



LEARNING IN HIGH-DIMENSIONAL FEATURE SPACES USING ANOVA-BASED FAST MATRIX-VECTOR MULTIPLICATION

FRANZISKA NESTLER

Department of Mathematics, Chair of Applied Functional Analysis
TU Chemnitz, Germany

MARTIN STOLL AND THERESA WAGNER*

Department of Mathematics, Chair of Scientific Computing
TU Chemnitz, Germany

(Communicated by Benjamin Peherstorfer)

ABSTRACT. Kernel matrices are crucial in many learning tasks such as support vector machines or kernel ridge regression. The kernel matrix is typically dense and large-scale. Depending on the dimension of the feature space even the computation of all of its entries in reasonable time becomes a challenging task. For such dense matrices the cost of a matrix-vector product scales quadratically with the dimensionality N , if no customized methods are applied. We propose the use of an ANOVA kernel, where we construct several kernels based on lower-dimensional feature spaces for which we provide fast algorithms realizing the matrix-vector products. We employ the non-equispaced fast Fourier transform (NFFT), which is of linear complexity for fixed accuracy. Based on a feature grouping approach, we then show how the fast matrix-vector products can be embedded into a learning method choosing kernel ridge regression and the conjugate gradient solver. We illustrate the performance of our approach on several data sets.

1. Introduction. Kernel methods [31, 15] are an important class of machine learning algorithms ranging from unsupervised to supervised learning. At the heart of these algorithms lies the kernel matrix $K \in \mathbb{R}^{N,N}$ with entries $K_{ij} := \kappa(\mathbf{x}_i, \mathbf{x}_j)$, which arises from evaluating the kernel function $\kappa(\cdot, \cdot)$ on the given data points $\mathbf{x}_i \in \mathbb{R}^d$. The kernel matrix is typically dense and large-scale, which makes computations with it computationally demanding [33, 28, 12]. In order to reduce the computational cost of working with dense kernel matrices many approaches have been considered in the literature. Among them fast transforms are a very promising group of methods. We here point to [20, 27] for the fast Gauss transform and to [2, 24] for the use of the non-equispaced fast Fourier transform (NFFT), which we will focus on in this paper. For a comparison of the fast Gauss transform with the NFFT-approach we refer to [2]. Alternatively, sketching has been considered as a technique for solving kernel ridge regression problems [3] both for the task of forming a suitable preconditioner and performing the matrix-vector products efficiently.

2020 *Mathematics Subject Classification.* Primary: 65F45, 65T50; Secondary: 62J10.

Key words and phrases. ANOVA kernel, kernel ridge regression, non-equispaced fast Fourier transform, fast summation, fast matrix-vector multiplication, multiple kernel learning.

*Corresponding author: Theresa Wagner.

Furthermore, fast parallel solvers for kernel matrices [37] have been proposed computing an approximate factorization of the kernel matrix in the fashion of N -body methods [18] with $\mathcal{O}(N \log N)$ work.

Fast transform-based methods are typically hit by a curse of dimensionality when the dimensionality d of the feature vectors $\mathbf{x}_i = (x_i^1, \dots, x_i^d)^\top \in \mathbb{R}^d$ grows without any additional assumptions made on the kernel function. In order to carry over the power of the NFFT-based fast matrix-vector products, we here propose the use of an ANOVA kernel that is a combination of kernels relying on a small number of features and hence enables fast multiplication methods. While we focus on kernel ridge regression, the proposed technique can be embedded into a variety of kernel-based learning schemes such as deep kernel methods [36], Gaussian processes [26] or graph-based learning [17].

Figure 1 indicates that the matrix-vector products using the NFFT will outperform the dense matrix-vector products even for small dimensions N . We use the remainder of this paper to discuss how this is achieved in detail. Naturally, the results will depend on the data, which for this example were generated randomly, and the choice of parameters. Figure 1 is meant to give an impression of the power of ANOVA-based fast matrix-vector multiplications.

We show in the remainder of this paper how such fast multiplications and a new feature grouping can be embedded into kernel-based learning. In Section 2 we introduce kernel ridge regression, which is followed by a discussion of the ANOVA kernel. We there show that this kernel allows a grouping of the features and after discussing the NFFT-based multiplication in Section 4, we introduce a novel multiple kernel grouping based on mutual information scores. Our experiments indicate that the new kernel results in much faster training and prediction times, especially for large data sets, without or with very benign loss of accuracy. These computations can be carried out on standard hardware without requiring additional high performance computing power.

2. Kernel ridge regression. As already eluded to, we focus on the case of kernel ridge regression as a learning method that utilizes positive kernels [31, 15]. It is targeted at detecting dependencies between features and the label to successfully predict class affiliations for new data points from the same source. The underlying theory of machine learning is well developed for the linear case indeed [31]. Real-world problems typically require non-linear methods though. By learning in a reproducing kernel Hilbert space (RKHS) [30], this difficulty can be overcome. Replacing dot products by a kernel evaluation allows us to utilize linear methods in this space without ever explicitly having to compute in it (cf. [31, 15, 2]).

Let \mathcal{H} be a RKHS

$$\mathcal{H} := \overline{\text{span} \{ \kappa(\mathbf{x}, \cdot), \mathbf{x} \in \mathcal{X} \}},$$

where \mathcal{X} is a non-empty set and κ is a reproducing kernel of \mathcal{H} on $\mathcal{X} \times \mathcal{X}$. Moreover, let

$$\langle f, g \rangle := \sum_{j=1}^{N'} \sum_{i=1}^N \alpha_i \beta_j \kappa(\mathbf{x}_j', \mathbf{x}_i)$$

define the inner product of elements

$$f(\cdot) = \sum_{i=1}^N \alpha_i \kappa(\cdot, \mathbf{x}_i) \quad \text{and} \quad g(\cdot) = \sum_{j=1}^{N'} \beta_j \kappa(\cdot, \mathbf{x}_j')$$




FIGURE 1. Kernel-vector multiplication with and without NFFT-approach on various balanced samples of different size on synthetic data, with $\sigma = 100$ and $n_{\text{runs}} = 100$

and

$$\|f\|_{\mathcal{H}} := \sqrt{\langle f, f \rangle}$$

the norm in \mathcal{H} . Then, with κ being the reproducing kernel of \mathcal{H} , the evaluation functional

$$\langle \kappa(\cdot, \mathbf{x}), f \rangle = \sum_{i=1}^N \alpha_i \kappa(\mathbf{x}_i, \mathbf{x}) = f(\mathbf{x})$$

holds for all $f \in \mathcal{H}$, $\mathbf{x} \in \mathcal{X}$ [30]. Now, let a data set be given that is represented by feature vectors $\mathbf{x}_i \in \mathbb{R}^d$, $i = 1, \dots, N$, where $N \in \mathbb{N}$ denotes the number of data points and $d \in \mathbb{N}$ the number of features. Furthermore, let each of those vectors possess a label $y_i \in \{-1, 1\}$, so that we obtain pairs (\mathbf{x}_i, y_i) . As explained above, our aim is to train our model on the given labeled data, to be able to predict class affiliations for new unlabeled data later.

