

# Fast robust kernel regression through sign gradient descent with early stopping

Oskar Allerbo<sup>1</sup>

<sup>1</sup>*Department of Mathematics, KTH Royal Institute of Technology, e-mail: [osallerbo@kth.se](mailto:osallerbo@kth.se)*

**Abstract:** Kernel ridge regression, KRR, is a generalization of linear ridge regression that is non-linear in the data, but linear in the model parameters. Here, we introduce an equivalent formulation of the objective function of KRR, which opens up for replacing the ridge penalty with the  $\ell_\infty$  and  $\ell_1$  penalties.

Using the  $\ell_\infty$  and  $\ell_1$  penalties, we obtain robust and sparse kernel regression, respectively. We study the similarities between explicitly regularized kernel regression and the solutions obtained by early stopping of iterative gradient-based methods, where we connect  $\ell_\infty$  regularization to sign gradient descent,  $\ell_1$  regularization to forward stagewise regression (also known as coordinate descent), and  $\ell_2$  regularization to gradient descent, and, in the last case, theoretically bound for the differences. We exploit the close relations between  $\ell_\infty$  regularization and sign gradient descent, and between  $\ell_1$  regularization and coordinate descent to propose computationally efficient methods for robust and sparse kernel regression.

We finally compare robust kernel regression through sign gradient descent to existing methods for robust kernel regression on five real data sets, demonstrating that our method is one to two orders of magnitude faster, without compromised accuracy.

**MSC2020 subject classifications:** Primary 62G35; secondary 62J07.

**Keywords and phrases:** Robust regression, kernel regression, sign gradient descent, gradient descent, gradient flow.

Received September 2024.

## 1. Introduction

Kernel ridge regression, KRR, is a generalization of linear ridge regression that is non-linear in the data, but linear in the model parameters. As for linear ridge regression, KRR has a closed-form solution, but at the cost of solving a system of  $n$  linear equations, where  $n$  is the number of training observations. The KRR estimate coincides with the posterior mean of kriging, or Gaussian process regression, (Krige, 1951; Matheron, 1963) and has successfully been applied within a wide range of applications (Zahrt et al., 2019; Ali et al., 2020; Chen and Leclair, 2021; Fan et al., 2021; Le et al., 2021; Safari and Rahimzadeh Arashloo, 2021; Shahsavari et al., 2021; Singh Alvarado et al., 2021; Wu et al., 2021; Chen et al., 2022).

Robust regression is often implemented by replacing the standard  $\ell_2$  loss function with a loss function that is less sensitive to outliers in the data. Robust methods for kernel regression tend to rely either on M-estimators (De Brabanter et al., 2009; Wibowo, 2009; Debruyne et al., 2010; Hwang et al., 2015) or quantile regression (Hwang and Shim, 2005; Takeuchi et al., 2006; Li et al., 2007). In contrast to KRR, these methods are all iterative and thus tend to be computationally heavy, even though more computationally efficient methods for kernel quantile regression have recently been proposed by Zheng (2022) and Tang et al. (2024).

In the linear case, many alternatives to the ridge penalty have been proposed, including the lasso penalty (Tibshirani, 1996), which is known for creating sparse models. By replacing the ridge penalty of KRR with the lasso penalty Roth (2004), Guigue et al. (2005) and Feng et al. (2016) have obtained non-linear regression that is sparse in the observations. Rather than using the lasso, or  $\ell_1$ , penalty, Russu et al. (2016) and Demontis et al. (2017) trained kernelized support vector machines with max (or infinity,  $\ell_\infty$ ) norm regularization to obtain models that are robust against adversarial data, which in the context of regression translates to outliers.

An alternative to explicit regularization is to use an iterative optimization algorithm and to stop the training before the algorithm converges, something that is known as early stopping. A well-known example of this is the connection between the explicitly regularized lasso model and the iterative method forward stagewise regression, which is also known as coordinate descent. The similarities between the solutions are well studied by e.g. Efron et al. (2004), Hastie et al. (2007) and Tibshirani (2015). There are also striking similarities between explicitly regularized ridge regression and gradient descent with early stopping (Friedman and Popescu, 2004; Ali et al., 2019). Replacing explicit infinity norm regularization with an iterative optimization method with early stopping does not seem to be as well studied. However, as is shown in this paper, the solutions are similar to those of sign gradient descent with early stopping.

One benefit of replacing explicit regularization with early stopping is that the entire solution path, which consists of all different levels of regularizations between 0 and  $\infty$ , is obtained by running the iterative optimization algorithm to convergence only once. In contrast, for explicit regularization, the solution has to be calculated once for every considered regularization strength. In general, no closed-form solutions exist for the lasso and infinity norm regularized problems, and iterative optimization methods have to be run to convergence once for every candidate value of the regularization strength, something that tends to be computationally heavy.

In this paper, we present an equivalent formulation of KRR, which we use to obtain kernel regression with the  $\ell_\infty$  and  $\ell_1$  norms, in addition to the default,  $\ell_2$ , norm. We also use the equivalent formulation to solve kernel regression using the three different gradient-based optimization algorithms gradient descent, sign gradient descent, and coordinate descent (forward stagewise regression). We relate each iterative algorithm to one explicitly regularized solution and use gradient-based optimization with early stopping to obtain computationally effi-

cient robust, and sparse, kernel regression. Even though robust kernel regression has more evident applications than sparse kernel regression does, since the calculations are analogous for  $\ell_\infty$  and  $\ell_1$  regularization we include both cases. In the results section, we however focus on robust regression.

The rest of the paper is structured as follows. In Section 2, we review kernel ridge regression, before introducing our equivalent formulation in Section 3. In Section 4, we review kernel gradient descent, KGD, and present theoretical comparisons between KGD with infinitesimal step size and KRR. In Section 5, we generalize KRR into robust and sparse kernel regression by replacing the  $\ell_2$  penalty with the  $\ell_\infty$  and  $\ell_1$  penalties, respectively. Similarly to in Section 4, we also introduce gradient-based optimization algorithms that we relate to the explicitly regularized methods; kernel sign gradient descent for robust kernel regression, and kernel coordinate descent for sparse kernel regression. Finally, in Section 6, we demonstrate our findings with experiments on real and synthetic data.

Our main contributions are listed below.

- We present an equivalent objective function for KRR and use this formulation to generalize KRR to the  $\ell_\infty$  and  $\ell_1$  penalties.
- We theoretically show that kernel sign gradient descent, and kernel coordinate descent, with early stopping, correspond to robust and sparse kernel regression, respectively. We use this to introduce computationally efficient regularization by replacing explicit  $\ell_\infty$  and  $\ell_1$  regularization by implicit regularization through early stopping.
- We demonstrate on five real data sets that robust kernel regression through kernel sign gradient descent is one to two orders of magnitude faster than existing robust kernel regression methods, without compromised performance in terms of prediction.

All proofs are deferred to Appendix D.

## 2. Review of kernel ridge regression

For a positive semi-definite kernel function,  $k(\mathbf{x}, \mathbf{x}') \in \mathbb{R}$ , and  $n$  paired observations,  $(\mathbf{x}_i, y_i)_{i=1}^n \in \mathbb{R}^p \times \mathbb{R}$ , presented in a design matrix,  $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times p}$ , and a response vector,  $\mathbf{y} \in \mathbb{R}^n$ , and for a given regularization strength,  $\lambda \geq 0$ , the objective function of kernel ridge regression, KRR, is given by

$$\hat{\boldsymbol{\alpha}} = \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{K}\boldsymbol{\alpha}\|_2^2 + \frac{\lambda}{2} \|\boldsymbol{\alpha}\|_{\mathbf{K}}^2 \quad (1)$$

with predictions given by

$$\begin{bmatrix} \hat{\mathbf{f}} \\ \hat{\mathbf{f}}^* \end{bmatrix} = \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} \hat{\boldsymbol{\alpha}}.$$

Here,  $\hat{\mathbf{f}} = \hat{\mathbf{f}}(\mathbf{X}) \in \mathbb{R}^n$  and  $\hat{\mathbf{f}}^* = \hat{\mathbf{f}}^*(\mathbf{X}^*, \mathbf{X}) \in \mathbb{R}^{n^*}$  denote model predictions for the training data,  $\mathbf{X}$ , and new data,  $\mathbf{X}^* = [\mathbf{x}_1^*, \mathbf{x}_2^*, \dots, \mathbf{x}_{n^*}^*]^\top \in \mathbb{R}^{n^* \times p}$ , and  $\mathbf{K} =$

$\mathbf{K}(\mathbf{X}) \in \mathbb{R}^{n \times n}$  and  $\mathbf{K}^* = \mathbf{K}^*(\mathbf{X}^*, \mathbf{X}) \in \mathbb{R}^{n^* \times n}$  denote two kernel matrices defined according to  $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$  and  $\mathbf{K}_{ij}^* = k(\mathbf{x}_i^*, \mathbf{x}_j)$ . The weighted norm,  $\|\mathbf{v}\|_{\mathbf{A}}$ , is defined according to  $\|\mathbf{v}\|_{\mathbf{A}}^2 = \mathbf{v}^\top \mathbf{A} \mathbf{v}$  for any symmetric positive definite matrix  $\mathbf{A}$ .

The closed-form solution for  $\hat{\boldsymbol{\alpha}}$  is given by

$$\hat{\boldsymbol{\alpha}} = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{y}, \quad (2)$$

and consequently

$$\begin{bmatrix} \hat{f} \\ \hat{f}^* \end{bmatrix} = \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{y}. \quad (3)$$

An alternative interpretation of KRR is as linear regression for a non-linear feature expansion of  $\mathbf{X}$ . According to Mercer's Theorem (Mercer, 1909), every kernel can be written as the inner product of feature expansions of its two arguments:  $k(\mathbf{x}, \mathbf{x}') = \boldsymbol{\varphi}(\mathbf{x})^\top \boldsymbol{\varphi}(\mathbf{x}')$  for  $\boldsymbol{\varphi}(\mathbf{x}) \in \mathbb{R}^q$ . Thus, denoting the feature expansions of the design matrix and the new data with  $\boldsymbol{\Phi} = \boldsymbol{\Phi}(\mathbf{X}) \in \mathbb{R}^{n \times q}$  and  $\boldsymbol{\Phi}^* = \boldsymbol{\Phi}^*(\mathbf{X}^*) \in \mathbb{R}^{n^* \times q}$ , the two kernel matrices can be expressed as  $\mathbf{K} = \boldsymbol{\Phi} \boldsymbol{\Phi}^\top$  and  $\mathbf{K}^* = \boldsymbol{\Phi}^* \boldsymbol{\Phi}^\top$ . For  $\boldsymbol{\beta} = \boldsymbol{\Phi}^\top \boldsymbol{\alpha}$ , Equations 1 and 3 become

$$\begin{aligned} \hat{\boldsymbol{\beta}} &= \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^q} \frac{1}{2} \|\mathbf{y} - \boldsymbol{\Phi} \boldsymbol{\beta}\|_2^2 + \frac{\lambda}{2} \|\boldsymbol{\beta}\|_2^2 \\ \begin{bmatrix} \hat{f} \\ \hat{f}^* \end{bmatrix} &= \begin{bmatrix} \boldsymbol{\Phi} \\ \boldsymbol{\Phi}^* \end{bmatrix} \hat{\boldsymbol{\beta}}, \end{aligned} \quad (4)$$

which is exactly linear ridge regression for the feature expansion of the kernel.

### 3. Equivalent formulations of kernel ridge regression

In this section, we present an equivalent formulation of the objective function in Equation 1, which provides the same solution for  $\hat{\boldsymbol{\alpha}}$ . This formulation opens up for generalizing KRR by using penalties other than the ridge penalty and also provides an interesting connection to functional gradient descent (Mason et al., 1999).

In Appendix A, we take this one step further by reformulating Equation 1 directly in the model predictions,  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$ , by presenting two equivalent objectives, that, when minimized with respect to  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$ , generate the solution in Equation 3. In this and the following sections, all calculations are done with respect to  $\boldsymbol{\alpha}$ , and the corresponding expressions for  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$  are obtained through multiplication with  $[\mathbf{K}^\top, \mathbf{K}^{*\top}]^\top$ , but, as shown in the appendix, the expressions for  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$  can be obtained directly without taking the detour over  $\boldsymbol{\alpha}$ .

In Proposition 3.1, we show how we can move the weighted norm in Equation 1 from the penalty term to the reconstruction term.

**Proposition 3.1.**

$$\hat{\alpha} = \arg \min_{\alpha \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{K}\alpha\|_2^2 + \frac{\lambda}{2} \|\alpha\|_{\mathbf{K}}^2 \quad (5a)$$

$$\begin{aligned} &= \arg \min_{\alpha \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{K}\alpha\|_{\mathbf{K}^{-1}}^2 + \frac{\lambda}{2} \|\alpha\|_2^2 \\ &= (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{y}. \end{aligned} \quad (5b)$$

*Remark 3.1.* Despite the appearance of  $\mathbf{K}^{-1}$  in Equation 5b, this inverse never needs to be calculated, since in all calculations,  $\mathbf{K}^{-1}$  cancels through multiplication by  $\mathbf{K}$ .

The alternative formulation in Equation 5b, where the reconstruction term, rather than the regularization term, is weighted by the kernel matrix, has two interesting implications. First, a standard ridge penalty on the regularization term opens up for using other penalties than the  $\ell_2$  norm, such as the  $\ell_1$  and  $\ell_\infty$  norms. Although these penalties have previously been used in the kernel setting, it has been done by replacing  $\|\alpha\|_{\mathbf{K}}^2$  in Equation 1 by  $\|\alpha\|_1$  or  $\|\alpha\|_\infty$ , hence acting as if the objective of KRR were to minimize  $\|\mathbf{y} - \mathbf{K}\alpha\|_2^2 + \lambda \|\alpha\|_2^2$ , and losing the connection to linear regression in feature space.

Second, the gradient of Equation 5b with respect to  $\alpha$  is

$$\mathbf{K}\alpha - \mathbf{y} + \lambda\alpha. \quad (6)$$

Multiplying by  $[\mathbf{K}^\top, \mathbf{K}^{*\top}]^\top$ , we obtain

$$\begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\mathbf{f} - \mathbf{y}) + \lambda \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix}, \quad (7)$$

which is the gradient used in functional gradient descent. Functional gradient descent is generally derived by differentiating functionals. Here, however, the expression is a simple consequence of the equivalent objective function of KRR.

#### 4. Kernel gradient descent and kernel gradient flow

Before replacing the Euclidean norm in Equation 5b with the  $\ell_\infty$  and  $\ell_1$  norms in Section 5, we investigate solving kernel regression iteratively with gradient descent. We also use gradient descent with infinitesimal step size, known as gradient flow, to obtain a closed-form solution which we use for direct comparisons to kernel ridge regression.

The similarities between ridge regression and gradient descent with early stopping are well studied for linear regression (Friedman and Popescu, 2004; Ali et al., 2019; Allerbo et al., 2023). When starting at zero, optimization time can be thought of as an inverse penalty, where longer optimization time corresponds to weaker regularization. When applying gradient descent to kernel regression, something we refer to as kernel gradient descent, KGD, we replace explicit

regularization with implicit regularization through early stopping. That is, we use  $\lambda = 0$  and consider training time,  $t$ , as a regularizer.

With  $\lambda = 0$  in Equation 6, starting at  $\mathbf{0}$ , the KGD update becomes, for step size  $\eta$ ,

$$\hat{\alpha}_{k+1} = \hat{\alpha}_k + \eta \cdot (\mathbf{y} - \mathbf{K}\hat{\alpha}_k), \quad \hat{\alpha}_0 = \mathbf{0}. \quad (8)$$

To compare the regularization injected by early stopping to that of ridge regression, we let the optimization step size go to zero to obtain a closed-form solution, which we refer to as kernel gradient flow, KGF. Then, Equation 8 can be thought of as the Euler forward formulation of the differential equation in Equation 9,

$$\frac{d\hat{\alpha}(t)}{dt} = \mathbf{y} - \mathbf{K}\hat{\alpha}(t), \quad \hat{\alpha}(0) = \mathbf{0}, \quad (9)$$

the solution to which is given by Equation 10.

$$\hat{\alpha}_{\text{KGF}}(t) = (\mathbf{I} - \exp(-t\mathbf{K}))\mathbf{K}^{-1}\mathbf{y}, \quad (10)$$

where  $\exp$  denotes the matrix exponential.

*Remark 4.1.* Note that  $(\mathbf{I} - \exp(-t\mathbf{K}))\mathbf{K}^{-1} = \mathbf{K}^{-1}(\mathbf{I} - \exp(-t\mathbf{K}))$  is well-defined even if  $\mathbf{K}$  is singular. The matrix exponential is defined through its Taylor expansion and from  $\mathbf{I} - \exp(-t\mathbf{K}) = t\mathbf{K} - \frac{1}{2!}t^2\mathbf{K}^2 + \dots$ , a matrix  $\mathbf{K}$  factors out, that cancels  $\mathbf{K}^{-1}$ .

