natural follow up questions and possibilities of future extensions. For example, our current symmetric matrix fitting is bound to have ones in the diagonal. A simple correction of the functional form, such as multiplying the symmetric outer product to the RBF component matrix, can be used to approximate symmetric matrices with non-constant diagonal while keeping the number of parameters the same. More generally, changing the functional form brings in more flexibility to the method. It would be even more interesting to learn the correct functional form from the data. Another direction of a follow-up work is to formalize similar RBF-based approximations for multi-order tensor data, *e.g.*, a multi-time series dataset, a color image, or a video clip.

Our current optimization strategy is quite generic and not specialized for a particular functional form. If the nonlinear components degenerate into standard outer products, *e.g.* as in the case of SVD and tensor CP decomposition, then better ways are known to solve the approximation problem much faster than gradient descent based approaches. Unfortunately, the application of linear solver-based approaches are likely ruled out by the Gaussian nonlinearity. Nonetheless, it would be very interesting to leverage more general optimization strategies beyond gradient descent that can still benefit from some structural assumptions about the model.

Compressed data in the form of an RBF approximation can be used as an intermediate representation in many machine learning pipelines. For example, an RBF approximation can compress neural networks models or massive training datasets. Finally, it is a very natural future direction to further investigate particular simple data features that nonlinear approximation preserves in various applications.

8. ACKNOWLEDGMENT

This work was supported by the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under U.S. Department of Energy Contract No. DE-AC02-05CH11231.

REFERENCES

- [1] J. Douglas Carroll and Jih-Jie Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of "Eckart-Young" decomposition. *Psychometrika*, 35(3):283–319, September 1970.
- [2] Richard A. Harshman. Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multimodal factor analysis. *UCLA Working Papers in Phonetics*, 16:1–84, 1970.
- [3] Thomas Hofmann, Bernhard Schölkopf, and Alexander J. Smola. Kernel Methods in Machine Learning. *The Annals of Statistics*, 36(3):1171–1220, 2008.
- [4] Vladimir Koltchinskii, Evarist Giné, et al. Random matrix approximation of spectra of integral operators. *Bernoulli*, 6(1):113–167, 2000.
- [5] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [6] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nathan Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. *arXiv preprint arXiv:1705.08292*, 2017.
- [7] Michael O'Neill and Stephen J Wright. Behavior of accelerated gradient methods near critical points of nonconvex functions. *Mathematical Programming*, 176(1):403–427, 2019.
- [8] Soham De, Anirbit Mukherjee, and Enayat Ullah. Convergence guarantees for rmsprop and adam in non-convex optimization and an empirical comparison to nesterov acceleration. *arXiv preprint arXiv:1807.06766*, 2018.
- [9] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [10] Xiang Wang, Chenwei Wu, Jason D Lee, Tengyu Ma, and Rong Ge. Beyond lazy training for over-parameterized tensor decomposition. *arXiv preprint arXiv:2010.11356*, 2020.
- [11] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

¹Moreover, multi-dimensional kernels can be approximated by one-dimensional ones, just like any other data matrix, so, considering heavier multi-dimensional components in our framework is probably redundant.

- [12] Hisashi Kashima, Koji Tsuda, and Akihiro Inokuchi. Marginalized kernels between labeled graphs. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 321–328. AAAI Press, 2003.
- [13] Yu-Hang Tang and Wibe A. de Jong. Prediction of atomization energy using graph kernel and active learning. *The Journal of Chemical Physics*, 150(4):044107, January 2019.
- [14] Yan Xiang, Yu-Hang Tang, Hongyi Liu, Guang Lin, and Huai Sun. Predicting Single-Substance Phase Diagrams: A Kernel Approach on Graph Representations of Molecules. *The Journal of Physical Chemistry A*, May 2021.
- [15] Lorenz C. Blum and Jean-Louis Reymond. 970 Million Druglike Small Molecules for Virtual Screening in the Chemical Universe Database GDB-13. *Journal of the American Chemical Society*, 131(25):8732–8733, July 2009.
- [16] Emmanuel Abbe. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research*, 18(1):6446–6531, 2017.
- [17] Tom Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.
- [18] Ravindran Kannan and Santosh Vempala. *Spectral algorithms*. Now Publishers Inc, 2009.
- [19] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [20] Andrzej Cichocki, Rafal Zdunek, Anh Huy Phan, and Shun-ichi Amari. *Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation*. John Wiley & Sons, 2009.
- [21] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Proceedings of the 20th International Conference on Neural Information Processing Systems, NIPS’07*, pages 1177–1184, Red Hook, NY, USA, December 2007. Curran Associates Inc.
- [22] Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir Zandieh. Random Fourier features for kernel ridge regression: Approximation bounds and statistical guarantees. In *International Conference on Machine Learning*, pages 253–262. PMLR, 2017.
- [23] Ruoxi Wang, Yingzhou Li, and Eric Darve. On the Numerical Rank of Radial Basis Function Kernels in High Dimensions. *SIAM Journal on Matrix Analysis and Applications*, 39(4):1810–1835, January 2018.
- [24] David Duvenaud, James Lloyd, Roger Grosse, Joshua Tenenbaum, and Ghahramani Zoubin. Structure Discovery in Nonparametric Regression through Compositional Kernel Search. In *International Conference on Machine Learning*, pages 1166–1174. PMLR, May 2013.
- [25] Michal Smolik and Vaclav Skala. Radial basis function and multi-level 2d vector field approximation. *Mathematics and Computers in Simulation*, 181:522–538, 2021.
- [26] Marek Zabran. Radial basis function approximation optimal shape parameters estimation: preliminary experimental results. 2020.
- [27] J. Park and I. W. Sandberg. Universal Approximation Using Radial-Basis-Function Networks. *Neural Computation*, 3(2):246–257, June 1991.