For that, we can formulate the learning task

$$\hat{f} = \arg \min_{f \in \mathcal{H}} \sum_{i=1}^N (y_i - f(\mathbf{x}_i))^2 + \lambda \|f\|_{\mathcal{H}}^2 \tag{1}$$

in a RKHS \mathcal{H} , where $\lambda > 0$ is the regularization parameter (cf. [7]). It corresponds to the linear regression problem in \mathcal{X} . By the representer theorem [15], the solution to (1) is of the form

$$f = \sum_{i=1}^N \alpha_i \kappa(\mathbf{x}_i, \cdot), \tag{2}$$

where $\mathbf{x}_i \in \mathbb{R}^d$ are the labeled training samples [30]. Substituting (2) to (1), we obtain the finite-dimensional quadratic program

$$\hat{\boldsymbol{\alpha}} = \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^N} \sum_{i=1}^N \left(y_i - \sum_{j=1}^N \alpha_j \kappa(\mathbf{x}_j, \mathbf{x}_i) \right)^2 + \lambda \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j \kappa(\mathbf{x}_i, \mathbf{x}_j)$$

in $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_N]^\top$, see [15] for details. This can be rewritten as

$$\hat{\boldsymbol{\alpha}} = \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^N} \|\mathbf{y} - K\boldsymbol{\alpha}\|_2^2 + \lambda \boldsymbol{\alpha}^\top K \boldsymbol{\alpha}, \quad (3)$$

where $\mathbf{y} = [y_1, \dots, y_N]^\top$ is the target vector and

$$K = \begin{bmatrix} \kappa(\mathbf{x}_1, \mathbf{x}_1) & \dots & \kappa(\mathbf{x}_1, \mathbf{x}_N) \\ \vdots & & \vdots \\ \kappa(\mathbf{x}_N, \mathbf{x}_1) & \dots & \kappa(\mathbf{x}_N, \mathbf{x}_N) \end{bmatrix}$$

the kernel matrix, being symmetric and positive semi-definite. In a straightforward way, we obtain the gradient of the objective function as

$$K(K + \lambda I)\boldsymbol{\alpha} - K\mathbf{y} \stackrel{!}{=} \mathbf{0}$$

and we get for the optimal $\hat{\boldsymbol{\alpha}}$ that

$$(K + \lambda I)\hat{\boldsymbol{\alpha}} - \mathbf{y} \in \ker(K),$$

which is true for

$$\hat{\boldsymbol{\alpha}} = (K + \lambda I)^{-1} \mathbf{y} + \mathbf{z},$$

where $K\mathbf{z} = \mathbf{0}$. Due to the fact that for $\boldsymbol{\alpha}' = \boldsymbol{\alpha} + \mathbf{z}$ and $f' = \sum_{i=1}^N \alpha'_i \kappa(\mathbf{x}_i, \cdot)$

$$\|f - f'\|_{\mathcal{H}}^2 = (\boldsymbol{\alpha} - \boldsymbol{\alpha}')^\top K(\boldsymbol{\alpha} - \boldsymbol{\alpha}') = 0,$$

we have that $f = f'$ and the only solution we care for is obtained as

$$\hat{\boldsymbol{\alpha}} = (K + \lambda I)^{-1} \mathbf{y}.$$

For $\lambda > 0$ the system to be solved is symmetric and positive definite. Hence, we will rely on the conjugate gradient (CG) method [14] for its solution. The CG method requires the computation of one matrix-vector product per iteration and, thus, the multiplication with the kernel matrix K

$$K\boldsymbol{\alpha} = \left[\sum_{j=1}^N \alpha_j \kappa(\mathbf{x}_i, \mathbf{x}_j) \right]_{i=1}^N \in \mathbb{R}^N \quad (4)$$

is the most expensive task within this procedure. A naive direct computation requires $\mathcal{O}(N^2)$ arithmetic operations, which we will reduce drastically by applying a fast summation approach, as introduced in Section 4.2.




FIGURE 2. Embedding 2-dimensional data into a 3-dimensional space via an embedding map.

3. **The ANOVA kernel.** As indicated above, kernel functions are derived from the use of inner products when the similarity between data points needs to be evaluated. Thereby, the data is embedded into a new, higher-dimensional feature space such that non-linear relations between features can be modeled in a linear way as illustrated in Figure 2, see [31] for details. In this higher-dimensional space the separation of the data is easier to perform, which again can be seen from the illustration in Figure 2.

But the right choice of a kernel is by no means trivial. In this section we design a favorable kernel function for speeding up matrix-vector multiplications for large-scale data.

Let a data set be represented by feature vectors $\mathbf{x}_i = (x_i^1, \dots, x_i^d)^\top \in \mathbb{R}^d$, where d is the number of features and N is the number of data points or samples. After introducing an embedding map

$$\phi : \mathbf{x} \in \mathbb{R}^d \mapsto \phi(\mathbf{x}) \in \mathbb{R}^{d_\phi},$$

with $d < d_\phi$, we recode our pairs of data points from (\mathbf{x}_i, y_i) to $(\phi(\mathbf{x}_i), y_i)$. This enables us to transform non-linear relations to linear ones. Substituting all inner products $\mathbf{x}_i^\top \mathbf{x}_j$ between two feature vectors $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^d$ with the kernel evaluation

$$\kappa(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle$$

is referred to as the *kernel trick* [31].

The kernel crucial for our investigations in this paper is the well-known ANOVA kernel [31]

$$\kappa_{\text{anova}}^{d_A}(\mathbf{x}_i, \mathbf{x}_j) = \sum_{1 \leq l_1 < l_2 < \dots < l_{d_A} \leq d} \prod_{t=1}^{d_A} \kappa_{l_t}^{\text{base}}(\mathbf{x}_i, \mathbf{x}_j)$$

of degree d_A , with base kernels $\kappa_{l_t}^{\text{base}}(\mathbf{x}_i, \mathbf{x}_j)$ extending the kernel evaluation.

In comparison to other kernels, it provides more freedom regarding the choice of considered features and the weighting (cf. [31]).

Given the importance of the Gaussian kernel [26]

$$\kappa_{\text{gauss}}(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{\sigma^2}\right),$$

we use weighted Gaussian kernels for all our base kernels. Moreover, if we set the degree $d_A = 1$, this yields the so-called *simple Gaussian ANOVA kernel*

$$\kappa_{\text{anova}}^{\text{1d-gauss}}(\mathbf{x}_i, \mathbf{x}_j) = \sum_{l=1}^d \eta_l \exp\left(-\frac{\|x_i^l - x_j^l\|_2^2}{\sigma^2}\right),$$

in which the l -th Gaussian base kernel only depends on the l -th feature. This yields a sum of d weighted kernels. Each feature is taken into account in this kernel definition, which seems reasonable. But including particular features, which are either redundant or do not interact with the label at all, attributes more influence to those features than they actually have. By this, relations might be detected, which do not exist in reality. Excluding these detrimental features instead, can help us increase the model's accuracy. Examining each feature's importance separately instead neglects relevant relations between them. Additionally, we must not forget that we aim to minimize the computational complexity. We therefore are well-advised to keep the number of summands low. This gives rise to the *extended Gaussian ANOVA kernel*

$$\kappa_{\text{anova}}^{\text{3d-gauss}}(\mathbf{x}_i, \mathbf{x}_j) = \sum_{l=1}^P \eta_l \underbrace{\exp\left(-\frac{\|\mathbf{x}_i^{\mathcal{W}_l} - \mathbf{x}_j^{\mathcal{W}_l}\|_2^2}{\sigma^2}\right)}_{\kappa_l(\mathbf{x}_i^{\mathcal{W}_l}, \mathbf{x}_j^{\mathcal{W}_l})}, \quad (5)$$