*Remark 4.2.* It is possible to generalize KGF to Nesterov accelerated gradient descent with momentum (Nesterov, 1983; Polyak, 1964). In this case  $\exp(-t\mathbf{K})$  in Equation 10 generalizes to  $\exp\left(-\frac{t}{1-\gamma}\mathbf{K}\right)$ , where  $\gamma \in [0, 1)$  is the strength of the momentum. See Appendix D for details.

To facilitate the comparisons between KGF and KRR, we rewrite Equation 2 as

$$\hat{\alpha}_{\text{KRR}}(\lambda) = \left(\mathbf{I} - (\mathbf{I} + 1/\lambda \cdot \mathbf{K})^{-1}\right)\mathbf{K}^{-1}\mathbf{y}, \quad (11)$$

where we have used that

$$(\mathbf{K} + \lambda\mathbf{I})^{-1} = \left(\mathbf{I} - (\mathbf{I} + 1/\lambda \cdot \mathbf{K})^{-1}\right)\mathbf{K}^{-1},$$

which is a consequence of the Woodbury matrix identity.

Since  $\exp(-t\mathbf{K}) = \exp(t\mathbf{K})^{-1}$ , the KGF and KRR solutions differ only in the factor

$$\exp(t\mathbf{K}) \text{ v.s. } \mathbf{I} + 1/\lambda\mathbf{K}.$$

Thus, for  $t = 1/\lambda$ , we can think of the ridge penalty as a first-order Taylor approximation of the gradient flow penalty.

#### 4.1. Comparisons between kernel ridge regression and kernel gradient flow

In this section, we compare the KRR and KGF solutions for  $\lambda = 1/t$ . To do this, we introduce the following notation, where  $\mathbf{k}(\mathbf{x}^*)^\top = \mathbf{k}(\mathbf{x}^*, \mathbf{X})^\top \in \mathbb{R}^n$  is

the row in  $\mathbf{K}^*$  corresponding to  $\mathbf{x}^*$ :

$$\begin{aligned}\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t) &:= \mathbf{K} \hat{\boldsymbol{\alpha}}_{\text{KGF}}(t), & \hat{f}_{\text{KGF}}(\mathbf{x}^*, t) &:= \mathbf{k}(\mathbf{x}^*)^\top \hat{\boldsymbol{\alpha}}_{\text{KGF}}(t), \\ \hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, \lambda) &:= \mathbf{K} \hat{\boldsymbol{\alpha}}_{\text{KRR}}(\lambda), & \hat{f}_{\text{KRR}}(\mathbf{x}^*, \lambda) &:= \mathbf{k}(\mathbf{x}^*)^\top \hat{\boldsymbol{\alpha}}_{\text{KRR}}(\lambda), \\ \hat{\boldsymbol{\alpha}}^0 &:= \hat{\boldsymbol{\alpha}}_{\text{KGF}}(t = \infty) = \hat{\boldsymbol{\alpha}}_{\text{KRR}}(\lambda = 0) = \mathbf{K}^{-1} \mathbf{y}, \\ \hat{\mathbf{f}}^0(\mathbf{X}) &:= \hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t = \infty) = \hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, \lambda = 0) = \mathbf{K} \hat{\boldsymbol{\alpha}}^0 = \mathbf{y}, \\ \hat{f}^0(\mathbf{x}^*) &:= \hat{f}_{\text{KGF}}(\mathbf{x}^*, t = \infty) = \hat{f}_{\text{KRR}}(\mathbf{x}^*, \lambda = 0) = \mathbf{k}(\mathbf{x}^*)^\top \hat{\boldsymbol{\alpha}}^0 = \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \mathbf{y}, \\ \hat{\alpha}_i^{\text{KGF}} &:= (\hat{\boldsymbol{\alpha}}_{\text{KGF}})_i, & \hat{\alpha}_i^{\text{KRR}} &:= (\hat{\boldsymbol{\alpha}}_{\text{KRR}})_i, & \hat{\alpha}_i^0 &:= (\hat{\boldsymbol{\alpha}}^0)_i.\end{aligned}$$

Even though not explicitly stated, all estimates depend on the training data  $(\mathbf{X}, \mathbf{y})$ . We further use the notation  $\boldsymbol{\xi} \sim (\boldsymbol{\mu}, \boldsymbol{\Sigma})$  to denote that the random variable  $\boldsymbol{\xi}$  follows a distribution with  $\mathbb{E}(\boldsymbol{\xi}) = \boldsymbol{\mu}$  and  $\text{Cov}(\boldsymbol{\xi}) = \boldsymbol{\Sigma}$ .

In Proposition 4.1, we bound the differences between the KGF and KRR solutions in terms of the non-regularized solutions,  $\boldsymbol{\alpha}_0$ ,  $\hat{\mathbf{f}}^0(\mathbf{X})$  and  $\hat{f}^0(\mathbf{x}^*)$ . In parts (a) and (b), where we bound the differences between the parameter and in-sample prediction vectors, no further assumptions are needed. In parts (c) and (d), with bounds on individual parameters and predictions, including out-of-sample predictions, some very reasonable assumptions have to be made on the data. For all four bounds, the larger the non-regularized value is, the larger the difference between the KGF and KRR estimates is allowed to be.

**Proposition 4.1.** For  $t \geq 0$ ,  $\mathbf{y} \in \mathbb{R}^n$ ,

$$\begin{aligned}(a) \quad & \|\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t) - \hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t)\|_2^2 \leq 0.0415 \cdot \|\hat{\boldsymbol{\alpha}}^0\|_2^2, \\ (b) \quad & \left\| \hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t) - \hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, 1/t) \right\|_2^2 \leq 0.0415 \cdot \|\hat{\mathbf{f}}^0(\mathbf{X})\|_2^2 = 0.0415 \cdot \|\mathbf{y}\|_2^2.\end{aligned}$$

For  $t \geq 0$ ,  $\mathbf{y} = \mathbf{K} \boldsymbol{\alpha}_0 + \boldsymbol{\varepsilon}$ ,  $\boldsymbol{\alpha}_0 \sim (\mathbf{0}, \boldsymbol{\Sigma}_\alpha)$ ,  $\boldsymbol{\varepsilon} \sim (\mathbf{0}, \sigma_\varepsilon^2 \mathbf{I})$ , where  $\boldsymbol{\Sigma}_\alpha$  and  $\mathbf{K}$  are simultaneously diagonalizable,

$$\begin{aligned}(c) \quad & \mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \hat{\alpha}_i^{\text{KGF}}(t) - \hat{\alpha}_i^{\text{KRR}}(1/t) \right)^2 \right) \leq 0.0415 \cdot \mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( (\hat{\alpha}_i^0)^2 \right), \quad i = 1, 2, \dots, n, \\ (d) \quad & \mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \hat{f}_{\text{KGF}}(\mathbf{x}^*, t) - \hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t) \right)^2 \right) \leq 0.0415 \cdot \mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \hat{f}^0(\mathbf{x}^*)^2 \right).\end{aligned}$$

Two typical options for  $\boldsymbol{\Sigma}_\alpha$  are  $\sigma_\alpha^2 \mathbf{I}$  and  $\sigma_\beta^2 \mathbf{K}^{-1}$ , where the second formulation implies  $\boldsymbol{\beta}_0 \sim (\mathbf{0}, \sigma_\beta^2 \mathbf{I})$  in the feature space formulation of KRR from Equation 4.

In Proposition 4.2, we compare the distances to the observation vector,  $\mathbf{y}$ , of the KGF and KRR solutions. According to parts (a) and (c), for a given regularization, the furthest possible (among all values of  $\mathbf{y}$ , or expected  $y_i^2$ ) normalized KGF solution lies closer to the observations than the furthest possible normalized KRR solution does. Analogously, according to parts (b) and (d), for a given regularization, the furthest possible KRR solution lies closer to zero than the furthest possible KGF solution does. These results agree with those of Richards et al. (2021), who state that, for linear regression, when the signal-to-noise ratio is high, gradient descent tends to outperform ridge regression: For low noise, a solution that lies closer to the observations, as provided by KGF, is beneficial.

**Proposition 4.2.** For  $t \geq 0$ ,  $\mathbf{y} \in \mathbb{R}^n$ ,

$$(a) \quad \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t) - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) \leq \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, 1/t) - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right),$$

$$(b) \quad \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, 1/t)\|_2}{\|\mathbf{y}\|_2} \right) \leq \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t)\|_2}{\|\mathbf{y}\|_2} \right).$$

For  $t \geq 0$ ,  $\mathbf{y} = \mathbf{K}\boldsymbol{\alpha}_0 + \boldsymbol{\varepsilon}$ ,  $\boldsymbol{\alpha}_0 \sim (\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\alpha}})$ ,  $\boldsymbol{\varepsilon} \sim (\mathbf{0}, \sigma_{\varepsilon}^2 \mathbf{I})$ , where  $\boldsymbol{\Sigma}_{\boldsymbol{\alpha}}$  and  $\mathbf{K}$  are simultaneously diagonalizable, and  $i = 1, 2, \dots, n$ ,

$$(c) \quad \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \hat{f}_{\text{KGF}}(\mathbf{x}_i, t) - y_i \right)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2)} \right)$$

$$\leq \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \hat{f}_{\text{KRR}}(\mathbf{x}_i, 1/t) - y_i \right)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2)} \right),$$

$$(d) \quad \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \hat{f}_{\text{KRR}}(\mathbf{x}_i, 1/t)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2)} \right)$$

$$\leq \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \hat{f}_{\text{KGF}}(\mathbf{x}_i, t)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2)} \right).$$

Finally, if we assume that a true function exists, parameterized by the true  $\boldsymbol{\alpha}_0$  as

$$\begin{bmatrix} \mathbf{f}_0(\mathbf{X}) \\ \mathbf{f}_0(\mathbf{X}^*) \end{bmatrix} = \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} \boldsymbol{\alpha}_0,$$

and with observations according to  $\mathbf{y} = \mathbf{f}_0(\mathbf{X}) + \boldsymbol{\varepsilon}$ , we can calculate the expected squared differences between the estimated and true models, something that is often referred to as the risk:  $\text{Risk}(\hat{\boldsymbol{\theta}}; \boldsymbol{\theta}_0) = \mathbb{E} \left( \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2^2 \right)$ . In Proposition 4.3, which is an adaptation of Theorems 1 and 2 by Ali et al. (2019), the KGF estimation and prediction risks are bounded in terms of the corresponding KRR risks. In all three cases, the KGF risk is less than 1.69 times the KRR risk.

**Proposition 4.3.** For  $t \geq 0$ ,  $\mathbf{y} = \mathbf{f}_0(\mathbf{X}) + \boldsymbol{\varepsilon} = \mathbf{K}\boldsymbol{\alpha}_0 + \boldsymbol{\varepsilon}$ ,  $\boldsymbol{\varepsilon} \sim (\mathbf{0}, \sigma_{\varepsilon}^2 \mathbf{I})$ ,

- (a) For any  $\boldsymbol{\alpha}_0$ ,  $\text{Risk}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t); \boldsymbol{\alpha}_0) \leq 1.6862 \cdot \text{Risk}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t); \boldsymbol{\alpha}_0)$ .
- (b) For any  $\boldsymbol{\alpha}_0$ ,

$$\text{Risk}(\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t); \mathbf{f}_0(\mathbf{X})) \leq 1.6862 \cdot \text{Risk}(\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, 1/t); \mathbf{f}_0(\mathbf{X})).$$



(c) For  $\alpha_0 \sim (\mathbf{0}, \Sigma_\alpha)$ , where  $\Sigma_\alpha$  and  $\mathbf{K}$  are simultaneously diagonalizable,

$$\begin{aligned} & \mathbb{E}_{\alpha_0} \left( \text{Risk}(\hat{f}_{\text{KGF}}(\mathbf{x}^*, t); f_0(\mathbf{x}^*)) \right) \\ & \leq 1.6862 \cdot \mathbb{E}_{\alpha_0} \left( \text{Risk}(\hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t); f_0(\mathbf{x}^*)) \right), \text{ where } f_0(\mathbf{x}^*) \in \mathbf{f}_0(\mathbf{X}^*). \end{aligned}$$

To summarize this section, the KGF and KRR solutions tend to agree well both in terms of the parameter and prediction vectors and in terms of the risks. Which algorithm actually performs better, depends on the specific data, sometimes KGF performs slightly better than KRR and sometimes vice versa. Proposition 4.2 suggests that the KGF solution lies closer to the observations than KRR solution does, which can often be an advantage, however not in the presence of extreme outliers. These conclusions are supported by the experiments in Section 6.

## 5. Kernel regression with the $\ell_\infty$ and $\ell_1$ norms

In this section, we replace the squared  $\ell_2$  norm of KRR in Equation 5b by the  $\ell_\infty$  and  $\ell_1$  norms, respectively, to obtain  $\ell_\infty$  and  $\ell_1$  regularized kernel regression, which we abbreviate as  $\text{K}\ell_\infty\text{R}$  and  $\text{K}\ell_1\text{R}$ , respectively. We also connect the explicitly regularized algorithms  $\text{K}\ell_\infty\text{R}$  and  $\text{K}\ell_1\text{R}$  to kernel sign gradient descent, KSGD, and kernel coordinate descent, KCD, with early stopping, similarly to how we related KRR and KGF in Section 4. The six algorithms and their abbreviations are stated in Table 1.

The objective functions for  $\text{K}\ell_\infty\text{R}$  and  $\text{K}\ell_1\text{R}$  are, respectively

$$\arg \min_{\alpha \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{K}\alpha\|_{\mathbf{K}^{-1}}^2 + \lambda \|\alpha\|_\infty \quad (12)$$

and

$$\arg \min_{\alpha \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{K}\alpha\|_{\mathbf{K}^{-1}}^2 + \lambda \|\alpha\|_1. \quad (13)$$

In contrast to KRR, unless the data is uncorrelated, no closed-form solutions exist for Equations 12 and 13. However, the problems are still convex (strictly convex if  $\mathbf{K}$  is strictly positive definite, which is the case for all kernels with infinite-dimensional feature expansions) and solutions can be obtained using the

TABLE 1  
Kernel regression algorithms.

Algorithm	Abbreviation
Kernel ridge regression ( $\ell_2$ penalty)	KRR
Kernel regression with the $\ell_\infty$ penalty (robust)	$\text{K}\ell_\infty\text{R}$
Kernel regression with the $\ell_1$ penalty (sparse)	$\text{K}\ell_1\text{R}$
Kernel gradient descent	KGD
Kernel sign gradient descent (robust)	KSGD
Kernel coordinate descent (sparse)	KCD

iterative optimization method proximal gradient descent (Rockafellar, 1976), which, in contrast to standard gradient descent, can handle the discontinuities of the gradients of the  $\ell_\infty$  and  $\ell_1$  norms.

Applying  $\ell_\infty$  regularization, which penalizes the largest parameter(s), promotes a solution with no extreme parameter values. Thus, Equation 12 promotes a solution where the impact of extreme observations is alleviated, which might be beneficial if outliers are present in the data. Vice versa,  $\ell_1$ , or lasso, regularization promotes a sparse solution, where parameters are added sequentially to the model, with the most significant parameters included first. We further demonstrate this in Section 6.1.

### 5.1. Regularization through early stopping

The regularization strength must be carefully selected for the model to generalize well on new data. While a too large regularization results in poor performance due to a too simple model, a too small regularization might result in a too complex model that incorporates the noise in the training data, and thus generalizes poorly to new data (overfitting). This is sometimes referred to as the bias-variance tradeoff. In practice, the regularization strength is usually selected by some sort of cross-validation, where parts of the data are kept aside during training and then used for validation once the model is trained. By training multiple models, all with different regularization strengths, one can select the model that performs best on the validation data.

However, applying the proximal gradient descent algorithm might be computationally heavy, especially when evaluating several different regularization strengths. For some classes of models, there exist computationally efficient versions of leave-one-out and generalized cross-validation, but, as demonstrated in Appendix B, these do not perform well in this setting. Instead, we use another approach: In Section 4, we showed how the solution of KGF with early stopping is very similar to that of KRR, with a later stopping time,  $t$ , corresponding to a smaller regularization strength,  $\lambda$ . Thus, running KGD until convergence once, all different stopping times, or regularization strengths, between  $t = 0$  (corresponding to  $\lambda = \infty$ ) and  $t = t_{\max}$  (corresponding to  $\lambda = 0$ ) are obtained through this single execution of the algorithm. (For more detailed descriptions of the connections between explicit regularization and early stopping, see the works by e.g. Friedman and Popescu (2004), Tibshirani (2015) or Allerbo et al. (2023).) Analogously, if we can find iterative optimization algorithms that correspond to the  $\ell_\infty$  and  $\ell_1$  penalties, instead of solving the problem using proximal gradient descent once for each value of  $\lambda$ , all different early stopping regularization strengths could be obtained by one single call of that algorithm.