where $P \leq \lceil \frac{d}{3} \rceil$, with index sets

$$\mathcal{W}_l = \{w_1^l, w_2^l, w_3^l\} \subseteq \{1, \dots, d\} \quad (6)$$

and corresponding

$$\mathbf{x}_i^{\mathcal{W}_l} = \begin{bmatrix} x_i^{w_1^l} & x_i^{w_2^l} & x_i^{w_3^l} \end{bmatrix}^\top \quad \text{and} \quad \mathbf{x}_j^{\mathcal{W}_l} = \begin{bmatrix} x_j^{w_1^l} & x_j^{w_2^l} & x_j^{w_3^l} \end{bmatrix}^\top, \quad (7)$$

where $i, j = 1, \dots, N$. By this, exactly 3 features embody the input for every Gaussian base kernel in (5). The motivation for using 3-dimensional inputs is down to the NFFT, see Section 4, which provides the underlying mathematical theory for the presented method for speeding up matrix-vector multiplications for large-scale data. In fact, this method runs our computations very efficiently as long as the input dimension is smaller than 4. Thus, \mathcal{W}_l holding exactly 3 features enables us to exploit the full computational power of the NFFT, while involving as many feature interactions as possible and keeping the number of summands low at the same time. This choice hence fulfills all crucial aspects. Also we choose the indices corresponding to each \mathcal{W}_l only once so that the kernels never share any features. We refer to Section 5 for details on determining the index sets \mathcal{W}_l and the kernel weights η_l .

4. NFFT-based fast matrix multiplication. In this section, we give a short introduction to the NFFT as well as the fast summation approach. With these methods, we are able to reduce the arithmetic complexity from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$.

4.1. NFFT. For an even grid size $\mathbf{M} = (M_1, \dots, M_d)^\top \in 2\mathbb{N}^d$, i.e. a d -dimensional vector with even integer components, we consider multivariate index sets of the form

$$\mathcal{I}_{\mathbf{M}} := \left\{ -\frac{M_1}{2}, \dots, \frac{M_1}{2} - 1 \right\} \times \dots \times \left\{ -\frac{M_d}{2}, \dots, \frac{M_d}{2} - 1 \right\}.$$

For a given set of Fourier coefficients $\hat{f}_{\mathbf{k}} \in \mathbb{C}$, $\mathbf{k} \in \mathcal{I}_{\mathbf{M}}$, on that grid and a set of nodes $\mathbf{x}_j \in \mathbb{T}^d$, $j = 1, \dots, N$, where the torus \mathbb{T} is defined as $\mathbb{T} = \mathbb{R} \bmod \mathbb{Z} \simeq [-\frac{1}{2}, \frac{1}{2})$, the NFFT [10, 6, 32, 9] is an efficient algorithm computing the sums

$$f_j := f(\mathbf{x}_j) := \sum_{\mathbf{k} \in \mathcal{I}_{\mathbf{M}}} \hat{f}_{\mathbf{k}} e^{2\pi i \mathbf{k}^\top \mathbf{x}_j} \quad \forall j = 1, \dots, N, \tag{8}$$

i.e., the function values of a 1-periodic trigonometric polynomial $f = f(\mathbf{x})$ at arbitrary, commonly irregularly distributed nodes \mathbf{x}_j . The arithmetic complexity is $\mathcal{O}(|\mathcal{I}_{\mathbf{M}}| \log |\mathcal{I}_{\mathbf{M}}| + N)$. The accuracy is controlled by several additional parameters which we do not introduce here. For a detailed discussion on parameters involved in applied algorithms we refer to [16]. Note that the size of the grid $\mathcal{I}_{\mathbf{M}}$ grows exponentially with d , since $\min(\mathbf{M})^d \leq |\mathcal{I}_{\mathbf{M}}|$. In other words, the curse of dimensionality makes the evaluation expensive in case of large dimensions. In order to overcome this problem, we exclusively deal with small dimensions ≤ 3 in every NFFT computation, where FFT-based methods on regular grids are still very efficient. Thus, the overall feature dimension, which we expect to be large, is decomposed into multiple ANOVA kernels of degree $d_A \leq 3$, see Section 5 for the multiple kernel approach.

A fast algorithm for the adjoint problem, which performs the summation

$$g_{\mathbf{k}} := \sum_{j=1}^N c_j e^{-2\pi i \mathbf{k}^\top \mathbf{x}_j} \quad \forall \mathbf{k} \in \mathcal{I}_{\mathbf{M}} \tag{9}$$

efficiently, is known as adjoint NFFT or sometimes type-2 NFFT. Here, the nodes $\mathbf{x}_j \in \mathbb{T}^d$ and corresponding coefficients $c_j \in \mathbb{C}$ are given. Using matrix-vector notation, equations (8) and (9) can be written as

$$\mathbf{f} := \Phi \hat{\mathbf{f}} \quad \text{and} \quad \mathbf{g} := \Phi^* \mathbf{c},$$

respectively, where $\Phi = \left(e^{2\pi i \mathbf{k}^\top \mathbf{x}_j} \right)_{j, \mathbf{k}} \in \mathbb{C}^{N, |\mathcal{I}_{\mathbf{M}}|}$ is a non-uniform Fourier matrix and Φ^* is simply its adjoint. The NFFT and its adjoint are the main ingredients of the fast summation approach, which is introduced below.

4.2. Fast summation. The fast summation approach computes sums of the form

$$s(\mathbf{z}_i) := \sum_{j=1}^{N_x} \alpha_j \kappa(\mathbf{z}_i, \mathbf{x}_j) \quad \forall i = 1, \dots, N_z, \tag{10}$$

for two given sets of nodes $\mathcal{Z} := \{\mathbf{z}_i \in \mathbb{R}^d, i = 1, \dots, N_z\}$, $\mathcal{X} := \{\mathbf{x}_j \in \mathbb{R}^d, j = 1, \dots, N_x\}$, coefficients $\{\alpha_j \in \mathbb{C}, j = 1, \dots, N_x\}$ and a given kernel κ efficiently based on combining the above described algorithms (NFFT and adjoint NFFT), cf. [24]. Note that the special case $\mathcal{X} = \mathcal{Z}$ is included. In the context of kernel

ridge regression, such sums have to be computed when evaluating the model (2) at new data points, i. e. \mathcal{X} represents the set of training data and \mathcal{Z} contains new test data. Within the kernel ridge regression we make use of the special case $\mathcal{Z} = \mathcal{X}$ with $N_x = N_z =: N$ in order to multiply with the kernel matrix K , see (4).

While in our setting the kernel κ in (10) represents a single summand κ_l of our specially designed extended Gaussian ANOVA kernel (5), the following discussion is also true for more general kernels.

Note that we restrict our considerations to radial kernels without singularities. In a more general setting, also kernels with a singularity in the origin, as for example $\|z - x\|^{-1}$ or $\log \|z - x\|$ can be handled [11, 25]. In this case we require $z_i \neq x_j$ for all i, j or rather to exclude all terms with $i = j$ in the case that the two sets of nodes coincide.

The basic idea of the fast summation is explained as follows. In the one-dimensional setting we approximate a univariate kernel by a trigonometric polynomial

$$\kappa(z, x) = \kappa(r) \approx \sum_{k=-M/2}^{M/2-1} \hat{c}_k e^{2\pi i k r / h}, \quad (11)$$

where $r := z - x$ and the period h has to be chosen appropriately. Note that we consider non-periodic functions or rather kernels κ and, thus, an approximation by a trigonometric polynomial is not straightforward. Instead of truncating the function and simply sticking together periodic images, which would give a non-smooth periodic function, we introduce a larger period h and construct a smooth transition in order to end up with a periodic function of some higher smoothness.

If the distances between the given nodes satisfy $|z_i - x_j| \leq L$, we choose $h \geq 2L$. In order to get a good approximation we finally embed κ into a smooth periodic function via

$$\tilde{\kappa}(r) = \begin{cases} \kappa(r) & : r \in [-L, L], \\ p(r) & : r \in (L, L + \ell), \end{cases}$$

where p is a polynomial matching the derivatives of κ up to a certain degree such that its periodic continuation $\tilde{\kappa}$ is smooth [11]. For a graphical illustration see Figure 3. We make heavy use of the fact that the kernel is symmetric. If the given nodes lie in an interval of length L , the resulting period h is more than twice as large as the original interval length. The Fourier coefficients \hat{c}_k in (11) are simply computed by applying the ordinary FFT to M equidistant samples of the resulting periodic function.

With some slight modifications this ansatz is also applicable in a multivariate setting with a radial kernel $\kappa(\|\cdot\|)$. In this case, the problem is, roughly speaking, at first considered in an appropriate one-dimensional setting and the multivariate approximation is obtained via rotation [11]. Defining the scaled nodes $\tilde{z}_i := h^{-1}z_i \in \mathbb{T}^d$ and $\tilde{x}_j := h^{-1}x_j \in \mathbb{T}^d$, we obtain

$$\begin{aligned} s(z_i) &\approx \sum_{j=1}^{N_x} \alpha_j \sum_{\mathbf{k} \in \mathcal{I}_M} \hat{c}_{\mathbf{k}} e^{2\pi i \mathbf{k}^\top (\tilde{z}_i - \tilde{x}_j)} \\ &= \sum_{\mathbf{k} \in \mathcal{I}_M} \hat{c}_{\mathbf{k}} \left(\sum_{j=1}^{N_x} \alpha_j e^{-2\pi i \mathbf{k}^\top \tilde{x}_j} \right) e^{2\pi i \mathbf{k}^\top \tilde{z}_i} \end{aligned} \quad (12)$$




FIGURE 3. The univariate kernel function κ , here a Gaussian, is truncated and a polynomial p is constructed such that we end up with a periodic function $\tilde{\kappa}$, which is smooth up to a certain degree. In higher-dimensional settings, the kernel function is truncated at $\|\mathbf{r}\| = L$, i. e., radially. The resulting periodic function will then have the period h with respect to each dimension.

by replacing the kernel κ in (10) by its Fourier representation (11). The inner sums are now obtained via the adjoint NFFT (9), followed by a multiplication with the Fourier coefficients. Finally, the outer sums are computed with the NFFT (8).

The quality of the approximation depends on several parameters, involving the grid size \mathbf{M} . If these parameters, and therewith the basic accuracy of the approximation (11), are kept fixed, the number of required arithmetic operations scales like $\mathcal{O}(N_x + N_z)$.

4.3. Fast multiplication with the kernel matrices. Now, it is easy to see that the multiplication with the kernel matrix K can be realized with the presented fast summation approach, since in our setting we simply consider a sum of radial kernels (5) acting in low dimensions. Here we are in the special setting that the set of nodes \mathcal{X} and \mathcal{Z} coincide, as it is the case within the kernel ridge regression.

In matrix-vector notation, we have the approximation

$$K\alpha \approx \Phi D \Phi^* \alpha,$$

where $\Phi = \left(e^{2\pi i \mathbf{k}^\top \tilde{\mathbf{x}}_j} \right)_{j,\mathbf{k}}$ is defined as before and $D = \text{diag}(\hat{c}_{\mathbf{k}})_{\mathbf{k}}$ is a diagonal matrix containing the precomputed Fourier coefficients of the periodized kernel $\tilde{\kappa}$.

After the training procedure, the predicted values for the testing data are obtained based on (12). The matrix-vector notation reads as

$$\mathbf{s} \approx \Phi_{\mathbf{z}} D \Phi^* \alpha,$$

where the Fourier matrix $\Phi_{\mathbf{z}} := \left(e^{2\pi i \mathbf{k}^\top \tilde{\mathbf{z}}_i} \right)_{i,\mathbf{k}}$ now includes the test nodes \mathbf{z}_i .

5. Multiple kernel learning. The NFFT-approach from Section 4 is very fast for input-dimensions up to 3. However, the presented method is targeted at large-scale problems with a large number of features. The ANOVA kernel, which was introduced in Section 3, helps us decompose the feature dimensions into multiple kernels. That enables us to exploit the speed-up of the NFFT-approach anyhow. Therefore, we are in need of combining multiple kernels. Thus, a strategy on how to identify important features and on how to separate them up into multiple kernels

is needed. Given its low computational complexity, we decide to rank all features by their mutual information score (MIS) [5].

Mutual information is a measure of uncertainty, i. e. it describes how much knowing the value of a feature reduces our uncertainty regarding the label. The mutual information score is a positive value that usually does not exceed 2. Since it is a univariate metric, it cannot detect relations between multiple features. Therefore, we must keep in mind that having a high mutual information score does not directly indicate a feature's real importance. However, mutual information has proven to yield reasonable results and is very attractive regarding its computational complexity. Thus, features with a score below a chosen threshold are dropped, all others are distributed to the index sets \mathcal{W}_l , see (6), for the kernel functions κ_l in (5). All but (possibly) the last kernel are built from exactly 3 features. For simplicity, they are arranged following the MIS ranking, such that the 3 features with the highest scores are included in the first kernel K_1 and so on.

In most applications, using multiple kernels is motivated in combining different similarity measures or including samples coming from various sources. A prominent example from computational biology is combining protein sequences, gene expressions and protein structures to predict protein-protein interactions [34]. But associating kernels based on such heterogeneous sets of data is by no means trivial. Gönen et al. [13] give an excellent overview of common combination methods.

However, in our case, working with multiple kernels is due to the model construction. Therefore, we are in a completely different setting. The kernels neither rely on different similarity measures, nor do our data points come from various sources. Using an additional weighting scheme is not necessary here. Because of this, all kernels are equally weighted, so that their weights sum up to 1. This yields the combined kernel

$$K = \eta_1 K_1 + \dots + \eta_P K_P,$$

with $\eta_l = \frac{1}{P}$ and

$$K_l = \begin{bmatrix} \kappa_l(\mathbf{x}_1^{\mathcal{W}_l}, \mathbf{x}_1^{\mathcal{W}_l}) & \dots & \kappa_l(\mathbf{x}_1^{\mathcal{W}_l}, \mathbf{x}_N^{\mathcal{W}_l}) \\ \vdots & & \vdots \\ \kappa_l(\mathbf{x}_N^{\mathcal{W}_l}, \mathbf{x}_1^{\mathcal{W}_l}) & \dots & \kappa_l(\mathbf{x}_N^{\mathcal{W}_l}, \mathbf{x}_N^{\mathcal{W}_l}) \end{bmatrix}$$

for $l = 1, \dots, P$. Here, $P \in \mathbb{N}$ is the number of kernels to combine and \mathcal{W}_l are the considered index sets as defined in (6), so that $\mathbf{x}_i^{\mathcal{W}_l}$ and $\mathbf{x}_j^{\mathcal{W}_l}$ are the corresponding data points restricted to those indices, see (7). Thereby, as requested, all inputs $\mathbf{x}_i^{\mathcal{W}_l}$, $i = 1, \dots, N$, $l = 1, \dots, P - 1$, are 3-dimensional. The last index set \mathcal{W}_P contains all features, which are left, i. e. either 1, 2 or 3 features.