The similarities between  $\ell_1$  regularization (lasso) and the iterative optimization method forward stagewise regression (also known as coordinate descent, or, for infinitesimal step size, coordinate flow) are well studied, for instance by Efron et al. (2004) and Hastie et al. (2007). Efron et al. (2004) show that under

certain conditions, including for uncorrelated data, the solutions paths of lasso and coordinate flow coincide, but even when these conditions do not hold, the solution paths tend to be very similar, where smaller correlations lead to larger similarities. Similarly, by comparing Equations 10 and 11, or Equations 6 and 10 by [Allerbo et al. \(2023\)](#) for standard linear regression, we note that also for  $\ell_2$  regularization and gradient flow, the solution paths exactly coincide for uncorrelated data.

The topic of a gradient-based algorithm similar to  $\ell_\infty$  regularization does not seem to be as well studied. However, in Proposition 5.1, we show that the solutions of  $\ell_\infty$  regularization and sign gradient flow coincide for uncorrelated data. This is shown using the constrained form of regularized regression, which, due to Lagrangian duality, is equivalent to the penalized form.

Before stating the proposition, we first review the three gradient-based algorithms and their corresponding infinitesimal step-size versions. For a loss function,  $L(\cdot)$ , that quantifies the reconstruction error, we denote the gradient and its maximum component at time step  $k$  as

$$\begin{aligned} \mathbf{g}_k &:= \frac{\partial L(\boldsymbol{\theta}_k; \mathbf{X}, \mathbf{y})}{\partial \boldsymbol{\theta}_k} \\ m_k &:= \arg \max_d (|\mathbf{g}_k|)_d, \end{aligned}$$

where the absolute value of the gradient is evaluated element-wise. The update rules of gradient descent/flow, sign gradient descent/flow, and coordinate descent/flow are stated in Equation 14, where  $\eta$  denotes the optimization step size, and where  $\mathbf{I}_m(t)$  is a diagonal matrix that defines which gradient component(s) that is updated at time  $t$  (for details about  $\mathbf{I}_m$ , see [Allerbo et al. \(2023\)](#)).

$$\begin{aligned} \text{Gradient Descent:} \quad & \boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \eta \cdot \mathbf{g}_k \\ \text{Gradient Flow:} \quad & \frac{\partial \boldsymbol{\theta}(t)}{\partial t} = -\mathbf{g}(t) \\ \text{Sign Gradient Descent:} \quad & \boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \eta \cdot \text{sign}(\mathbf{g}_k) \\ \text{Sign Gradient Flow:} \quad & \frac{\partial \boldsymbol{\theta}(t)}{\partial t} = -\text{sign}(\mathbf{g}(t)) \\ \text{Coordinate Descent:} \quad & (\boldsymbol{\theta}_{k+1})_{m_k} = (\boldsymbol{\theta}_k)_{m_k} - \eta \cdot \text{sign}(\mathbf{g}_k)_{m_k} \\ \text{Coordinate Flow:} \quad & \frac{\partial \boldsymbol{\theta}(t)}{\partial t} = -\mathbf{I}_m(t) \cdot \text{sign}(\mathbf{g}(t)). \end{aligned} \tag{14}$$

*Remark 5.1.* For coordinate descent (flow), in each iteration (at each time), only the coordinate(s) corresponding to the maximal absolute gradient value is updated.

*Remark 5.2.* The name “coordinate descent” is sometimes also used for other, related, algorithms.

**Proposition 5.1.**

(a) Let  $\hat{\beta}^\infty(c)$  denote the solution to

$$\min_{\beta \in \mathbb{R}^p} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 \quad \text{s.t.} \quad \|\beta\|_\infty \leq c,$$

and let  $\hat{\beta}^{SGF}(t)$  denote the solution to

$$\frac{d\hat{\beta}(t)}{dt} = -\text{sign} \left( \frac{\partial}{\partial \hat{\beta}(t)} \left( \|\mathbf{y} - \mathbf{X}\hat{\beta}(t)\|_2^2 \right) \right), \quad \hat{\beta}(0) = \mathbf{0}.$$

When  $\mathbf{X}^\top \mathbf{X}$  is diagonal, with elements  $\{s_{ii}\}_{i=1}^p$ , the two solutions decompose element-wise and coincide for  $c = t$ :

$$\hat{\beta}_i^\infty(t) = \hat{\beta}_i^{SGF}(t) = \text{sign}((\mathbf{X}^\top \mathbf{y})_i / s_{ii}) \cdot \min(t, |(\mathbf{X}^\top \mathbf{y})_i / s_{ii}|).$$

(b) Let  $\hat{\alpha}^\infty(c)$  denote the solution to

$$\min_{\alpha \in \mathbb{R}^n} \|\mathbf{y} - \mathbf{K}\alpha\|_{\mathbf{K}^{-1}}^2 \quad \text{s.t.} \quad \|\alpha\|_\infty \leq c,$$

and let  $\hat{\alpha}^{SGF}(t)$  denote the solution to

$$\frac{d\hat{\alpha}(t)}{dt} = -\text{sign} \left( \frac{\partial}{\partial \hat{\alpha}(t)} \left( \|\mathbf{y} - \mathbf{K}\hat{\alpha}(t)\|_{\mathbf{K}^{-1}}^2 \right) \right), \quad \hat{\alpha}(0) = \mathbf{0}.$$

When  $\mathbf{K}$  is diagonal, with elements  $\{k_{ii}\}_{i=1}^n$ , the two solutions decompose element-wise and coincide for  $c = t$ :

$$\hat{\alpha}_i^\infty(t) = \hat{\alpha}_i^{SGF}(t) = \text{sign}(y_i / k_{ii}) \cdot \min(t, |y_i / k_{ii}|).$$

*Remark 5.3.* For infinitesimal step size, sign gradient descent, Adam (Kingma and Ba, 2014), and RMSProp (Hinton et al., 2012) coincide, up to constants of order  $10^{-8}$  added for numerical stability (Balles and Hennig, 2018). Thus, in the uncorrelated case,  $\ell_\infty$  regularization also coincides with Adam and RMSProp with infinitesimal step size.

*Remark 5.4.* Similarly to ridge regression and gradient flow, and lasso and coordinate flow, the exact coincidence of  $\ell_\infty$  regularization and sign gradient flow only holds for uncorrelated data. However, due to continuity, unless the correlations are very large, the solution paths still tend to be very similar.

## 5.2. Properties of the early stopping algorithms

One might wonder whether the fact that the exact coincidence of the solution paths for explicit and early stopping regularization holds only for uncorrelated data, implies that the properties of the explicitly regularized methods do not

transfer to the early stopping methods on uncorrelated data. Or is it still the case that KSGD, just like  $K\ell_\infty R$ , provides a solution that is robust to outliers and that KCD, like  $K\ell_1 R$ , promotes a sparse solution that includes the most significant observations first? In Proposition 5.2, we show that this is indeed the case by using the equivalent interpretation of kernel regression as linear regression in feature space from Equation 4.

**Proposition 5.2.** *Solving  $\min_{\alpha} \|\mathbf{y} - \mathbf{K}\alpha\|_{\mathbf{K}^{-1}}^2$  for  $\alpha$  with*

- *gradient descent, is equivalent of solving  $\min_{\beta} \|\mathbf{y} - \Phi\beta\|_2^2$  for  $\beta$  with gradient descent.*
- *sign gradient descent, is equivalent of solving  $\min_{\beta} \|\mathbf{y} - \Phi\beta\|_1$  for  $\beta$  with gradient descent.*
- *coordinate descent, is equivalent of solving  $\min_{\beta} \|\mathbf{y} - \Phi\beta\|_\infty$  for  $\beta$  with gradient descent.*

According to Proposition 5.2, KSGD and KCD correspond to feature space regression with the  $\ell_1$  and  $\ell_\infty$  loss functions, respectively. (Note that while we previously discussed different norms for the *penalty term*, Proposition 5.2 considers different norms for *loss function*.) Regression with the  $\ell_1$  loss function, which is known as least absolute deviations, is a commonly used method for robust regression. Vice versa, regression with the  $\ell_\infty$  loss function minimizes the maximum residual, i.e. more extreme observations contribute more to the solution than less extreme observations do.

Even for convex problems, a too large optimization step size will cause the solution to diverge. In Proposition 5.3, we present upper bounds for the step sizes, which guarantee that the solutions do not diverge.

**Proposition 5.3.** *When solving  $\min_{\alpha} \|\mathbf{y} - \mathbf{K}\alpha\|_{\mathbf{K}^{-1}}^2$  for  $\alpha$  with*

- *gradient descent, i.e. when  $\hat{\alpha}_{k+1} = \hat{\alpha}_k + \eta \cdot (\mathbf{y} - \mathbf{K}\hat{\alpha}_k)$ , then*

$$\eta \leq \frac{2}{s_{\max}(\mathbf{K})} \implies \|\mathbf{y} - \mathbf{K}\hat{\alpha}_{k+1}\|_2 \leq \|\mathbf{y} - \mathbf{K}\hat{\alpha}_k\|_2,$$

*where  $s_{\max}(\mathbf{K})$  denotes the largest singular value of  $\mathbf{K}$ .*

- *sign gradient descent, i.e. when  $\hat{\alpha}_{k+1} = \hat{\alpha}_k + \eta \cdot \text{sign}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k)$ , then*

$$\eta \leq \min_{1 \leq i \leq n} (|\mathbf{y} - \mathbf{K}\hat{\alpha}_k|_i) \implies \|\mathbf{y} - \mathbf{K}\hat{\alpha}_{k+1}\|_1 \leq \|\mathbf{y} - \mathbf{K}\hat{\alpha}_k\|_1,$$

*where the absolute value is taken element-wise.*

- *coordinate descent, i.e. when  $(\hat{\alpha}_{k+1})_{m_k} = (\hat{\alpha}_k)_{m_k} + \eta \cdot \text{sign}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k)_{m_k}$ , then*

$$\eta \leq \min_{1 \leq i \leq n} (|\mathbf{y} - \mathbf{K}\hat{\alpha}_k|_i) \implies \|\mathbf{y} - \mathbf{K}\hat{\alpha}_{k+1}\|_\infty \leq \|\mathbf{y} - \mathbf{K}\hat{\alpha}_k\|_\infty,$$

*where  $m_k$  denotes the coordinate with the largest absolute value, and where the absolute value is taken element-wise.*

*Remark 5.5.* Note that depending on which algorithm is used, the convergence is obtained for different norms, similar to in Proposition 5.2.

## 6. Experiments

In this section, we demonstrate our methods on synthetic and real data. We first demonstrate the six methods in Table 1 on simple synthetic data in Section 6.1. Then, in Section 6.2, we compare KSGD and  $K\ell_\infty R$  to four other robust kernel regression methods on five real data sets.  $K\ell_\infty R$  and  $K\ell_1 R$  were implemented by solving Equations 12 and 13 using proximal gradient descent, with an optimization step size of 0.01, a step size that was also used for KGD, KSGD, and KCD. We consistently used the Gaussian kernel,  $k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2\sigma^2}\right)$ ; in Appendix C, we also present results for four additional kernels.

### 6.1. Demonstrations on synthetic data

To get a better intuition of the six methods discussed in Sections 4 and 5, in this section, we demonstrate them on simple, synthetic data.

In Figures 1 and 2, the similarities between explicit regularization and gradient-based optimization with early stopping are demonstrated. The solutions are very similar, although not exactly identical. In Figure 1, we observe that for




FIG 1. Comparisons of the effects of KGF/KRR, KSGD/ $K\ell_\infty R$ , and KCD/ $K\ell_1 R$  on  $\hat{f}(x)$ . In the top panel, a larger regularization, or a shorter training time, is used than in the bottom panel. For KGF/KRR, we use  $t = 1/\lambda$ , while for the other two cases,  $t$  and  $\lambda$  are chosen so that the functions coincide as well as possible.

As proposed by Proposition 4.1, the KGF and KRR solutions differ most where the non-regularized solution is large, and as suggested by Proposition 4.2, the KGF solution tends to lie closer to the observations than the KRR solution does.

For KSGD/ $K\ell_\infty R$ , the more extreme observations tend to be penalized harder, resulting in a more robust solution that is less sensitive to outliers. For KCD/ $K\ell_1 R$ , some observations do not contribute to the solution, resulting in peaks at the more significant observations.

The solutions obtained through early stopping are very similar to those obtained through explicit regularization, although not exactly identical.




FIG 2. Comparisons of the effects of KGF/KRR, KSGD/ $Kl_\infty R$ , and KCD/ $Kl_1 R$  on the seven elements in  $\hat{\alpha}$ . In the top panel, a larger regularization, or a shorter training time, is used than in the bottom panel. For KGF/KRR, we use  $t = 1/\lambda$ , while, for the two other cases,  $t$  and  $\lambda$  are chosen so that the functions coincide as well as possible.

For KSGD/ $Kl_\infty R$ , the more extreme observations tend to be penalized harder, resulting in no extreme  $\hat{\alpha}_i$ 's. For KCD/ $Kl_1 R$ , some observations do not contribute to the solution, resulting in the corresponding  $\hat{\alpha}_i$ 's being 0.

The solutions obtained through early stopping are very similar to those obtained through explicit regularization, although not exactly identical.

KGF and KRR, all residuals are treated equally and all parts of the function are updated at a pace that is proportional to the distance to the solution at convergence. For KSGD and  $Kl_\infty$ , which are less sensitive to outliers, all parts of the function are initially updated at the same pace, and, in contrast to more moderate observations, the most extreme observations are not fully incorporated until at the end of the training. For KCD and  $Kl_1 R$ , which minimize the maximum residual, the most extreme observations are incorporated first, while more moderate observations are initially ignored.

In Figure 3, we further compare KSGD/ $Kl_\infty R$  and KCD/ $Kl_1 R$  to KGD/KRR. For robust regression, to obtain data with outliers, 100 observations were sampled according to

$$x \sim \mathcal{U}(-10, 10), \quad y = \sin\left(\frac{\pi}{2} \cdot x\right) + \mathcal{C}(0, 0.1), \quad (15)$$

where  $\mathcal{U}(a, b)$  denotes the uniform distribution on  $(a, b)$ , and  $\mathcal{C}(0, \gamma)$  the centered Cauchy distribution with scale parameter  $\gamma$ .

To demonstrate sparse regression on a signal that is mostly zero, with a narrow but distinct peak, 100 observations were sampled according to

$$x \sim \mathcal{U}(-10, 10), \quad y = e^{-5 \cdot x^2} + \mathcal{N}(0, 0.1^2), \quad (16)$$

where again  $\mathcal{U}(a, b)$  denotes the uniform distribution and  $\mathcal{N}(\mu, \sigma^2)$  denotes the normal distribution. On the first (robust) data set, we compare KSGD,  $Kl_\infty R$ ,




FIG 3. Modelling the data generated by Equations 15 (top) and 16 (bottom) using KSGD/ $K\ell_\infty R$  (top) or KCD/ $K\ell_1 R$  (bottom) KGD/KRR (both panels). For the first data set, KSGD/ $K\ell_\infty R$  are less affected by the outliers than KGD/KRR are, and KGD is more affected by the outliers than KRR is. For the second data set, in contrast to KGD/KRR, KCD/ $K\ell_1 R$  are able to model the peak without incorporating the noise.

KGD, and KRR, and on the second (sparse) data set, we compare KCD,  $K\ell_1 R$ , KGD, and KRR. In both cases, bandwidth and regularization were selected by 10-fold cross-validation evaluating  $30 \times 30$  logarithmically spaced candidate values.

As expected, KGD/KRR are more sensitive to the outliers than KSGD/ $K\ell_\infty R$  are, with KGD being more sensitive than KRR, something that probably can be attributed to the fact that the KGF solution tends to lie closer to the observations than the KRR solution does, as suggested by Proposition 4.2. For the second data set, to perform well at the peak, KGD/KRR must also incorporate the noise in the regions where the true signal is zero. KCD/ $K\ell_1 R$  include the most significant observations first and are thus able to perform well at the peak without incorporating the noise in the zero regions.

## 6.2. Robust kernel regression on real data

In this section, we use five real data sets to compare robust kernel regression through KSGD and  $K\ell_\infty R$  to four state-of-the-art methods for robust kernel regression, and to KGD and KRR. The competing methods, for which we used the default implementations, are presented in Table 2, and the data sets in Table 3.



TABLE 2  
Robust kernel regression models used.