The multiplication with K is then performed using P NFFTs. Note that the applicability of the fast NFFT-approach is of course not limited to equal weights.

6. Numerical experiments. In this section, we demonstrate numerical results for the presented method on a couple of benchmark data sets. The corresponding Python implementations are available at <https://github.com/wagnertheresa/NFFT4ANOVA>. They rely on the FastAdjacency package¹ by Bünger, which was introduced in [2]. All experiments in this paper were run on a standard laptop computer

¹<https://github.com/dominikbuenger/FastAdjacency>

with $8 \times$ AMD Ryzen 7 4700U processors with Radeon Graphics and 15.0 GiB of RAM. I.e. the proposed method does not require expensive hardware or having access to supercomputers, what enables a wide and low-threshold application.

6.1. Data preprocessing. Operations such as computing the inner product or the norm obviously cannot be performed on every data type. Therefore, the data must be preprocessed in such a way that data points are free of missing values and categorical variables and solely consist of real values, before applying the classifier. In our implementations, we assume the data to already be of the required form. Therefore, (possibly) necessary preprocessing routines must be run individually beforehand.

For handling missing values, we suggest to replace missing entries by some valid number. For details on imputation, we refer to Donders et al. [8]. Categorical variables are usually incorporated by applying one of numerous encoding schemes, such as the bin-counting scheme or the feature hashing scheme [38].

While dealing with missing values and categorical variables is not covered in our implementations, handling imbalanced data sets is. Since the KRR problem (3) is targeted at maximizing the accuracy and minimizing the prediction-error, running the model on an imbalanced data set can have unwanted effects. Running the algorithm naively on a data set with one class being extremely over-represented might lead to a model, which assigns all future data to this class, no matter what the particular data points look like. In this case, the model did not discover anything during learning on the training data and all new data points from the under-represented class will be misclassified. However, these are precisely the cases, where classifying samples from the under-represented class correctly, i.e. achieving a high recall, is of utmost importance. Fraud detection is just one example for this. To ensure that this scenario never occurs and the model yields reasonable results, we use re-sampling techniques [19] to balance the class distribution of the data. Our model is targeted at large-scale data. Hence, for simplicity, we under-sample the over-represented class by randomly removing samples from the majority class, when a data set is imbalanced.

Another necessary aspect of data preprocessing is scaling [1, 21]. To make sure that our model does not perceive features with large scales as more relevant than features with small ranges, our implementation includes a function that z-score normalizes [29] the data set before starting any computations. It standardizes all columns, such that their mean is 0 and their standard deviation is 1. By this, the impact of outliers is reduced. To prevent train-test contamination, mean and standard deviation are computed solely with respect to the training data, so that the test data are scaled with the same statistics as the training data.

6.2. Accuracy and runtime of NFFT-based KRR learning. To demonstrate the power of the NFFT-based fast summation, we want to apply our method to kernel ridge regression learning now. As described in Chapter 2, computing the kernel-vector products within in CG algorithm, see (4), is most expensive within the learning task. We therefore want to replace the standard multiplication by the multiplication with NFFT-approach there to reduce the computational complexity.

For this, we consider a couple of benchmark data sets, see Table 1.

data set	N	d	number of data points in under-represented class	largest balanced sample possible
Telescope ²	19020	10	6688	13376
HIGGS ³	11000000	28	5170877	10341754
SUSY ⁴	5000000	18	2287827	4575654
YearPredictionMSD ⁵	515345	90	206945	413890
cod-rna ⁶	488565	8	162855	209884 ⁷

TABLE 1. Benchmark data sets for numerical experiments

In our experiments, we compare the performance of the proposed NFFT-based method with two state-of-the-art sklearn machine learning models, namely *sklearn.kernel_ridge.KernelRidge* and *sklearn.svm.SVC*. To the best of our knowledge both methods do not employ any tailored matrix routines such as sketching. In our computations, we fix the MIS-threshold to 0.0, which means we include all features, the FastAdjacency-setup to *default* and the tolerance of convergence in the CG algorithm to 1e-03. At this stage we did not observe a dramatic increase in the number of CG steps and so far have not considered any preconditioning. This might become necessary if more challenging data sets and parameter choices are made. The kernels are equally weighted and the data is balanced and z-score normalized. The results express mean values over 10 runs for the corresponding sample size, where the parameters are tuned using GridSearch for each and every run. For the considered classifiers, we choose the parameter-grids shown in Table 2.

classifier	kernel coefficient	regularization parameter
NFFTKernelRidge	$\sigma : [10^{-3}, 10^{-2}, 10^{-1}, 1, 10^1, 10^2, 10^3]$	$\lambda : [1, 10^1, 10^2, 10^3]$
sklearn KRR	$\gamma : [10^6, 10^4, 10^2, 1, 10^{-2}, 10^{-4}, 10^{-6}]$	$\alpha : [1, 10^1, 10^2, 10^3]$
sklearn SVC	$\gamma : [10^6, 10^4, 10^2, 1, 10^{-2}, 10^{-4}, 10^{-6}]$	$C : [1, 10^{-1}, 10^{-2}, 10^{-3}]$

TABLE 2. Parameter-grids for the considered classifiers

The respective parameter-grids for the distinct classifiers are equivalent, as the relations $\alpha = \frac{1}{C} = \lambda$ and $\gamma = \frac{1}{\sigma^2}$ hold. For large samples, only one run is performed, since the standard deviation is very small, when that many data points are included. Those samples are marked with * in the following tables. For all but the cod-rna data set, where we respect the train and test set distribution, the test set is obtained using random 50 percent of the sample. For the entire experiment, the train-test-split is 50:50. The objects of study are the accuracy and the runtimes for fitting and predicting within the learning methods.

²<https://www.kaggle.com/brsdincer/telescope-spectrum-gamma-or-hadron>

³<https://archive.ics.uci.edu/ml/datasets/HIGGS>, Whiteson et al. [4]

⁴<https://archive.ics.uci.edu/ml/datasets/SUSY>, Whiteson et al. [4]

⁵<https://archive.ics.uci.edu/ml/datasets/yearpredictionmsd>, transform from multiclass into binary class by dividing into classes “before 2000” and “as from 2000”

⁶<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>, Uzilov et al. [35]

⁷largest balanced sample possible $\neq 2 \cdot$ number of data points in underrepresented class, since we need to respect the training and test set distribution

In Table 3, the results for the Telescope Data Set are shown. We notice that the NFFT-approach’s runtime starts at quite a high level, especially in comparison with sklearn SVC. But with increasing sample size, our NFFTKernelRidge classifier overtakes the other ones, with the accuracy being very satisfying as we even achieve 0.1 and 0.3 percent more for the maximal balanced sample, respectively. While we beat the sklearn classifiers not until a sample size somewhat greater than 10000 in fitting time, we outrun them in prediction time with 5000 data points already. With this, the presented classifier is 1.85 or 1.65 times as fast as the sklearn classifiers for the largest balanced sample possible.

sample size	best accuracy			mean runtime fit			mean runtime predict			mean total runtime		
	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC
100	69.6	74.2	57.0	0.1091	0.0026	0.0008	0.0121	0.0011	0.0003	0.1212	0.0036	0.0011
500	80.4	78.1	76.6	0.1600	0.0075	0.0036	0.0168	0.0029	0.0021	0.1768	0.0104	0.0056
1000	81.3	79.7	78.3	0.2104	0.0121	0.0125	0.0225	0.0089	0.0082	0.2328	0.0210	0.0207
5000*	83.2	82.0	81.4	0.4207	0.2086	0.1694	0.0376	0.1138	0.1300	0.4583	0.3224	0.2994
10000*	83.9	83.8	83.7	0.8128	0.8929	0.7027	0.0675	0.3991	0.5210	0.8803	1.2929	1.2237
13376*	83.9	83.8	83.6	1.2122	1.7200	1.2932	0.0978	0.7063	0.8661	1.3100	2.4264	2.1594

TABLE 3. Telescope Data Set - Accuracy and runtime of GridSearch on various balanced samples of different size

Since the NFFT-based fast summation enables us to reduce the quadratic computational complexity, so that we nearly reach a linear scaling, this method is worth while even more for large data sets. Therefore, we want to consider the larger HIGGS Data Set next, see Table 4. As before, our method pays off with increasing sample size. Again, we overtake the competitors in prediction earlier than in fitting. But now, it takes larger samples to beat the sklearn classifiers. This is due to the HIGGS Data Set having 28 features, whereas the Telescope Data Set only had 10 features. And involving more features means having more windows, so that we must use more NFFTs for the approximation, which requires more time.

A huge drawback of sklearn KRR is that it does not work for large-scale samples at all. For 25000 training data points the kernel repeatedly dies and when starting it for 50000 or more training data points, the computations break, since the program is “unable to allocate 18.6 GiB for an array with shape (50000, 50000) and data type float64”. This is why some entries in the tables are missing. The same was experienced earlier, when computing kernel-vector products with the standard multiplication. In contrast, the proposed NFFT-approach runs kernel-vector multiplications for samples of this size absolutely trouble-free and in perfectly reasonable time. Hence, our method enables working with large-scale data, where the state-of-the-art methods completely fail without adding additional sparsity.

In our experiment, the largest sample size, where all classifiers are working, is 25000 data points, i.e. 12500 training and test data points each. Here, the NFFTKernelRidge classifier is 1.64 or 1.48 times as fast as the sklearn classifiers in fitting. In predicting, the proposed learning model is even 15.12 or 8.8 times as fast. However, for this sample size, our accuracy is 2.2 or 1.7 percent beneath. But with increasing sample size, the NFFTKernelRidge classifier achieves to catch up on the accuracy, so that it reaches 0.1 percent more than sklearn SVC for 200000 data points. In fitting, the NFFT-approach is 9.44 and in predicting even 80.92 times as fast.

sample size	best accuracy			mean runtime fit			mean runtime predict			mean total runtime		
	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC
100	55.8	54.4	40.4	0.2624	0.0020	0.0011	0.0344	0.0008	0.0006	0.2968	0.0028	0.0017
500	58.8	55.7	54.3	0.3561	0.0048	0.0031	0.0392	0.0020	0.0020	0.3953	0.0068	0.0051
1000	62.1	60.5	59.3	0.4325	0.0133	0.0107	0.0450	0.0040	0.0080	0.4776	0.0173	0.0187
5000*	63.3	62.9	62.0	1.0610	0.1914	0.2557	0.1056	0.1056	0.2084	1.1666	0.2970	0.4641
10000*	66.7	65.5	65.0	2.1512	0.8663	1.0339	0.1992	0.3673	0.7972	2.3504	1.2336	1.8312
25000*	64.8	67.0	66.5	6.3951	10.4802	9.4575	0.4987	7.5416	4.3869	6.8938	18.0218	13.8443
50000*	65.5	-	67.3	14.4947	-	56.6365	0.9397	-	20.4737	15.4344	-	77.1102
100000*	67.1	-	67.9	34.6103	-	197.9109	1.8479	-	84.8191	36.4582	-	282.7300
200000*	68.6	-	68.5	105.0926	-	991.7101	4.1666	-	337.1600	109.2592	-	1328.8701

TABLE 4. HIGGS Data Set - Accuracy and runtime of GridSearch on various balanced samples of different size

To confirm our observations once again, we run this experiment on 3 more benchmark data sets. For the SUSY Data Set, see Table 5, we trail the sklearn classifiers by 0.5 or 0.3 percent in accuracy, but are 5.31 or 4.78 times as fast in total runtime for a sample size of 50000. And when working with 200000 data points, the NFFT-based method is 0.4 percent worse than sklearn SVC in accuracy, but 16.7 times as fast.

sample size	best accuracy			mean runtime fit			mean runtime predict			mean total runtime		
	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC
100	62.2	62.0	55.6	0.1218	0.0009	0.0005	0.0159	0.0004	0.0001	0.1377	0.0014	0.0006
500	74.3	74.2	71.8	0.1754	0.0044	0.0028	0.0216	0.0029	0.0016	0.1969	0.0073	0.0044
1000	76.8	75.5	73.6	0.2427	0.0114	0.0091	0.0307	0.0039	0.0058	0.2734	0.0153	0.0149
5000*	77.6	77.3	76.9	0.8570	0.1879	0.2033	0.0775	0.0939	0.1614	0.9345	0.2817	0.3647
10000*	78.4	78.4	78.0	1.5382	1.0586	0.8214	0.1356	0.3487	0.6078	1.6737	1.4073	1.4291
50000*	78.9	79.4	79.2	9.1113	42.7143	31.1025	0.6305	9.0420	15.5091	9.7417	51.7563	46.6116
100000*	78.7	-	79.2	23.9190	-	148.1214	1.2657	-	62.3419	25.1847	-	210.4633
200000*	78.9	-	79.3	56.7895	-	732.4476	2.3603	-	255.1349	59.1497	-	987.5824

TABLE 5. SUSY Data Set - Accuracy and runtime of GridSearch on various balanced samples of different size

For the YearPredictionMSD Data Set, see Table 6, the accuracy yielded by our NFFTKernelRidge classifier is 2.4 percent worse than sklearn SVC's accuracy, when 413890 data points are included. Even though using the NFFT-method makes the learning task 15.9 times faster, the forfeitures in accuracy are not negligible. By construction, the NFFT-approach only covers relations between the features being in the same window. The sklearn classifiers involve all relations, instead. It is easily conceivable, that the NFFT-method might lack one or another feature-connection, so that it cannot always get at the sklearn classifiers' accuracy. But we observe that the NFFTKernelRidge classifier obtains great runtimes even for data with 90 features, i. e. with 30 windows. It highly depends on the exact data set and classification task, whether it is worth taking 16 times as long for 2 percent more accuracy.