Method	Abbreviation
Huber loss implemented by iteratively reweighted least squares	KMR-H
Tukey loss implemented by iteratively reweighted least squares	KMR-T
Fast quantile regression implemented by <a href="#">Zheng (2022)</a>	KQR-Z
Fast quantile regression implemented by <a href="#">Tang et al. (2024)</a>	KQR-T

TABLE 3  
Real data sets, with corresponding reference when existing, used for comparing the methods.

Data set	Size, $n \times p$
Sound pressure of airfoils <sup>1</sup>	$1502 \times 6$
House values in California ( <a href="#">Pace and Barry, 1997</a> ) <sup>2</sup>	$20640 \times 9$
Energy consumption in steel production ( <a href="#">Sathishkumar et al., 2021</a> ) <sup>3</sup>	$35040 \times 7$
Critical temperature of superconductors ( <a href="#">Hamidieh, 2018</a> ) <sup>4</sup>	$21263 \times 82$
Daily temperature in the U.K. in the year 2000 ( <a href="#">Wood et al., 2017</a> ) <sup>5</sup>	$45568 \times 5$

Each data set was standardized to zero mean and unit variance. Then 50 random splits were created by randomly selecting 100 observations, which in turn were split 80%/20% into training and testing data.

Hyper-parameter selection was performed through 10-fold cross-validation, with 30 logarithmically spaced values for the kernel bandwidth,  $\sigma$ , and regularization strength,  $\lambda$ , respectively. Moreover, all four methods in Table 2 come with an additional hyper-parameter, for which we evaluated three different values. For quantile regression, where the parameter  $\tau \in (0, 1)$  defines which quantile to use, we used  $\tau \in [0.25, 0.5, 0.75]$ . For the two M-estimators, which both come with a parameter  $k$ , which tunes the sensitivity toward outliers, we used  $k \in [0.5k_d, k_d, 2k_d]$ , where  $k_d$  is the default tuning constant. For Huber regression,  $k_d = 1.345\hat{\sigma}$ , and for Tukey bisquare regression,  $k_d = 4.685\hat{\sigma}$ , where  $\hat{\sigma}$  denotes the standard deviation of the residuals.

For the early stopping algorithms, rather than using an explicit regularization parameter  $\lambda$ , the stopping times were selected by monitoring performance on a validation data set, for which 10% of the training set was used. Once the performance started to decrease on the validation set, the algorithm was terminated. The experiments were run on a cluster with AMD EPYC Zen 2, 2.25 GHz cores, using one core per experiment.

In Tables 4 and 5, we present the computation time and  $R^2$  on test data for the experiments.<sup>6</sup> The difference between the two tables is that in Table 5,

<sup>1</sup>The data set is available at <https://archive.ics.uci.edu/dataset/291/airfoil+self+noise>.

<sup>2</sup>The data set is available at [https://www.dcc.fc.up.pt/~ltorgo/Regression/cal\\_housing.html](https://www.dcc.fc.up.pt/~ltorgo/Regression/cal_housing.html).

<sup>3</sup>The data set is available at <https://archive.ics.uci.edu/dataset/851/steel+industry+energy+consumption>.

<sup>4</sup>The data set is available at <https://archive.ics.uci.edu/dataset/464/superconductivity+data>.

<sup>5</sup>The data set is available at <https://www.maths.ed.ac.uk/~swood34>.

<sup>6</sup> $R^2 := 1 - \|\mathbf{y} - \hat{\mathbf{f}}\|_2^2 / \|\mathbf{y} - \bar{y}\|_2^2 \leq 1$ , where  $\mathbf{y}$  and  $\bar{y}$  denotes the vector of observations

TABLE 4

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, *without* amplified outliers. The six robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	9.36, (5.41, 13.8)	0.53, (−0.47, 0.82)
	$K\ell_\infty R$	383, (370, 404)	0.56, (−0.54, 0.87)
	KMR-H	172, (158, 192)	0.51, (−0.64, 0.84)
	KMR-T	161, (151, 178)	0.53, (−0.45, 0.84)
	KQR-A	289, (275, 311)	0.55, (−0.55, 0.84)
	KQR-B	314, (185, 345)	0.41, (−0.13, 0.72)
	KGD	14.7, (10.9, 18.4)	0.57, (−0.49, 0.85)
	KRR	2.92, (2.87, 3.04)	0.53, (−0.44, 0.84)
California House Values	KSGD	4.14, (2.20, 8.39)	0.63, (−0.09, 0.86)
	$K\ell_\infty R$	325, (311, 348)	0.64, (−0.17, 0.88)
	KMR-H	137, (124, 153)	0.60, (−0.30, 0.89)
	KMR-T	120, (112, 135)	0.58, (−0.31, 0.89)
	KQR-A	271, (259, 294)	0.61, (−0.19, 0.86)
	KQR-B	206, (115, 218)	0.33, (−0.10, 0.62)
	KGD	12.6, (9.20, 15.7)	0.60, (−0.17, 0.89)
	KRR	2.95, (2.90, 3.14)	0.60, (−0.34, 0.88)
Steel Energy Consumption	KSGD	12.0, (9.85, 20.3)	0.98, (0.92, 0.99)
	$K\ell_\infty R$	416, (401, 436)	0.98, (0.94, 0.99)
	KMR-H	166, (146, 199)	0.99, (0.98, 1.00)
	KMR-T	156, (148, 191)	0.99, (0.98, 1.00)
	KQR-A	267, (256, 332)	0.99, (0.97, 1.00)
	KQR-B	541, (475, 620)	0.90, (0.60, 0.98)
	KGD	22.2, (16.8, 26.0)	0.99, (0.94, 1.00)
	KRR	2.95, (2.90, 3.29)	0.99, (0.97, 1.00)
Superconductor Critical Temperature	KSGD	12.6, (8.54, 17.7)	0.59, (−0.01, 0.84)
	$K\ell_\infty R$	409, (401, 425)	0.64, (0.14, 0.89)
	KMR-H	169, (155, 193)	0.66, (0.05, 0.86)
	KMR-T	160, (135, 189)	0.65, (0.04, 0.89)
	KQR-A	288, (255, 329)	0.61, (−0.04, 0.88)
	KQR-B	482, (338, 632)	0.46, (−0.08, 0.69)
	KGD	14.8, (10.9, 17.4)	0.67, (0.21, 0.89)
	KRR	3.44, (3.37, 3.88)	0.66, (0.17, 0.89)
U.K. Temperature	KSGD	6.17, (2.99, 9.76)	0.40, (−0.07, 0.69)
	$K\ell_\infty R$	370, (358, 388)	0.37, (−0.06, 0.67)
	KMR-H	182, (161, 206)	0.38, (−1.13, 0.68)
	KMR-T	180, (157, 192)	0.39, (−1.03, 0.67)
	KQR-A	310, (282, 340)	0.36, (−1.24, 0.66)
	KQR-B	321, (258, 343)	0.42, (−0.06, 0.66)
	KGD	11.3, (8.31, 15.1)	0.40, (−0.31, 0.71)
	KRR	2.97, (2.87, 3.14)	0.40, (−1.13, 0.68)

TABLE 5

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, *with* amplified outliers. The six robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster. The two non-robust methods, KGD/KRR, do not perform as well as the robust methods.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	10.5, (5.29, 13.8)	0.44, (−0.05, 0.74)
	$K\ell_\infty$ R	395, (375, 431)	0.42, (−0.76, 0.75)
	KMR-H	168, (101, 199)	0.38, (−1.12, 0.71)
	KMR-T	154, (115, 182)	0.36, (−1.08, 0.73)
	KQR-A	304, (284, 346)	0.38, (−1.10, 0.74)
	KQR-B	346, (287, 714)	0.25, (−5.95, 0.61)
	KGD	12.4, (3.23, 16.5)	0.25, (−82.33, 0.71)
	KRR	2.99, (2.88, 3.22)	0.27, (−1.61, 0.70)
California House Values	KSGD	4.65, (2.51, 8.68)	0.46, (−0.06, 0.80)
	$K\ell_\infty$ R	352, (326, 384)	0.42, (−0.50, 0.76)
	KMR-H	130, (80.1, 170)	0.36, (−0.74, 0.81)
	KMR-T	118, (91.1, 143)	0.37, (−0.63, 0.77)
	KQR-A	297, (268, 327)	0.30, (−0.44, 0.82)
	KQR-B	259, (160, 283)	0.23, (−2.20, 0.56)
	KGD	11.0, (4.98, 14.2)	0.17, (−4.39, 0.76)
	KRR	2.99, (2.92, 3.35)	0.29, (−0.22, 0.77)
Steel Energy Consumption	KSGD	14.5, (11.1, 23.3)	0.89, (0.06, 0.97)
	$K\ell_\infty$ R	426, (404, 458)	0.85, (−1.13, 0.98)
	KMR-H	168, (122, 220)	0.86, (−1.59, 0.98)
	KMR-T	174, (141, 214)	0.86, (−0.54, 0.98)
	KQR-A	308, (263, 340)	0.88, (−2.11, 0.98)
	KQR-B	707, (578, 1.220)	0.67, (−1.68, 0.95)
	KGD	13.3, (5.84, 17.8)	0.74, (−33.96, 0.97)
	KRR	2.96, (2.87, 3.34)	0.78, (−3.34, 0.97)
Superconductor Critical Temperature	KSGD	14.3, (10.7, 21.3)	0.45, (−0.22, 0.81)
	$K\ell_\infty$ R	416, (401, 441)	0.48, (−0.49, 0.84)
	KMR-H	159, (114, 202)	0.48, (−0.35, 0.79)
	KMR-T	146, (118, 201)	0.48, (−0.33, 0.82)
	KQR-A	290, (258, 326)	0.44, (−0.62, 0.83)
	KQR-B	586, (406, 1.190)	0.28, (−0.96, 0.69)
	KGD	12.1, (5.00, 15.4)	0.37, (−1.41, 0.83)
	KRR	3.46, (3.37, 3.89)	0.47, (−0.51, 0.82)
U.K. Temperature	KSGD	7.29, (3.14, 11.6)	0.33, (−0.13, 0.63)
	$K\ell_\infty$ R	385, (368, 423)	0.25, (−1.13, 0.62)
	KMR-H	164, (99.7, 213)	0.29, (−3.08, 0.66)
	KMR-T	156, (117, 186)	0.27, (−2.31, 0.65)
	KQR-A	305, (282, 339)	0.22, (−3.07, 0.63)
	KQR-B	328, (277, 532)	0.20, (−0.64, 0.60)
	KGD	9.97, (3.15, 15.7)	0.11, (−20.91, 0.61)
	KRR	2.92, (2.87, 3.16)	0.17, (−2.46, 0.61)

the presence of outliers is amplified by multiplying each element,  $y_i$ , in the response vector by  $(1 + |\varepsilon_i|)$ , where  $\varepsilon_i \sim \mathcal{C}(0, 0.01)$  is Cauchy distributed. We note that amplifying the outliers only affects the relative performance between the robust and non-robust methods, not the relative performance among the robust methods.

For all five data sets, KSGD and  $\text{K}\ell_\infty\text{R}$  perform at least as well as the competing four robust methods in terms of test  $R^2$ . However, KSGD is significantly faster than the competing methods, performing between one and two orders of magnitude faster.

Comparing Tables 4 and 5, we note that KGD and KRR are much more sensitive to outliers than the robust methods are. We also note with amplified outliers, KRR tends to perform better than KGD in terms of test  $R^2$ , while in Table 4, the tendency is the opposite. This can probably be attributed to the fact that the KGF solution tends to lie closer to the observations than the KRR solution does, as suggested by Proposition 4.2. This is often beneficial, however not in the presence of outliers.

## 7. Conclusions

We introduced an equivalent formulation of kernel ridge regression and used it to define kernel regression with the  $\ell_\infty$  and  $\ell_1$  penalties, and for solving kernel regression with gradient-based optimization methods. We introduced the methods kernel sign gradient descent and kernel coordinate descent and utilized the similarities between  $\ell_\infty$  regularization and sign gradient descent, and between  $\ell_1$  regularization and coordinate descent (forward stagewise regression), to obtain computationally efficient algorithms for robust and sparse kernel regression, respectively. We demonstrated on five real data sets that our implementation of robust regression is one to two orders of magnitude faster than existing robust kernel regression methods.

Our generalizations of kernel ridge regression, together with regularization through early stopping, enable computationally efficient, kernelized robust, and sparse, regression. Although we investigated only kernel regression with the  $\ell_2$ ,  $\ell_1$ , and  $\ell_\infty$  penalties, many other penalties could be considered, such as the adaptive lasso (Zou, 2006), the group lasso (Yuan and Lin, 2006), the exclusive lasso (Zhou et al., 2010) and OSCAR (Bondell and Reich, 2008). This is, however, left for future work.

Code is available at [https://github.com/allerbo/fast\\_robust\\_kernel\\_regression](https://github.com/allerbo/fast_robust_kernel_regression).

---

and its mean, and  $\hat{\mathbf{f}}$  denotes the vector of model predictions, measures the proportion of the variance in  $\mathbf{y}$  that is explained by the model.  $R^2 = 1$  corresponds to a perfect fit, while a model that predicts  $\hat{f}_i = \bar{y}$  for all  $i$  results in  $R^2 = 0$ ; thus a negative value of  $R^2$  is possible, but it corresponds to a model that performs worse than always predicting the mean of the observed data.

## Appendix A: Calculations for $[\hat{\mathbf{f}}^\top, \hat{\mathbf{f}}^{*\top}]^\top$

In this section, we revisit some of the calculations in the main manuscript and reformulate them directly in terms of  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$ , rather than obtaining them by multiplying  $\boldsymbol{\alpha}$  by  $[\mathbf{K}^\top, \mathbf{K}^{*\top}]^\top$ .

### A.1. Equivalent formulations of kernel ridge regression

In Proposition A.1, we present three different objective functions, that all render the same solution for  $[\hat{\mathbf{f}}^\top, \hat{\mathbf{f}}^{*\top}]^\top$  as in Equation 3. When expressing KRR in terms of  $\boldsymbol{\alpha} \in \mathbb{R}^n$ , it is enough to include  $\mathbf{K} \in \mathbb{R}^{n \times n}$  in the expression to capture the dynamics of the kernel. However, when expressing KRR in terms of  $[\hat{\mathbf{f}}^\top, \hat{\mathbf{f}}^{*\top}]^\top \in \mathbb{R}^{n+n^*}$ , we need to introduce the larger kernel matrix

$$\mathbf{K}^{**} := \begin{bmatrix} \mathbf{K}(\mathbf{X}, \mathbf{X}) & \mathbf{K}(\mathbf{X}, \mathbf{X}^*) \\ \mathbf{K}(\mathbf{X}^*, \mathbf{X}) & \mathbf{K}(\mathbf{X}^*, \mathbf{X}^*) \end{bmatrix} \in \mathbb{R}^{(n+n^*) \times (n+n^*)}$$

in order to let the kernel affect not only  $\mathbf{f} \in \mathbb{R}^n$ , but also  $\mathbf{f}^* \in \mathbb{R}^{n^*}$ . Furthermore, we need the extended response vector  $[\mathbf{y}^\top, \tilde{\mathbf{y}}^\top]^\top$ , where  $\tilde{\mathbf{y}}$  is a copy of  $\mathbf{f}^*$ , so that  $\tilde{\mathbf{y}} - \mathbf{f}^* = \mathbf{0}$ .

**Proposition A.1.**

$$\begin{bmatrix} \hat{\mathbf{f}} \\ \hat{\mathbf{f}}^* \end{bmatrix} = \arg \min_{[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top \in \mathbb{R}^{n+n^*}} \frac{1}{2} \|\mathbf{y} - \mathbf{f}\|_2^2 + \frac{\lambda}{2} \left\| \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_{(\mathbf{K}^{**})^{-1}}^2 \quad (17a)$$

$$= \arg \min_{[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top \in \mathbb{R}^{n+n^*}} \frac{1}{2} \left\| \begin{bmatrix} \mathbf{y} \\ \tilde{\mathbf{y}} \end{bmatrix} - \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_2^2 + \frac{\lambda}{2} \left\| \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_{(\mathbf{K}^{**})^{-1}}^2 \quad (17b)$$

$$= \arg \min_{[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top \in \mathbb{R}^{n+n^*}} \frac{1}{2} \left\| \begin{bmatrix} \mathbf{y} \\ \tilde{\mathbf{y}} \end{bmatrix} - \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_{\mathbf{K}^{**}}^2 + \frac{\lambda}{2} \left\| \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_2^2 \quad (17c)$$

$$= \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{y}.$$

*Remark A.1.* We assert that  $\mathbf{f}^*$  does not affect the reconstruction error by requiring that  $\tilde{\mathbf{y}} - \mathbf{f}^* = \mathbf{0}$ . We may, however, not define  $\tilde{\mathbf{y}}$  to equal  $\mathbf{f}^*$ , since we do not want  $\tilde{\mathbf{y}}$  to be considered when differentiating Equations 17b and 17c with respect to  $\mathbf{f}^*$ .