For the cod-rna Data Set, see Table 7, NFFTKernelRidge's accuracy is only 0.5 percent worse than sklearn SVC's for 209884 data points, the largest balanced sample possible. This slight loss will be acceptable in most cases, when being 19.77 times as fast in return.

sample size	best accuracy			mean runtime fit			mean runtime predict			mean total runtime		
	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC
100	46.0	50.8	42.6	0.6578	0.0011	0.0007	0.0899	0.0005	0.0006	0.7476	0.0016	0.0013
500	61.4	64.2	64.2	0.8981	0.0051	0.0064	0.1122	0.0022	0.0054	1.0103	0.0074	0.0118
1000	66.0	69.0	68.2	1.1817	0.0107	0.0218	0.1405	0.0066	0.0196	1.3221	0.0173	0.0413
5000*	70.3	71.7	71.7	3.2502	0.1861	0.5164	0.3882	0.0966	0.4705	3.6384	0.2827	0.9869
10000*	72.3	73.4	73.1	6.7636	0.8400	2.1104	0.7016	0.3625	1.8575	7.4651	1.2025	3.9679
50000*	73.4	74.8	74.8	53.3774	43.2971	122.4035	3.9307	9.1279	50.3554	57.3081	52.4250	172.7589
100000*	74.2	-	75.8	102.0414	-	426.8239	6.1424	-	201.7097	108.1838	-	628.5336
200000*	74.4	-	76.4	245.1812	-	2006.6694	12.3315	-	785.3145	257.5127	-	2791.9839
413890*	74.3	-	76.7	669.4028	-	7696.3063	24.4887	-	3335.6310	693.8915	-	11031.9373

TABLE 6. YearPredictionMSD Data Set - Accuracy and runtime of GridSearch on various balanced samples of different size

sample size	best accuracy			mean runtime fit			mean runtime predict			mean total runtime		
	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC	NFFT KRR	sklearn KRR	SVC
100	78.0	77.2	72.2	0.0459	0.0010	0.0005	0.0066	0.0006	0.0002	0.0524	0.0016	0.0007
500	89.2	92.7	89.6	0.0671	0.0039	0.0021	0.0074	0.0023	0.0012	0.0745	0.0063	0.0034
1000	92.5	94.7	93.7	0.0882	0.0110	0.0071	0.0100	0.0037	0.0044	0.0982	0.0147	0.0116
5000*	95.4	95.7	95.6	0.3046	0.1788	0.1564	0.0274	0.0952	0.1031	0.3320	0.2740	0.2595
10000*	95.7	95.8	95.9	0.5864	0.8309	0.6044	0.0472	0.3723	0.4101	0.6336	1.2032	1.0145
50000*	95.7	95.9	96.0	3.9697	43.4380	17.9735	0.2376	9.5084	9.8435	4.2073	52.9464	27.8170
100000*	95.9	-	96.2	10.1743	-	73.8995	0.4686	-	39.0615	10.6430	-	112.9610
209884*	96.0	-	96.5	28.9638	-	422.5391	1.0058	-	170.0743	29.9696	-	592.6134

TABLE 7. cod-rna Data Set - Accuracy and runtime of GridSearch on various balanced samples of different size

In Figure 4, the results from Tables 3 to 7 are visualized. While the graphs representing the accuracy almost align for all classifiers on all data sets considered, the graphs showing the runtimes give a totally different picture. Again, it is clearly visible that the NFFTKernelRidge classifier avails to nothing for small data sets. But for large-scale data its runtime beats the competitive classifiers by several magnitudes. This applies for prediction in particular. Notice that the runtimes highly depend on the number of features, i.e. windows, involved and also on the choice of parameters.

7. Summary and outlook. We have studied the applicability of NFFT-based fast summation on kernel methods with a high-dimensional feature space. The ANOVA kernel has proved to be a viable tool to group the features into smaller pieces that are then amenable to the NFFT-based summation technique. We embedded this approach into a kernel ridge regression scheme that requires matrix-vector products as part of the CG algorithm to solve the linear system. We have illustrated the competitiveness of our method on several large-scale examples.

As a future topic we would like to examine the relationship between our method and the algorithms proposed by Potts and Schmischke [22, 23], which allow for a sensitivity analysis by studying truncated ANOVA decompositions of functions with low-dimensional structures in terms of the Fourier coefficients.

Additionally, we believe our scheme can be embedded into methods that rely on kernel matrices or even graph Laplacians. Examples are support vector machines or graph neural networks but also other loss functions that rely on kernel formulations and can then benefit from having the fast matrix-vector product available.

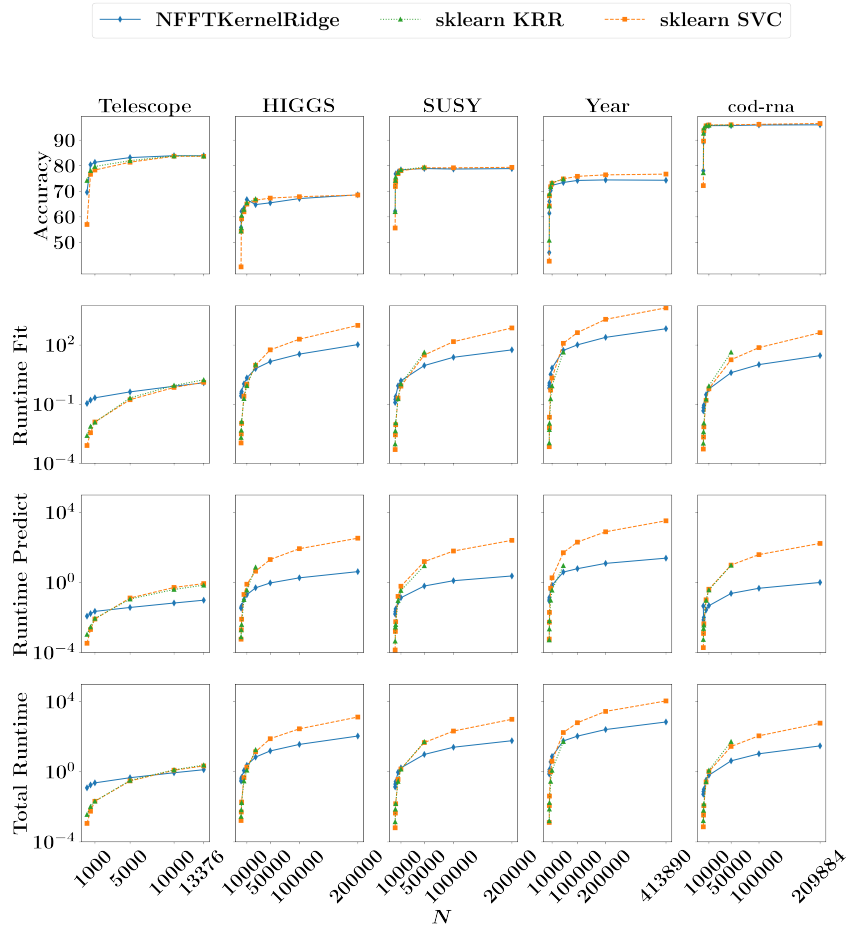


FIGURE 4. Accuracy and runtime of GridSearch on various balanced samples of different size

Acknowledgments. The authors gratefully acknowledge their support by the BMBF grant 01|S20053A (project SA ℓ E). We thank our colleagues in the project SA ℓ E for valuable discussions on the contents of this paper.