*Remark A.2.* The term  $\left\| \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_{(\mathbf{K}^{**})^{-1}}^2$ , might appear a bit peculiar, including the inverse of the extended kernel matrix (which never needs to be calculated due to cancellations). However, it is very closely related to the expression for the reproducing kernel Hilbert space norm,  $\|f\|_{\mathcal{H}_k}^2$ , obtained when expressing the functions  $f$  and  $k$  in terms of the orthogonal basis given by Mercer's Theorem:  $f(\mathbf{x}) = \sum_{i=1}^{\infty} f_i \phi_i(\mathbf{x})$  and  $k(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^{\infty} k_i \phi_i(\mathbf{x}) \phi_i(\mathbf{x}')$ , with  $\|f\|_{\mathcal{H}_k}^2 := \sum_{i=1}^{\infty} \frac{f_i^2}{k_i}$ . Defining the vector  $\tilde{\mathbf{f}}$  and the diagonal matrix  $\tilde{\mathbf{K}}$  according to  $\tilde{\mathbf{f}}_i = f_i$  and  $\tilde{\mathbf{K}}_{ii} = k_i$ , we obtain  $\|f\|_{\mathcal{H}_k}^2 = \sum_{i=1}^{\infty} \frac{f_i^2}{k_i} = \tilde{\mathbf{f}}^\top \tilde{\mathbf{K}}^{-1} \tilde{\mathbf{f}} = \|\tilde{\mathbf{f}}\|_{\tilde{\mathbf{K}}^{-1}}^2$ .

### A.2. Kernel gradient descent and kernel gradient flow

The gradient of Equation 17c with respect to  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$  is

$$\begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\mathbf{f} - \mathbf{y}) + \lambda \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix},$$

which coincides with Equation 7, as expected. Thus, the KGD update in  $[\hat{\mathbf{f}}^\top, \hat{\mathbf{f}}^{*\top}]^\top$  is

$$\begin{bmatrix} \hat{\mathbf{f}}_{k+1} \\ \hat{\mathbf{f}}_{k+1}^* \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{f}}_k \\ \hat{\mathbf{f}}_k^* \end{bmatrix} + \eta \cdot \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\mathbf{y} - \hat{\mathbf{f}}_k), \quad \begin{bmatrix} \hat{\mathbf{f}}_0 \\ \hat{\mathbf{f}}_0^* \end{bmatrix} = \mathbf{0},$$

with the corresponding differential equation

$$\frac{\partial}{\partial t} \begin{bmatrix} \hat{\mathbf{f}}(t) \\ \hat{\mathbf{f}}^*(t) \end{bmatrix} = \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\mathbf{y} - \hat{\mathbf{f}}(t)), \quad \begin{bmatrix} \hat{\mathbf{f}}(0) \\ \hat{\mathbf{f}}^*(0) \end{bmatrix} = \mathbf{0}, \quad (18)$$

whose solution is given by Equation 19.

$$\begin{bmatrix} \hat{\mathbf{f}}(t) \\ \hat{\mathbf{f}}^*(t) \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{K}^* \mathbf{K}^{-1} \end{bmatrix} (\mathbf{I} - \exp(-t\mathbf{K}))\mathbf{y}. \quad (19)$$

### A.3. Kernel regression with the $\ell_\infty$ and $\ell_1$ norms

If we replace the squared  $\ell_2$  norm in Equation 17c by the  $\ell_\infty$  and  $\ell_1$  norms, we obtain

$$\arg \min_{[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top \in \mathbb{R}^{n+n^*}} \frac{1}{2} \left\| \begin{bmatrix} \mathbf{y} \\ \tilde{\mathbf{y}} \end{bmatrix} - \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_{\mathbf{K}^{**}}^2 + \lambda \left\| \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_\infty \quad (20)$$

and

$$\arg \min_{[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top \in \mathbb{R}^{n+n^*}} \frac{1}{2} \left\| \begin{bmatrix} \mathbf{y} \\ \tilde{\mathbf{y}} \end{bmatrix} - \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_{\mathbf{K}^{**}}^2 + \lambda \left\| \begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \right\|_1. \quad (21)$$

In contrast to ridge regression, where the formulations in  $\boldsymbol{\alpha}$  and  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$  are equivalent in the sense that the latter can always be obtained from the first through multiplication with  $[\mathbf{K}^\top, \mathbf{K}^{*\top}]^\top$ , the solutions of Equations 20 and 21 cannot be obtained from the solutions of Equations 12 and 13.

Since  $\ell_\infty$  regularization promotes a solution with no extreme parameters, many of the predictions in the solution to Equation 20 are exactly equal. For  $\ell_1$  regularization, which promotes a sparse solution with the most significant parameters included first, in the solution to Equation 21, many of the predictions equal exactly zero, with only the most extreme predictions being non-zero. In Figure 4, we exemplify this for explicit regularization for two different regularization strengths. Even though it is technically possible to use gradient-based optimization with early stopping for  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$ , since the reconstruction error, and thus the gradient, is affected only by  $\mathbf{f}$  and not by  $\mathbf{f}^*$ , this does not work well for the sign gradient and coordinate descent algorithms.

Finally, for completeness, we present the analog of Proposition 5.1 for  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$  in Proposition A.2.

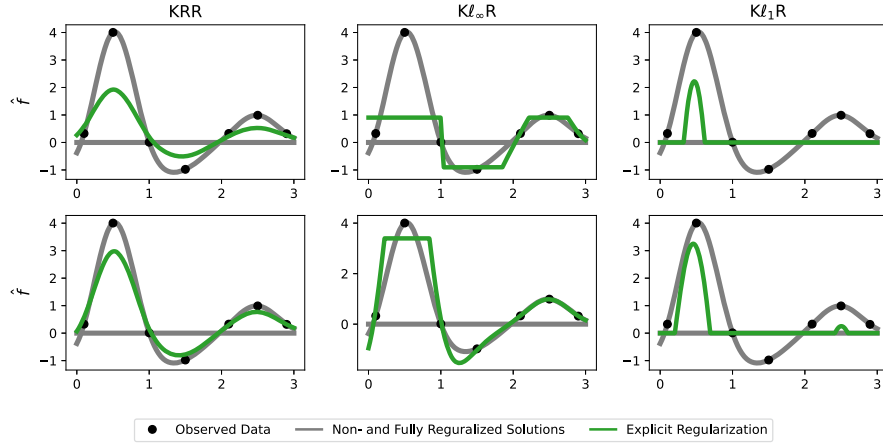


FIG 4. Comparisons of the effects of  $\ell_2$ ,  $\ell_\infty$ , and  $\ell_1$  regularization on  $[\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$ . In the top panel, a larger regularization is used than in the bottom panel. For  $\ell_\infty$  regularization, the absolute values of many predictions are exactly equal. To obtain this, compared to the non-regularized solution, some predictions are shifted away from zero, while some predictions are shifted toward zero. For  $\ell_1$  regularization, many predictions are set exactly to zero, with peaks at some extreme observations.

**Proposition A.2.** For  $\mathbf{f}^+ := [\mathbf{f}^\top, \mathbf{f}^{*\top}]^\top$  and  $\mathbf{y}^+ := [\mathbf{y}^\top, \tilde{\mathbf{y}}^\top]^\top$ , where  $\tilde{\mathbf{y}} - \mathbf{f}^* = \mathbf{0}$ , let  $\hat{\mathbf{f}}^{+\infty}(c)$  denote the solution to

$$\min_{\mathbf{f}^+ \in \mathbb{R}^{n+n^*}} \left\| \mathbf{y}^+ - \mathbf{f}^+ \right\|_{\mathbf{K}^{**}}^2 \quad \text{s.t.} \quad \|\mathbf{f}^+\|_\infty \leq c,$$

and let  $\hat{\mathbf{f}}^{+SGF}(t)$  denote the solution to

$$\frac{\partial \hat{\mathbf{f}}^+(t)}{\partial t} = -\text{sign} \left( \frac{\partial}{\partial \hat{\mathbf{f}}^+(t)} \left( \left\| \mathbf{y}^+ - \hat{\mathbf{f}}^+(t) \right\|_{\mathbf{K}^{**}}^2 \right) \right), \quad \hat{\mathbf{f}}^+(0) = \mathbf{0}.$$

When  $\mathbf{K}^{**}$  is diagonal, with elements  $\{k_{ii}\}_{i=1}^{n+n^*}$ ,  $k_{ii} > 0$ , the two solutions decompose element-wise and coincide for  $c = t$ :

$$\hat{f}_i^{+\infty}(t) = \hat{f}_i^{+SGF}(t) = \begin{cases} \text{sign}(y_i) \cdot \min(t, |y_i|) & \text{if } i \leq n \\ 0 & \text{if } i > n. \end{cases}$$

## Appendix B: Computationally efficient cross-validation

For models of the type  $\hat{\mathbf{f}} = \mathbf{H}\mathbf{y}$ , where  $\mathbf{H} \in \mathbb{R}^{n \times n}$  is a matrix that does not depend on  $\mathbf{y}$ , the fast leave-one-out cross-validation, LOOCV, and generalized

cross-validation, GCV, statistics are defined according to

$$\begin{aligned}\text{LOOCV} &:= \frac{1}{n} \sum_{i=1}^n \left( \frac{y_i - \hat{f}(\mathbf{x}_i)}{1 - H_{ii}} \right)^2 \\ \text{GCV} &:= \frac{1}{n} \sum_{i=1}^n \left( \frac{y_i - \hat{f}(\mathbf{x}_i)}{1 - \text{Tr}(\mathbf{H})/n} \right)^2.\end{aligned}$$

These two methods tend to be faster than k-fold cross-validation and provide similar performance. For KRR, where  $\hat{\mathbf{f}} = \mathbf{K}(\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{y}$ , we trivially obtain  $\mathbf{H} = \mathbf{K}(\mathbf{K} + \lambda \mathbf{I})^{-1}$ . For other penalties than the ridge penalty, it is however not trivial to obtain  $\mathbf{H}$ . Tibshirani (1996) noted that the lasso objective,  $\frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1$ , can be reformulated as  $\frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \frac{\lambda}{2} \boldsymbol{\beta}^\top \mathbf{B} \boldsymbol{\beta}$ , where  $\mathbf{B} \in \mathbb{R}^{p \times p}$  is a diagonal matrix with  $B_{ii} = 2/|\beta_i|$ . If  $\mathbf{B}$  is treated as a constant, this formulation becomes a weighted ridge regression problem which is minimized by  $\hat{\boldsymbol{\beta}} = \mathbf{X}(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{B})^{-1} \mathbf{X}^\top \mathbf{y}$ , and thus, for a given  $\hat{\boldsymbol{\beta}}$ , the author proposed approximating  $\mathbf{H}$  as  $\mathbf{X}(\mathbf{X}^\top \mathbf{X} + \lambda \hat{\mathbf{B}})^{-1} \mathbf{X}^\top$ , where  $\hat{\mathbf{B}}$  is a diagonal matrix with  $\hat{B}_{ii} = 2/|\hat{\beta}_i|$ .

Applying these ideas in the kernel formulation, we obtain, for an  $\mathbf{A}$  such that  $\frac{1}{2} \boldsymbol{\alpha}^\top \mathbf{A} \boldsymbol{\alpha} = \|\boldsymbol{\alpha}\|$ ,

$$\begin{aligned}\hat{\boldsymbol{\alpha}} &= \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{K} \boldsymbol{\alpha}\|_{\mathbf{K}^{-1}}^2 + \lambda \|\boldsymbol{\alpha}\| \\ &= \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{K} \boldsymbol{\alpha}\|_{\mathbf{K}^{-1}}^2 + \frac{\lambda}{2} \boldsymbol{\alpha}^\top \mathbf{A} \boldsymbol{\alpha} = (\mathbf{K} + \lambda \mathbf{A})^{-1} \mathbf{y}.\end{aligned}$$

Thus  $\mathbf{H}$  can be approximated by  $\mathbf{K}(\mathbf{K} + \lambda \hat{\mathbf{A}})^{-1}$ , where  $\hat{\mathbf{A}} \in \mathbb{R}^{n \times n}$  is a matrix such that  $\frac{1}{2} \hat{\boldsymbol{\alpha}}^\top \hat{\mathbf{A}} \hat{\boldsymbol{\alpha}} = \|\hat{\boldsymbol{\alpha}}\|$ . For the  $\ell_1$  norm,  $\|\hat{\boldsymbol{\alpha}}\|_1 = \sum_{i=1}^n |\hat{\alpha}_i|$ , this is obtained if  $\hat{\mathbf{A}}$  is a diagonal matrix with  $\hat{A}_{ii} = 2/|\hat{\alpha}_i|$ . For the  $\ell_\infty$  norm,  $\|\hat{\boldsymbol{\alpha}}\|_\infty = \max_i |\hat{\alpha}_i|$ , it is obtained if  $\hat{\mathbf{A}}$  is a matrix with only zeros, except for  $\hat{A}_{mm} = 2/|\hat{\alpha}_m|$ , where  $m = \arg \max_i |\hat{\alpha}_i|$ .

In Tables 6 and 7, we present the analogs of Tables 4 and 5 for LOOCV and GCV. As expected, the two methods perform faster than 10-fold cross-validation. However, for  $\text{K}\ell_\infty\text{R}$ , the performance in terms of test  $R^2$  is much lower than for 10-fold cross-validation. For KRR, LOOCV performs well in terms of test  $R^2$ .

## Appendix C: Additional experiments

In this section, we extend the results of Tables 4 and 5 by repeating the experiments for the four kernels in Table 8. Since, at the time of writing, the KQRT-T implementation only supported the Gaussian kernel, this algorithm is omitted in this section. The results are presented in Tables 9 to 16 and agree with those of Section 6.2.



TABLE 6

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for  $K\ell_\infty R$  in combination with LOOCV and GCV on the different data sets, *without* amplified outliers. Compared to 10-fold cross validation LOOCV and GCV perform faster, but the cost of decreased predictive performance, especially for  $K\ell_\infty R$ .

Data	Method	Type of Cross-Validation	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	$K\ell_\infty R$	LOOCV	49.1, (48.2, 57.2)	-0.01, (-0.38, 0.04)
		GCV	50.0, (48.7, 57.9)	-0.01, (-0.37, 0.15)
	KRR	LOOCV	1.21, (0.891, 3.11)	0.55, (-0.29, 0.84)
		GCV	1.43, (0.910, 3.35)	0.33, (-4.74, 0.75)
California House Values	$K\ell_\infty R$	LOOCV	43.1, (41.2, 50.3)	-0.02, (-0.36, 0.05)
		GCV	43.5, (41.7, 52.8)	-0.03, (-0.37, 0.00)
	KRR	LOOCV	1.26, (0.928, 2.25)	0.60, (-0.13, 0.88)
		GCV	1.26, (0.841, 2.16)	0.47, (-0.39, 0.83)
Steel Energy Consumption	$K\ell_\infty R$	LOOCV	54.1, (51.7, 63.1)	0.00, (-0.38, 0.71)
		GCV	54.3, (51.5, 68.5)	0.57, (-0.08, 0.98)
	KRR	LOOCV	0.891, (0.793, 1.63)	0.99, (0.97, 1.00)
		GCV	0.892, (0.797, 1.57)	0.98, (0.71, 1.00)
Superconductor Critical Temperature	$K\ell_\infty R$	LOOCV	52.4, (51.2, 56.0)	-0.00, (-0.42, 0.25)
		GCV	52.8, (51.6, 56.1)	0.44, (0.01, 0.74)
	KRR	LOOCV	1.79, (1.02, 2.91)	0.69, (0.12, 0.87)
		GCV	2.30, (1.15, 3.83)	0.46, (-2.71, 0.79)
U.K. Temperature	$K\ell_\infty R$	LOOCV	48.5, (46.6, 64.7)	-0.05, (-0.38, 0.03)
		GCV	48.6, (46.8, 61.3)	-0.03, (-0.38, 0.10)
	KRR	LOOCV	1.01, (0.860, 1.59)	0.38, (-0.30, 0.68)
		GCV	0.993, (0.796, 1.65)	0.28, (-6.81, 0.71)

TABLE 7

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for  $Kl_\infty R$  in combination with LOOCV and GCV on the different data sets, *with* amplified outliers. Compared to 10-fold cross validation LOOCV and GCV perform faster, but the cost of decreased predictive performance, especially for  $Kl_\infty R$ .

Data	Method	Type of Cross-Validation	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	$Kl_\infty R$	LOOCV	52.8, (48.9, 60.4)	-0.02, (-14.57, 0.01)
		GCV	56.1, (49.7, 64.9)	-0.02, (-14.33, 0.08)
	KRR	LOOCV	0.840, (0.760, 0.925)	0.30, (-1.12, 0.66)
		GCV	0.844, (0.754, 0.997)	-0.05, (-189.78, 0.69)
California House Values	$Kl_\infty R$	LOOCV	44.7, (42.4, 50.7)	-0.02, (-0.21, 0.01)
		GCV	45.0, (43.0, 50.3)	-0.03, (-1.28, 0.29)
	KRR	LOOCV	0.985, (0.832, 3.58)	0.29, (-0.62, 0.77)
		GCV	1.01, (0.833, 4.69)	0.16, (-37.61, 0.72)
Steel Energy Consumption	$Kl_\infty R$	LOOCV	54.5, (52.3, 68.9)	-0.04, (-0.94, 0.12)
		GCV	55.1, (52.3, 65.7)	0.01, (-0.97, 0.88)
	KRR	LOOCV	0.869, (0.753, 1.04)	0.79, (-0.80, 0.97)
		GCV	0.889, (0.762, 1.05)	0.13, (-405.97, 0.96)
Superconductor Critical Temperature	$Kl_\infty R$	LOOCV	52.9, (51.3, 61.0)	-0.04, (-0.90, 0.02)
		GCV	53.3, (51.5, 62.1)	0.47, (-0.19, 0.78)
	KRR	LOOCV	1.04, (0.993, 1.23)	0.54, (-0.46, 0.75)
		GCV	1.06, (0.987, 1.29)	0.13, (-37.16, 0.62)
U.K. Temperature	$Kl_\infty R$	LOOCV	51.2, (48.2, 59.1)	-0.05, (-4.55, 0.00)
		GCV	51.2, (48.3, 58.7)	-0.05, (-1.12, 0.00)
	KRR	LOOCV	0.906, (0.818, 1.13)	0.18, (-2.03, 0.56)
		GCV	0.887, (0.817, 1.14)	-0.01, (-61.69, 0.58)

TABLE 8  
Additional kernels used.

Name	Equation
Matérn, $\nu = \frac{1}{2}$ (Laplace)	$\exp\left(-\frac{\ \mathbf{x}-\mathbf{x}'\ _2}{\sigma}\right)$
Matérn, $\nu = \frac{3}{2}$	$\left(1 + \frac{\sqrt{3} \cdot \ \mathbf{x}-\mathbf{x}'\ _2}{\sigma}\right) \cdot \exp\left(-\frac{\sqrt{3} \cdot \ \mathbf{x}-\mathbf{x}'\ _2}{\sigma}\right)$
Matérn, $\nu = \frac{5}{2}$	$\left(1 + \frac{\sqrt{5} \cdot \ \mathbf{x}-\mathbf{x}'\ _2}{\sigma} + \frac{5 \cdot \ \mathbf{x}-\mathbf{x}'\ _2^2}{3 \cdot \sigma^2}\right) \cdot \exp\left(-\frac{\sqrt{5} \cdot \ \mathbf{x}-\mathbf{x}'\ _2}{\sigma}\right)$
Cauchy	$\left(1 + \frac{\ \mathbf{x}-\mathbf{x}'\ _2^2}{\sigma^2}\right)^{-1}$

TABLE 9

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, *without* amplified outliers, for  $\nu = 0.5$  (Laplace kernel). The five robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	9.81, (4.94, 14.7)	0.49, (−0.39, 0.78)
	$K\ell_\infty R$	383, (372, 416)	0.50, (−0.39, 0.81)
	KMR-H	113, (96.3, 127)	0.52, (−0.50, 0.83)
	KMR-T	84.2, (73.5, 96.8)	0.50, (−0.44, 0.83)
	KQR-A	235, (219, 253)	0.50, (−0.44, 0.83)
	KGD	20.2, (12.7, 26.7)	0.51, (−0.39, 0.84)
California House Values	KRR	2.64, (2.59, 3.24)	0.54, (−0.38, 0.84)
	KSGD	3.58, (2.63, 4.97)	0.62, (−0.09, 0.87)
	$K\ell_\infty R$	331, (312, 361)	0.63, (−0.10, 0.84)
	KMR-H	95.6, (81.9, 118)	0.66, (−0.20, 0.86)
	KMR-T	69.3, (57.7, 86.6)	0.65, (−0.21, 0.87)
	KQR-A	222, (191, 241)	0.67, (−0.18, 0.86)
Steel Energy Consumption	KGD	17.9, (12.4, 25.1)	0.66, (−0.13, 0.85)
	KRR	2.97, (2.64, 4.08)	0.66, (−0.12, 0.85)
	KSGD	17.9, (14.6, 28.4)	0.98, (0.85, 0.99)
	$K\ell_\infty R$	464, (408, 478)	0.98, (0.93, 0.99)
	KMR-H	119, (91.7, 139)	0.99, (0.95, 1.00)
	KMR-T	97.5, (78.7, 107)	0.98, (0.95, 1.00)
Superconductor Critical Temperature	KQR-A	254, (208, 268)	0.99, (0.95, 1.00)
	KGD	26.2, (19.3, 33.4)	0.98, (0.93, 0.99)
	KRR	2.98, (2.64, 3.34)	0.99, (0.95, 1.00)
	KSGD	18.0, (13.1, 22.1)	0.58, (0.04, 0.88)
	$K\ell_\infty R$	443, (404, 491)	0.65, (0.10, 0.90)
	KMR-H	129, (106, 161)	0.65, (−0.01, 0.87)
U.K. Temperature	KMR-T	99.4, (76.0, 124)	0.67, (0.06, 0.89)
	KQR-A	275, (213, 335)	0.65, (0.06, 0.90)
	KGD	15.3, (10.6, 22.8)	0.67, (0.12, 0.90)
	KRR	3.38, (3.16, 4.13)	0.65, (0.14, 0.88)
	KSGD	7.43, (4.34, 11.5)	0.43, (0.04, 0.67)
	$K\ell_\infty R$	382, (363, 420)	0.48, (0.02, 0.70)
U.K. Temperature	KMR-H	118, (96.8, 142)	0.46, (0.01, 0.68)
	KMR-T	87.2, (69.6, 110)	0.46, (0.01, 0.69)
	KQR-A	250, (226, 297)	0.47, (−0.10, 0.69)
	KGD	18.6, (14.2, 25.4)	0.47, (0.01, 0.69)
	KRR	2.74, (2.61, 3.03)	0.46, (0.02, 0.69)

TABLE 10

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, **with** amplified outliers, for  $\nu = 0.5$  (Laplace kernel). The five robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster. The two non-robust methods, KGD/KRR, do not perform as well as the robust methods.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	11.7, (4.46, 15.2)	0.41, (−0.07, 0.70)
	$K\ell_\infty R$	386, (375, 417)	0.43, (−0.73, 0.73)
	KMR-H	82.1, (38.3, 116)	0.41, (−0.39, 0.75)
	KMR-T	70.9, (36.9, 91.1)	0.40, (−0.14, 0.76)
	KQR-A	251, (223, 283)	0.43, (−0.07, 0.74)
	KGD	15.0, (3.83, 22.1)	0.31, (−57.39, 0.75)
California House Values	KRR	2.64, (2.60, 2.89)	0.32, (−1.62, 0.72)
	KSGD	4.05, (2.77, 6.99)	0.49, (−0.20, 0.77)
	$K\ell_\infty R$	344, (326, 382)	0.49, (−0.26, 0.79)
	KMR-H	72.2, (38.5, 111)	0.46, (−0.32, 0.80)
	KMR-T	58.4, (35.8, 79.1)	0.48, (−0.22, 0.80)
	KQR-A	226, (202, 266)	0.48, (−0.30, 0.79)
Steel Energy Consumption	KGD	12.7, (3.83, 18.7)	0.34, (−2.69, 0.80)
	KRR	2.73, (2.66, 3.18)	0.37, (−0.33, 0.80)
	KSGD	22.1, (17.5, 34.1)	0.89, (0.07, 0.97)
	$K\ell_\infty R$	433, (403, 516)	0.84, (−0.76, 0.97)
	KMR-H	81.8, (45.7, 125)	0.87, (−1.19, 0.97)
	KMR-T	72.2, (46.7, 107)	0.88, (−0.07, 0.97)
Superconductor Critical Temperature	KQR-A	263, (230, 330)	0.88, (−1.58, 0.97)
	KGD	11.6, (6.34, 17.5)	0.75, (−31.91, 0.96)
	KRR	2.73, (2.65, 2.92)	0.79, (−0.36, 0.96)
	KSGD	21.2, (15.0, 27.7)	0.47, (−0.10, 0.82)
	$K\ell_\infty R$	445, (400, 476)	0.50, (−0.28, 0.85)
	KMR-H	112, (55.7, 163)	0.48, (−0.38, 0.84)
U.K. Temperature	KMR-T	89.1, (52.6, 158)	0.49, (−0.30, 0.84)
	KQR-A	282, (222, 332)	0.48, (−0.10, 0.84)
	KGD	11.5, (4.98, 17.9)	0.44, (−0.96, 0.82)
	KRR	3.26, (3.13, 3.76)	0.46, (−0.44, 0.81)
	KSGD	9.33, (5.15, 14.9)	0.36, (−0.07, 0.63)
	$K\ell_\infty R$	383, (367, 461)	0.27, (−0.47, 0.63)
U.K. Temperature	KMR-H	78.5, (38.0, 135)	0.32, (−0.10, 0.65)
	KMR-T	65.9, (36.2, 98.6)	0.35, (−0.20, 0.64)
	KQR-A	249, (217, 312)	0.37, (−0.25, 0.69)
	KGD	13.8, (3.81, 22.6)	0.18, (−19.57, 0.62)
	KRR	2.66, (2.62, 3.22)	0.16, (−1.83, 0.63)

TABLE 11

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, *without* amplified outliers, for  $\nu = 1.5$ . The five robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	10.2, (5.42, 14.7)	0.50, (−0.42, 0.83)
	$K\ell_\infty R$	411, (375, 435)	0.53, (−0.52, 0.84)
	KMR-H	184, (145, 208)	0.54, (−0.41, 0.81)
	KMR-T	163, (135, 186)	0.55, (−0.43, 0.85)
	KQR-A	329, (264, 344)	0.52, (−0.42, 0.85)
	KGD	17.4, (11.9, 21.2)	0.56, (−0.50, 0.85)
California House Values	KRR	3.03, (2.95, 3.52)	0.56, (−0.41, 0.85)
	KSGD	3.74, (2.08, 7.04)	0.63, (−0.11, 0.85)
	$K\ell_\infty R$	330, (314, 351)	0.64, (−0.19, 0.89)
	KMR-H	134, (120, 148)	0.62, (−0.19, 0.87)
	KMR-T	110, (98.8, 118)	0.61, (−0.25, 0.87)
	KQR-A	279, (245, 289)	0.61, (−0.07, 0.88)
Steel Energy Consumption	KGD	14.6, (10.6, 18.2)	0.64, (−0.35, 0.89)
	KRR	3.24, (3.06, 3.41)	0.63, (−0.10, 0.85)
	KSGD	15.0, (12.4, 23.4)	0.98, (0.90, 1.00)
	$K\ell_\infty R$	420, (404, 449)	0.98, (0.94, 0.99)
	KMR-H	162, (147, 175)	0.99, (0.97, 1.00)
	KMR-T	145, (133, 157)	0.99, (0.97, 1.00)
Superconductor Critical Temperature	KQR-A	278, (244, 289)	0.99, (0.97, 1.00)
	KGD	22.9, (18.8, 26.2)	0.99, (0.94, 1.00)
	KRR	3.11, (3.04, 3.28)	0.99, (0.97, 1.00)
	KSGD	16.2, (11.4, 21.6)	0.58, (0.07, 0.86)
	$K\ell_\infty R$	415, (399, 495)	0.65, (0.17, 0.91)
	KMR-H	179, (142, 228)	0.64, (0.16, 0.90)
U.K. Temperature	KMR-T	156, (119, 204)	0.65, (0.14, 0.88)
	KQR-A	306, (252, 373)	0.61, (−0.00, 0.90)
	KGD	16.3, (12.1, 22.4)	0.68, (0.09, 0.90)
	KRR	4.06, (3.53, 4.52)	0.65, (0.13, 0.89)
	KSGD	6.36, (3.65, 9.72)	0.42, (0.01, 0.72)
	$K\ell_\infty R$	367, (358, 379)	0.43, (−0.11, 0.71)
U.K. Temperature	KMR-H	162, (144, 181)	0.44, (−0.27, 0.68)
	KMR-T	148, (129, 164)	0.44, (−0.25, 0.72)
	KQR-A	285, (259, 333)	0.39, (−0.23, 0.70)
	KGD	12.8, (9.70, 18.0)	0.43, (−0.20, 0.71)
	KRR	3.05, (2.98, 3.51)	0.46, (−0.25, 0.68)

TABLE 12

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, **with** amplified outliers, for  $\nu = 1.5$ . The five robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster. The two non-robust methods, KGD/KRR, do not perform as well as the robust methods.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	10.9, (6.35, 14.8)	0.40, (−0.25, 0.74)
	$K\ell_\infty R$	387, (378, 412)	0.44, (−0.77, 0.74)
	KMR-H	141, (82.1, 169)	0.39, (−0.14, 0.74)
	KMR-T	131, (93.0, 153)	0.41, (−0.10, 0.77)
	KQR-A	289, (271, 304)	0.39, (−0.15, 0.75)
	KGD	13.6, (3.52, 18.5)	0.25, (−80.63, 0.72)
California House Values	KRR	3.03, (2.99, 4.03)	0.27, (−1.61, 0.71)
	KSGD	4.78, (2.53, 8.28)	0.48, (−0.04, 0.81)
	$K\ell_\infty R$	363, (326, 394)	0.48, (−0.55, 0.80)
	KMR-H	109, (69.8, 146)	0.44, (−0.62, 0.81)
	KMR-T	98.2, (75.5, 127)	0.48, (−0.54, 0.82)
	KQR-A	290, (245, 326)	0.38, (−0.81, 0.83)
Steel Energy Consumption	KGD	11.9, (5.15, 14.8)	0.24, (−4.82, 0.75)
	KRR	3.12, (3.06, 3.47)	0.33, (−0.28, 0.78)
	KSGD	20.8, (15.9, 31.9)	0.89, (0.06, 0.97)
	$K\ell_\infty R$	495, (413, 530)	0.83, (−0.59, 0.98)
	KMR-H	161, (123, 211)	0.88, (−1.32, 0.98)
	KMR-T	169, (134, 204)	0.87, (−0.02, 0.98)
Superconductor Critical Temperature	KQR-A	344, (264, 373)	0.87, (−1.30, 0.98)
	KGD	13.8, (6.79, 18.3)	0.74, (−28.49, 0.97)
	KRR	3.09, (3.02, 3.75)	0.80, (−0.03, 0.97)
	KSGD	18.1, (13.9, 23.1)	0.46, (−0.29, 0.83)
	$K\ell_\infty R$	426, (409, 472)	0.48, (−0.50, 0.85)
	KMR-H	156, (95.1, 201)	0.49, (−0.43, 0.82)
U.K. Temperature	KMR-T	137, (96.7, 200)	0.50, (−0.44, 0.86)
	KQR-A	289, (247, 325)	0.44, (−0.55, 0.83)
	KGD	12.5, (4.72, 16.7)	0.37, (−1.33, 0.83)
	KRR	3.61, (3.56, 4.25)	0.46, (−0.46, 0.82)
	KSGD	9.54, (4.72, 14.7)	0.35, (−0.08, 0.63)
	$K\ell_\infty R$	460, (367, 491)	0.27, (−0.60, 0.62)
U.K. Temperature	KMR-H	153, (92.4, 226)	0.33, (−0.70, 0.66)
	KMR-T	146, (103, 204)	0.30, (−0.71, 0.65)
	KQR-A	363, (266, 393)	0.32, (−0.88, 0.66)
	KGD	11.6, (4.35, 17.7)	0.14, (−24.32, 0.61)
	KRR	3.11, (3.01, 4.03)	0.21, (−1.45, 0.61)