REFERENCES

- [1] A. Adeyemo, H. Wimmer and L. M. Powell, Effects of normalization techniques on logistic regression in data science, *Journal of Information Systems Applied Research*, **12** (2019), 37.
- [2] D. Alfke, D. Potts, M. Stoll and T. Volkmer, [NFFT meets krylov methods: Fast matrix-vector products for the graph Laplacian of fully connected networks](#), *Frontiers in Applied Mathematics and Statistics*, **4** (2018), 61.
- [3] H. Avron, K. L. Clarkson and D. P. Woodruff, [Faster kernel ridge regression using sketching and preconditioning](#), *SIAM J. Matrix Anal. Appl.*, **38** (2017), 1116–1138.
- [4] P. Baldi, P. Sadowski and D. Whiteson, [Searching for exotic particles in high-energy physics with deep learning](#), *Nature Communications*, **5** (2014), 1–9.
- [5] R. Battiti, [Using mutual information for selecting features in supervised neural net learning](#), *IEEE Transactions on Neural Networks*, **5** (1994), 537–550.

- [6] G. Beylkin, [On the Fast Fourier Transform of Functions with Singularities](#), *Applied Computational Harmonic Analysis*, **2** (1995), 363–381.
- [7] C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.
- [8] A. R. T. Donders, G. J. M. G. Van Der Heijden, T. Stijnen and K. G. M. Moons, [A gentle introduction to imputation of missing values](#), *Journal of Clinical Epidemiology*, **59** (2006), 1087–1091.
- [9] A. J. W. Duijndam and M. A. Schonewille, [Nonuniform fast Fourier transform](#), *GEO-PHYSICS*, **64** (1999), 539–551.
- [10] A. Dutt and V. Rokhlin, [Fast Fourier transforms for nonequispaced data](#), *SIAM J. Sci. Comput.*, **14** (1993), 1368–1393.
- [11] M. Fenn and G. Steidl, [Fast NFFT based summation of radial functions](#), *Sampling Theory in Signal and Image Processing*, **3** (2004), 1–28.
- [12] G. H. Golub and C. F. Van Loan, *Matrix Computations*, JHU Press, 2013.
- [13] M. Gönen and E. Alpaydm, [Multiple kernel learning algorithms](#), *J. Mach. Learn. Res.*, **12** (2011), 2211–2268.
- [14] M. R. Hestenes and E. Stiefel, [Methods of conjugate gradients for solving linear systems](#), *Journal of Research of the National Bureau of Standards*, **49** (1952), 409–436.
- [15] T. Hofmann, B. Schölkopf and A. J. Smola, [Kernel methods in machine learning](#). *Ann. Statist.*, **36** (2008), 1171–1220.
- [16] J. Keiner, S. Kunis and D. Potts, [Using NFFT3- a software library for various nonequispaced fast Fourier transforms](#), *ACM Trans. Math. Software*, **36** (2009), Article 19, 1–30.
- [17] T. N. Kipf and M. Welling, [Semi-supervised classification with graph convolutional networks](#), arXiv preprint, [arXiv:1609.02907](#), 2016.
- [18] W. March, B. Xiao and G. Biros, [ASKIT: Approximate skeletonization kernel-independent treecode in high dimensions](#), *SIAM J. Sci. Comput.*, **37** (2014), A1089–A1110.
- [19] A. I. Marqués, V. García and J. S. Sánchez, [On the suitability of resampling techniques for the class imbalance problem in credit scoring](#), *Journal of the Operational Research Society*, **64** (2013), 1060–1070.
- [20] V. I. Morariu, B. V. Srinivasan, V. C. Raykar, R. Duraiswami and L. S. Davis, [Automatic online tuning for fast Gaussian summation](#), *Advances in Neural Information Processing Systems*, **21** (2008).
- [21] S. G. K. Patro and K. K. Sahu, [Normalization: A preprocessing stage](#), *International Advanced Research Journal in Science, Engineering and Technology*, **2** (2015), 20–22. arXiv preprint, [arXiv:1503.06462](#), 2015.
- [22] D. Potts and M. Schmischke, [Approximations of high-dimensional periodic functions with Fourier-Based methods](#), *SIAM J. Numer. Anal.*, **59** (2021), 2393–2429.
- [23] D. Potts and M. Schmischke, [Learning multivariate functions with low-dimensional structures using polynomial bases](#), *J. Comput. Appl. Math.*, **403** (2022), 113821, 19 pp.
- [24] D. Potts and G. Steidl, [Fast summation at nonequispaced knots by NFFT](#), *SIAM J. Sci. Comput.*, **24** (2003), 2013–2037.
- [25] D. Potts, G. Steidl and A. Nieslony, [Fast convolution with radial kernels at nonequispaced knots](#), *Numer. Math.*, **98** (2004), 329–351.
- [26] C. E. Rasmussen, [Gaussian processes in machine learning](#), In *Summer School on Machine Learning*, Springer, 2003, 63–71.
- [27] V. C. Raykar and R. Duraiswami, [Fast large scale Gaussian process regression using approximate matrix-vector products](#), In *Learning Workshop*, 2007.
- [28] Y. Saad, *Iterative Methods for Sparse Linear Systems*, SIAM, 2003.
- [29] W. Sarle, [comp. ai. neural-nets FAQ, Part 2 of 7: Learning](#), <http://www.faqs.org/faqs/ai-faq/neural-nets/part2>, (1997), (accessed: 22 February 2021).
- [30] B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press, 2002.
- [31] J. Shawe-Taylor and N. Cristianini, *Kernel Methods for Pattern Analysis*, Cambridge University Press, 2004.
- [32] G. Steidl, [A note on fast Fourier transforms for nonequispaced grids](#), *Adv. Comput. Math.*, **9** (1998), 337–353.
- [33] M. Stoll, [A literature survey of matrix methods for data science](#), *GAMM-Mitt.*, **43** (2020), e202000013, 26 pp.
- [34] H. Tanabe, T. B. Ho, C. H. Nguyen and S. Kawasaki, [Simple but effective methods for combining kernels in computational biology](#), In *2008 IEEE International Conference on Research*,

- Innovation and Vision for the Future in Computing and Communication Technologies*, IEEE, 2008, 71–78.
- [35] A. V. Uzilov, J. M. Keegan and D. H. Mathews, [Detection of non-coding RNAs on the basis of predicted secondary structure formation free energy change](#), *BMC Bioinformatics*, **7** (2006), 1–30.
- [36] A. G. Wilson, Z. Hu, R. Salakhutdinov and E. P. Xing, Deep kernel learning, In *Artificial Intelligence and Statistics*, Proc. Mach. Learn. Res. (PMLR), 2016, 370–378.
- [37] C. Yu, W. March, B. Xiao and G. Biros, [INV-ASKIT: A parallel fast direct solver for kernel matrices](#), *2016 IEEE International Parallel and Distributed Processing Symposium*, 2016, 161–171.
- [38] A. Zheng and A. Casari, *Feature Engineering for Machine Learning: Principles and Techniques for Data Scientists*, O’Reilly Media, Inc., 2018.

Received November 2021; revised March 2022; early access July 2022.

E-mail address: franziska.nestler@math.tu-chemnitz.de

E-mail address: martin.stoll@math.tu-chemnitz.de

E-mail address: theresa.wagner@math.tu-chemnitz.de