TABLE 13

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, *without* amplified outliers, for  $\nu = 2.5$ . The five robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	9.84, (5.66, 15.7)	0.50, (−0.51, 0.82)
	$K\ell_\infty R$	409, (371, 455)	0.53, (−0.50, 0.85)
	KMR-H	190, (152, 209)	0.49, (−0.61, 0.82)
	KMR-T	177, (144, 197)	0.52, (−0.43, 0.82)
	KQR-A	327, (270, 352)	0.51, (−0.60, 0.83)
	KGD	16.7, (11.5, 20.6)	0.56, (−0.50, 0.86)
California House Values	KRR	3.27, (3.18, 3.68)	0.54, (−0.43, 0.85)
	KSGD	4.35, (2.34, 7.86)	0.60, (0.02, 0.86)
	$K\ell_\infty R$	343, (317, 383)	0.63, (−0.11, 0.89)
	KMR-H	155, (122, 176)	0.62, (−0.18, 0.87)
	KMR-T	130, (106, 147)	0.59, (−0.29, 0.87)
	KQR-A	301, (255, 328)	0.60, (−0.07, 0.88)
Steel Energy Consumption	KGD	14.4, (10.4, 19.0)	0.63, (−0.14, 0.89)
	KRR	3.56, (3.29, 3.84)	0.62, (−0.06, 0.85)
	KSGD	14.7, (12.2, 23.6)	0.98, (0.91, 0.99)
	$K\ell_\infty R$	439, (406, 481)	0.98, (0.95, 1.00)
	KMR-H	178, (151, 200)	0.99, (0.98, 1.00)
	KMR-T	167, (147, 178)	0.99, (0.97, 1.00)
Superconductor Critical Temperature	KQR-A	290, (251, 311)	0.99, (0.98, 1.00)
	KGD	23.1, (18.6, 27.0)	0.99, (0.95, 1.00)
	KRR	3.36, (3.27, 3.64)	0.99, (0.98, 1.00)
	KSGD	14.9, (10.9, 19.9)	0.58, (0.01, 0.87)
	$K\ell_\infty R$	431, (406, 491)	0.63, (0.16, 0.91)
	KMR-H	185, (158, 231)	0.64, (0.05, 0.91)
U.K. Temperature	KMR-T	166, (135, 209)	0.64, (0.10, 0.88)
	KQR-A	315, (264, 366)	0.62, (−0.05, 0.89)
	KGD	15.5, (12.6, 20.5)	0.68, (0.16, 0.90)
	KRR	4.15, (3.77, 4.42)	0.65, (0.13, 0.89)
	KSGD	7.07, (3.47, 11.4)	0.40, (−0.01, 0.73)
	$K\ell_\infty R$	394, (362, 460)	0.41, (−0.14, 0.70)
U.K. Temperature	KMR-H	183, (162, 228)	0.41, (−0.58, 0.66)
	KMR-T	177, (152, 205)	0.43, (−0.58, 0.68)
	KQR-A	314, (284, 380)	0.35, (−0.95, 0.69)
	KGD	12.6, (9.63, 17.3)	0.42, (−0.22, 0.67)
	KRR	3.53, (3.24, 4.04)	0.42, (−0.58, 0.66)

TABLE 14

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, **with** amplified outliers, for  $\nu = 2.5$ . The five robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster. The two non-robust methods, KGD/KRR, do not perform as well as the robust methods.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	10.8, (6.34, 14.3)	0.41, (−0.25, 0.71)
	$K\ell_\infty R$	389, (376, 410)	0.43, (−0.67, 0.75)
	KMR-H	150, (90.2, 177)	0.39, (−0.52, 0.68)
	KMR-T	140, (105, 160)	0.37, (−0.57, 0.79)
	KQR-A	292, (272, 310)	0.39, (−0.86, 0.70)
	KGD	13.4, (3.63, 17.6)	0.23, (−92.96, 0.68)
California House Values	KRR	3.26, (3.19, 3.35)	0.27, (−1.61, 0.70)
	KSGD	4.89, (2.78, 10.2)	0.46, (−0.05, 0.82)
	$K\ell_\infty R$	352, (326, 391)	0.43, (−0.41, 0.80)
	KMR-H	119, (80.8, 157)	0.40, (−0.70, 0.81)
	KMR-T	114, (88.8, 133)	0.44, (−0.55, 0.80)
	KQR-A	282, (266, 328)	0.34, (−0.78, 0.82)
Steel Energy Consumption	KGD	11.8, (5.21, 15.1)	0.20, (−3.93, 0.79)
	KRR	3.34, (3.29, 4.09)	0.31, (−0.27, 0.79)
	KSGD	17.4, (13.2, 25.6)	0.90, (0.06, 0.97)
	$K\ell_\infty R$	416, (401, 573)	0.85, (−0.70, 0.98)
	KMR-H	149, (116, 173)	0.87, (−1.92, 0.98)
	KMR-T	151, (132, 167)	0.87, (−0.02, 0.98)
Superconductor Critical Temperature	KQR-A	275, (259, 311)	0.86, (−2.41, 0.98)
	KGD	13.8, (6.41, 19.5)	0.76, (−31.99, 0.97)
	KRR	3.34, (3.30, 3.59)	0.80, (−5.16, 0.97)
	KSGD	16.6, (12.4, 22.5)	0.46, (−0.40, 0.82)
	$K\ell_\infty R$	417, (399, 478)	0.49, (−0.36, 0.85)
	KMR-H	163, (99.4, 211)	0.49, (−0.45, 0.80)
U.K. Temperature	KMR-T	146, (105, 204)	0.49, (−0.41, 0.85)
	KQR-A	287, (258, 370)	0.48, (−0.47, 0.82)
	KGD	12.3, (4.91, 17.2)	0.39, (−1.28, 0.83)
	KRR	3.84, (3.76, 4.70)	0.47, (−0.51, 0.80)
	KSGD	8.31, (4.60, 12.8)	0.33, (−0.12, 0.62)
	$K\ell_\infty R$	384, (368, 455)	0.25, (−1.11, 0.60)
U.K. Temperature	KMR-H	149, (98.7, 196)	0.27, (−2.30, 0.65)
	KMR-T	139, (107, 179)	0.30, (−2.36, 0.66)
	KQR-A	291, (272, 354)	0.26, (−1.52, 0.69)
	KGD	10.8, (3.44, 17.3)	0.11, (−20.02, 0.61)
	KRR	3.28, (3.22, 3.88)	0.19, (−2.27, 0.61)



TABLE 15

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, *without* amplified outliers, for the Cauchy kernel. The five robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	10.6, (6.23, 15.9)	0.49, (−0.46, 0.81)
	$K\ell_\infty$ R	403, (375, 445)	0.53, (−0.54, 0.82)
	KMR-H	189, (156, 209)	0.48, (−0.59, 0.84)
	KMR-T	177, (145, 196)	0.52, (−0.43, 0.82)
	KQR-A	326, (273, 348)	0.53, (−0.53, 0.83)
	KGD	17.7, (12.2, 22.6)	0.54, (−0.51, 0.87)
California House Values	KRR	2.26, (2.20, 2.57)	0.52, (−0.42, 0.84)
	KSGD	4.83, (2.93, 7.88)	0.63, (0.02, 0.89)
	$K\ell_\infty$ R	352, (313, 405)	0.64, (−0.19, 0.89)
	KMR-H	154, (131, 175)	0.60, (−0.13, 0.87)
	KMR-T	122, (106, 145)	0.59, (−0.31, 0.86)
	KQR-A	288, (253, 334)	0.56, (−0.21, 0.88)
Steel Energy Consumption	KGD	17.2, (13.5, 23.2)	0.62, (−0.26, 0.88)
	KRR	2.51, (2.29, 2.88)	0.61, (−0.11, 0.87)
	KSGD	21.4, (16.4, 28.5)	0.98, (0.89, 0.99)
	$K\ell_\infty$ R	454, (421, 518)	0.98, (0.94, 0.99)
	KMR-H	186, (150, 208)	1.00, (0.98, 1.00)
	KMR-T	176, (145, 195)	0.99, (0.97, 1.00)
Superconductor Critical Temperature	KQR-A	302, (258, 331)	0.99, (0.97, 1.00)
	KGD	24.4, (20.1, 29.0)	0.99, (0.94, 1.00)
	KRR	2.36, (2.26, 2.68)	0.99, (0.97, 1.00)
	KSGD	23.7, (17.4, 31.2)	0.57, (0.05, 0.86)
	$K\ell_\infty$ R	465, (400, 500)	0.66, (0.19, 0.91)
	KMR-H	188, (147, 231)	0.63, (0.10, 0.90)
U.K. Temperature	KMR-T	165, (130, 209)	0.64, (0.18, 0.88)
	KQR-A	313, (253, 351)	0.59, (−0.01, 0.89)
	KGD	18.3, (13.6, 24.6)	0.68, (0.10, 0.89)
	KRR	2.76, (2.64, 3.25)	0.65, (0.17, 0.89)
	KSGD	7.46, (4.06, 11.4)	0.44, (−0.02, 0.72)
	$K\ell_\infty$ R	410, (364, 421)	0.42, (−0.14, 0.69)
U.K. Temperature	KMR-H	198, (158, 220)	0.42, (−1.12, 0.65)
	KMR-T	189, (144, 202)	0.44, (−1.11, 0.67)
	KQR-A	338, (269, 356)	0.37, (−1.22, 0.69)
	KGD	14.6, (9.97, 19.1)	0.41, (−0.18, 0.70)
	KRR	2.37, (2.25, 2.65)	0.45, (−1.12, 0.67)

TABLE 16

The 2.5th, 50th, and 97.5th percentiles of computation time and test  $R^2$  for the different methods and data sets, **with** amplified outliers, for the Cauchy kernel. The five robust methods perform very similarly in terms of test  $R^2$ , while KSGD performs one to two orders of magnitude faster. The two non-robust methods, KGD/KRR, do not perform as well as the robust methods.

Data	Method	Computation Time [s] 50%, (2.5%, 97.5%)	Test $R^2$ 50%, (2.5%, 97.5%)
Airfoil Sound Pressure	KSGD	12.6, (7.98, 17.6)	0.42, (−0.05, 0.74)
	$K\ell_\infty R$	386, (371, 411)	0.42, (−0.79, 0.75)
	KMR-H	151, (94.5, 178)	0.41, (−0.85, 0.68)
	KMR-T	140, (105, 158)	0.33, (−0.60, 0.81)
	KQR-A	282, (265, 310)	0.42, (−1.13, 0.75)
	KGD	14.2, (3.66, 19.7)	0.23, (−83.22, 0.70)
California House Values	KRR	2.24, (2.20, 2.30)	0.28, (−2.40, 0.70)
	KSGD	5.59, (3.24, 9.36)	0.49, (−0.05, 0.81)
	$K\ell_\infty R$	338, (322, 362)	0.46, (−0.54, 0.78)
	KMR-H	112, (69.8, 143)	0.40, (−0.76, 0.80)
	KMR-T	107, (77.4, 121)	0.41, (−0.59, 0.81)
	KQR-A	268, (251, 291)	0.29, (−0.80, 0.82)
Steel Energy Consumption	KGD	13.9, (6.24, 16.8)	0.20, (−3.98, 0.77)
	KRR	2.37, (2.28, 2.42)	0.30, (−0.21, 0.76)
	KSGD	22.7, (19.6, 34.0)	0.89, (0.06, 0.97)
	$K\ell_\infty R$	416, (402, 485)	0.83, (−0.75, 0.98)
	KMR-H	152, (113, 199)	0.86, (−1.63, 0.98)
	KMR-T	152, (127, 184)	0.86, (−0.02, 0.98)
Superconductor Critical Temperature	KQR-A	272, (252, 320)	0.88, (−1.72, 0.98)
	KGD	14.0, (6.36, 18.6)	0.74, (−34.70, 0.97)
	KRR	2.31, (2.20, 2.39)	0.80, (−5.36, 0.97)
	KSGD	27.6, (22.0, 34.7)	0.47, (−0.09, 0.82)
	$K\ell_\infty R$	469, (406, 502)	0.47, (−0.39, 0.84)
	KMR-H	182, (111, 236)	0.49, (−0.39, 0.84)
U.K. Temperature	KMR-T	173, (114, 250)	0.49, (−0.54, 0.85)
	KQR-A	329, (270, 358)	0.48, (−0.87, 0.83)
	KGD	15.4, (8.11, 19.6)	0.37, (−1.24, 0.83)
	KRR	2.73, (2.59, 3.18)	0.46, (−0.49, 0.81)
	KSGD	9.66, (5.27, 17.3)	0.38, (−0.08, 0.62)
	$K\ell_\infty R$	428, (373, 456)	0.29, (−0.93, 0.63)
U.K. Temperature	KMR-H	163, (111, 225)	0.28, (−3.21, 0.66)
	KMR-T	152, (115, 202)	0.29, (−3.33, 0.66)
	KQR-A	338, (273, 369)	0.22, (−1.46, 0.64)
	KGD	12.3, (3.28, 19.1)	0.10, (−19.91, 0.59)
	KRR	2.33, (2.23, 2.65)	0.14, (−2.69, 0.59)

## Appendix D: Proofs

*Proof of Proposition 3.1.* Differentiating Equation 5b with respect to  $\alpha$  and setting the gradient to  $\mathbf{0}$ , we obtain

$$\begin{aligned} \mathbf{0} &= \frac{\partial}{\partial \hat{\alpha}} \left( \frac{1}{2} \|\mathbf{y} - \mathbf{K} \hat{\alpha}\|_{\mathbf{K}^{-1}}^2 + \frac{\lambda}{2} \|\hat{\alpha}\|_2^2 \right) = -\mathbf{y} + \mathbf{K} \hat{\alpha} + \lambda \hat{\alpha} = -\mathbf{y} + (\mathbf{K} + \lambda \mathbf{I}) \hat{\alpha} \\ &\iff \hat{\alpha} = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{y}. \quad \square \end{aligned}$$

*Proof of Equation 10.* With  $\hat{\eta}(t) = \mathbf{K} \hat{\alpha}(t) - \mathbf{y}$ , we obtain  $\frac{d\hat{\eta}(t)}{dt} = \mathbf{K} \frac{d\hat{\alpha}(t)}{dt}$  and Equation 9 can be written as

$$\frac{d\hat{\eta}(t)}{dt} = -\mathbf{K} \hat{\eta}(t) \iff \hat{\eta}(t) = \exp(-t\mathbf{K}) \hat{\eta}_0.$$

Now,

$$\hat{\alpha}(0) = \mathbf{0} \implies \hat{\eta}_0 = \hat{\eta}(0) = -\mathbf{y} \implies \hat{\eta}(t) = -\exp(-t\mathbf{K}) \mathbf{y}.$$

Solving for  $\hat{\alpha}(t)$ , we obtain

$$\hat{\alpha}(t) = \mathbf{K}^{-1} (\mathbf{I} - \exp(-t\mathbf{K})) \mathbf{y} = (\mathbf{I} - \exp(-t\mathbf{K})) \mathbf{K}^{-1} \mathbf{y}. \quad \square$$

*Proof of Remark 4.2.* The differential equation in Equation 9 can be obtained by writing

$$\hat{\alpha}_{k+1} = \hat{\alpha}_k + \eta \cdot (\mathbf{y} - \mathbf{K} \hat{\alpha}_k)$$

as

$$\hat{\alpha}(t + \Delta t) = \hat{\alpha}(t) + \Delta t \cdot (\mathbf{y} - \mathbf{K} \hat{\alpha}(t)),$$

rearranging and letting  $\Delta t \rightarrow 0$ , to obtain

$$\frac{d\hat{\alpha}(t)}{dt} = \lim_{\Delta t \rightarrow 0} \left( \frac{\hat{\alpha}(t + \Delta t) - \hat{\alpha}(t)}{\Delta t} \right) = \mathbf{y} - \mathbf{K} \hat{\alpha}(t).$$

For momentum and Nesterov accelerated gradient, the update rule for  $\hat{\alpha}$  in Equation 8 generalizes to

$$\begin{aligned} \hat{\alpha}(t + \Delta t) &= \hat{\alpha}(t) + \gamma \cdot (\hat{\alpha}(t) - \hat{\alpha}(t - \Delta t)) \\ &\quad + \Delta t \cdot (\mathbf{y} - \mathbf{K} (\hat{\alpha}(t) + \mathbb{I}_{\text{NAG}} \cdot \gamma \cdot (\hat{\alpha}(t) - \hat{\alpha}(t - \Delta t)))), \end{aligned}$$

where  $\gamma \in [0, 1)$  is the strength of the momentum and  $\mathbb{I}_{\text{NAG}} \in \{0, 1\}$  decides whether to use Nesterov accelerated gradient or not. Rearranging, we obtain

$$\begin{aligned} \frac{\hat{\alpha}(t + \Delta t) - \hat{\alpha}(t)}{\Delta t} &- \gamma \cdot \frac{\hat{\alpha}(t) - \hat{\alpha}(t - \Delta t)}{\Delta t} = \\ \mathbf{y} - \mathbf{K} (\hat{\alpha}(t) + \mathbb{I}_{\text{NAG}} \cdot \gamma \cdot (\hat{\alpha}(t) - \hat{\alpha}(t - \Delta t))), \end{aligned}$$

and, when  $\Delta t \rightarrow 0$ ,

$$(1 - \gamma) \cdot \frac{d\hat{\alpha}(t)}{dt} = \mathbf{y} - \mathbf{K} \hat{\alpha}(t).$$

Solving the differential equations in the same way as for Equation 10, we obtain

$$\hat{\alpha}(t) = \left( \mathbf{I} - \exp \left( -\frac{t}{1 - \gamma} \mathbf{K} \right) \right) \mathbf{K}^{-1} \mathbf{y}. \quad \square$$

*Proof of Proposition 4.1.* By using the fact that  $\mathbf{K}$ ,  $\mathbf{K}^{-1}$ , and  $\Delta_{\mathbf{K}} := \left( (\mathbf{I} + t\mathbf{K})^{-1} - \exp(-t\mathbf{K}) \right)$  commute, we obtain

$$\begin{aligned}\hat{\alpha}_{\text{KGF}}(t) - \hat{\alpha}_{\text{KRR}}(1/t) &= (\mathbf{I} - \exp(-t\mathbf{K})) \mathbf{K}^{-1} \mathbf{y} - \left( \mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1} \right) \mathbf{K}^{-1} \mathbf{y} \\ &= \underbrace{\left( (\mathbf{I} + t\mathbf{K})^{-1} - \exp(-t\mathbf{K}) \right)}_{=:\Delta_{\mathbf{K}}} \mathbf{K}^{-1} \mathbf{y} = \Delta_{\mathbf{K}} \mathbf{K}^{-1} \mathbf{y}, \\ \hat{f}_{\text{KGF}}(\mathbf{X}, t) - \hat{f}_{\text{KRR}}(\mathbf{X}, 1/t) &= \mathbf{K} (\hat{\alpha}_{\text{KGF}} - \hat{\alpha}_{\text{KRR}}) = \mathbf{K} \Delta_{\mathbf{K}} \mathbf{K}^{-1} \mathbf{y} \\ &= \Delta_{\mathbf{K}} \mathbf{K} \mathbf{K}^{-1} \mathbf{y} = \Delta_{\mathbf{K}} \mathbf{y}, \\ \hat{f}_{\text{KGF}}(\mathbf{x}^*, t) - \hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t) &= \mathbf{k}(\mathbf{x}^*)^\top (\hat{\alpha}_{\text{KGF}} - \hat{\alpha}_{\text{KRR}}) \\ &= \mathbf{k}(\mathbf{x}^*)^\top \Delta_{\mathbf{K}} \mathbf{K}^{-1} \mathbf{y} = \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \Delta_{\mathbf{K}} \mathbf{y},\end{aligned}$$

We denote the singular value decomposition of the symmetric matrix  $\mathbf{K}$  by  $\mathbf{K} = \mathbf{U} \mathbf{S} \mathbf{U}^\top$ , where  $\mathbf{S}$  is a diagonal matrix with diagonal elements  $S_{ii} = s_i$ . Then  $\Delta_{\mathbf{K}} = \mathbf{U} \mathbf{D} \mathbf{U}^\top$ , where  $\mathbf{D}$  is a diagonal matrix with entries

$$D_{ii} = d_i = \frac{1}{1 + t \cdot s_i} - e^{t \cdot s_i}.$$

Since for any vector  $\mathbf{x}$  and any diagonal matrix  $\mathbf{D}$ ,

$$\mathbf{x}^\top \mathbf{D} \mathbf{x} = \sum_i x_i^2 \cdot d_i \leq \sum_i x_i^2 \cdot \max_i d_i = \mathbf{x}^\top \mathbf{x} \cdot \max_i d_i, \quad (22)$$

and since  $\mathbf{K}^{-1}$  and  $\Delta_{\mathbf{K}}$  are both symmetric,

$$\begin{aligned}\|\hat{\alpha}_{\text{KGF}}(t) - \hat{\alpha}_{\text{KRR}}(1/t)\|_2^2 &= \mathbf{y}^\top \mathbf{K}^{-1} \Delta_{\mathbf{K}} \Delta_{\mathbf{K}} \mathbf{K}^{-1} \mathbf{y} = \mathbf{y}^\top \mathbf{K}^{-1} \mathbf{U} \mathbf{D}^2 \mathbf{U}^\top \mathbf{K}^{-1} \mathbf{y} \\ &\leq \mathbf{y}^\top \mathbf{K}^{-1} \underbrace{\mathbf{U} \mathbf{U}^\top}_{=\mathbf{I}} \mathbf{K}^{-1} \mathbf{y} \cdot \max_i d_i^2 = \|\mathbf{K}^{-1} \mathbf{y}\|_2^2 \cdot \max_i d_i^2,\end{aligned}$$

and

$$\begin{aligned}\|\hat{f}_{\text{KGF}}(\mathbf{X}, t) - \hat{f}_{\text{KRR}}(\mathbf{X}, 1/t)\|_2^2 &= \mathbf{y}^\top \Delta_{\mathbf{K}} \Delta_{\mathbf{K}} \mathbf{y} = \mathbf{y}^\top \mathbf{U} \mathbf{D}^2 \mathbf{U}^\top \mathbf{y} \\ &\leq \mathbf{y}^\top \underbrace{\mathbf{U} \mathbf{U}^\top}_{=\mathbf{I}} \mathbf{y} \cdot \max_i d_i^2 = \|\mathbf{y}\|_2^2 \cdot \max_i d_i^2.\end{aligned}$$

For out-of-sample predictions, calculations become different, since  $\mathbf{k}(\mathbf{x}^*)$  and  $\mathbf{K}$  do not commute. Hence we need to take expectation over  $\mathbf{y}$ . We first note that for  $\mathbf{y} = \mathbf{K} \alpha_0 + \varepsilon$ , where  $\alpha_0 \sim (\mathbf{0}, \Sigma_\alpha)$ , and  $\varepsilon \sim (\mathbf{0}, \sigma_\varepsilon^2 \mathbf{I})$ , and where  $\mathbf{K}$  and  $\Sigma_\alpha$  commute,

$$\mathbb{E}_{\varepsilon, \alpha_0} (\mathbf{y} \mathbf{y}^\top) = \mathbf{K} \mathbb{E} (\alpha_0 \alpha_0^\top) \mathbf{K} + 2\mathbf{K} \mathbb{E} (\alpha_0) \mathbb{E} (\varepsilon)^\top + \mathbb{E} (\varepsilon \varepsilon^\top) = \Sigma_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}.$$

Now,

$$\mathbb{E}_{\varepsilon, \alpha_0} \left( \left( \hat{f}_{\text{KGF}}(\mathbf{x}^*, t) - \hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t) \right)^2 \right)$$

$$\begin{aligned}
&= \mathbb{E}_{\varepsilon, \alpha_0} (\mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \Delta_K \mathbf{y} \mathbf{y}^\top \Delta_K \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}^*)) \\
&= \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \Delta_K (\Sigma_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) \Delta_K \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}^*) \\
&= \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \Sigma_\alpha^{1/2} \mathbf{K} \mathbf{U} \mathbf{D}^2 \mathbf{U}^\top \mathbf{K} \Sigma_\alpha^{1/2} \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}^*) \\
&\quad + \sigma_\varepsilon^2 \cdot \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \mathbf{U} \mathbf{D}^2 \mathbf{U}^\top \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}^*) \\
&\leq \left( \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \Sigma_\alpha^{1/2} \mathbf{K} \underbrace{\mathbf{U} \mathbf{U}^\top}_{=\mathbf{I}} \mathbf{K} \Sigma_\alpha^{1/2} \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}^*) \right. \\
&\quad \left. + \sigma_\varepsilon^2 \cdot \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \underbrace{\mathbf{U} \mathbf{U}^\top}_{=\mathbf{I}} \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}^*) \right) \cdot \max_i d_i^2 \\
&= \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} (\Sigma_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}^*) \cdot \max_i d_i^2 \\
&= \mathbb{E}_{\varepsilon, \alpha_0} (\mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \mathbf{y} \mathbf{y}^\top \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}^*)) \cdot \max_i d_i^2 \\
&= \mathbb{E}_{\varepsilon, \alpha_0} ((\mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-1} \mathbf{y})^2) \cdot \max_i d_i^2 = \mathbb{E}_{\varepsilon, \alpha_0} (\hat{f}^0(\mathbf{x}^*)^2) \cdot \max_i d_i^2,
\end{aligned}$$

where we have again used Equation 22, and the fact that  $\mathbf{K}$ ,  $\Sigma_\alpha^{1/2}$  and  $\Delta_K$  commute.

If we repeat the calculations with  $\mathbf{k}(\mathbf{x}^*)$  replaced by  $\mathbf{e}_i \in \mathbb{R}^n$ , where  $(\mathbf{e}_i)_i = 1$  and remaining elements equal 0, so that  $\hat{\alpha}_i = \mathbf{e}_i^\top \hat{\boldsymbol{\alpha}}$ , we obtain

$$\begin{aligned}
\mathbb{E}_{\varepsilon, \alpha_0} \left( (\hat{\alpha}_i^{\text{KGF}}(t) - \hat{\alpha}_i^{\text{KRR}}(1/t))^2 \right) x_i d_i^2 &= \mathbb{E}_{\varepsilon, \alpha_0} ((\mathbf{e}_i^\top \mathbf{K}^{-1} \mathbf{y})^2) \cdot \max_i d_i^2 \\
&= \mathbb{E}_{\varepsilon, \alpha_0} ((\hat{\alpha}_i^0)^2) \cdot \max_i d_i^2.
\end{aligned}$$

What is left to do, is to calculate  $\max_i d_i^2$ . Let  $x = t \cdot s_i$ . Then, for  $x \geq 0$  (which holds, since both  $t$  and  $s_i$  are non-negative),  $d_i^2 = \left( \frac{1}{1+x} - e^{-x} \right)^2$  is obviously larger than 0, and equals 0 for  $x = 0$  and  $x = +\infty$ . To find the maximum of the expression we differentiate and set the derivative to zero.

$$\frac{\partial}{\partial x} \left( \left( \frac{1}{1+x} - e^{-x} \right)^2 \right) = 2 \left( e^{-x} - \frac{1}{(1+x)^2} \right) \left( \frac{1}{1+x} - e^{-x} \right).$$

For  $i \in \{1, 2\}$  we obtain

$$\begin{aligned}
e^{-x} - \frac{1}{(1+x)^i} = 0 &\iff -\frac{1+x}{i} e^{-\frac{1+x}{i}} = -\frac{e^{-\frac{1}{i}}}{i} \\
\iff -\frac{1+x}{i} = W_k \left( -\frac{e^{-\frac{1}{i}}}{i} \right) &\iff x = -1 - i \cdot W_k \left( -\frac{e^{-\frac{1}{i}}}{i} \right)
\end{aligned}$$

where  $W_k$  is the  $k$ -th branch of the Lambert W function, for  $k \in \{-1, 0\}$ . The only combination of  $i$  and  $k$  that yields an  $x \neq 0$  is  $i = 2$ ,  $k = -1$ , for which we obtain  $x \approx 2.51287$  and  $\left( \frac{1}{1+x} - e^{-x} \right)^2 \Big|_{x=2.51287} \approx 0.04166$ .  $\square$

*Proof of Proposition 4.2.* In all four parts, we will use the inequality  $e^x \geq 1 + x$  and the equalities  $\max_{\mathbf{v} \neq \mathbf{0}} \frac{\|\mathbf{A}\mathbf{v}\|_2}{\|\mathbf{v}\|_2} = \|\mathbf{A}\|_2 = s_{\max}(\mathbf{A})$ , where  $s_{\max}(\mathbf{A})$  denotes the largest singular value. We will denote the largest and smallest singular values of  $\mathbf{K}$  by  $s_{\max} = s_{\max}(\mathbf{K})$  and  $s_{\min} = s_{\min}(\mathbf{K})$ .

For part (a), we have

$$\begin{aligned} \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t) - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) &= \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\mathbf{K}\mathbf{K}^{-1}(\mathbf{I} - \exp(-t\mathbf{K}))\mathbf{y} - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) \\ &= \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\exp(-t\mathbf{K})\mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) = \|\exp(-t\mathbf{K})\|_2 = e^{-ts_{\min}}, \end{aligned}$$

where  $s_{\min}$  denotes the smallest singular value of  $\mathbf{K}$ .

Similarly,

$$\begin{aligned} \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, 1/t) - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) &= \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\mathbf{K}\mathbf{K}^{-1}(\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1})\mathbf{y} - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) \\ &= \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|(\mathbf{I} + t\mathbf{K})^{-1}\mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) = \|(\mathbf{I} + t\mathbf{K})^{-1}\|_2 = \frac{1}{1 + ts_{\min}}, \end{aligned}$$

where  $e^{-ts_{\min}} \leq \frac{1}{1 + ts_{\min}}$ , since  $e^x \geq 1 + x$ .

For part (b),

$$\begin{aligned} \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, 1/t)\|_2}{\|\mathbf{y}\|_2} \right) &= \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\mathbf{K}\mathbf{K}^{-1}(\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1})\mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) \\ &= \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|(\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1})\mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) = \|\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1}\|_2 = 1 - \frac{1}{1 + ts_{\max}} \end{aligned}$$

and

$$\begin{aligned} \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t)\|_2}{\|\mathbf{y}\|_2} \right) &= \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|\mathbf{K}\mathbf{K}^{-1}(\mathbf{I} - \exp(-t\mathbf{K}))\mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) \\ &= \max_{\mathbf{y} \neq \mathbf{0}} \left( \frac{\|(\mathbf{I} - \exp(-t\mathbf{K}))\mathbf{y}\|_2}{\|\mathbf{y}\|_2} \right) = \|\mathbf{I} - \exp(-t\mathbf{K})\|_2 = 1 - e^{-ts_{\max}}, \end{aligned}$$

where  $1 - \frac{1}{1 + ts_{\max}} \leq 1 - e^{-ts_{\max}}$ , since  $e^x \geq 1 + x$ .

For parts (c) and (d), we first note that for  $\mathbf{y} = \mathbf{K}\boldsymbol{\alpha}_0 + \boldsymbol{\varepsilon}$ , where  $\boldsymbol{\alpha}_0 \sim (\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\alpha}})$ , and  $\boldsymbol{\varepsilon} \sim (\mathbf{0}, \sigma_{\varepsilon}^2 \mathbf{I})$ , and where  $\mathbf{K}$  and  $\boldsymbol{\Sigma}_{\boldsymbol{\alpha}}$  commute,

$$\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (\mathbf{y}\mathbf{y}^{\top}) = \mathbf{K}\mathbb{E}(\boldsymbol{\alpha}_0\boldsymbol{\alpha}_0^{\top})\mathbf{K} + 2\mathbf{K}\mathbb{E}(\boldsymbol{\alpha}_0)\mathbb{E}(\boldsymbol{\varepsilon})^{\top} + \mathbb{E}(\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}^{\top}) = \boldsymbol{\Sigma}_{\boldsymbol{\alpha}}\mathbf{K}^2 + \sigma_{\varepsilon}^2\mathbf{I},$$

and that this quantity commutes with other functions of  $\mathbf{K}$ . Furthermore, let  $\mathbf{e}_i \in \mathbb{R}^n$ , denote the  $i$ -th base vector, where  $(\mathbf{e}_i)_i = 1$  and remaining elements equal 0, so that  $\mathbf{e}_i^{\top} \mathbf{v} = v_i$ .

For part (c),

$$\begin{aligned}
& \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \hat{f}_{\text{KGF}}(\mathbf{x}_i, t) - y_i \right)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \mathbf{e}_i^\top \exp(-t\mathbf{K}) \mathbf{y} \right)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \mathbf{e}_i^\top \mathbf{y} \right)^2 \right)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \mathbf{e}_i^\top \exp(-t\mathbf{K}) \mathbf{y} \mathbf{y}^\top \exp(-t\mathbf{K}) \mathbf{e}_i \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \mathbf{e}_i^\top \mathbf{y} \mathbf{y}^\top \mathbf{e}_i \right)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbf{e}_i^\top \exp(-t\mathbf{K}) (\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) \exp(-t\mathbf{K}) \mathbf{e}_i}{\mathbf{e}_i^\top (\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) \mathbf{e}_i} \right) \\
&= \max_{\|\sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} \mathbf{e}_i\|_2 \neq 0} \left( \frac{\left( \left\| \exp(-t\mathbf{K}) \sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} \mathbf{e}_i \right\|_2 \right)^2}{\left\| \sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} \mathbf{e}_i \right\|_2^2} \right) = e^{-2ts_{\min}}
\end{aligned}$$

and

$$\begin{aligned}
& \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \hat{f}_{\text{KRR}}(\mathbf{x}_i, 1/t) - \mathbf{y} \right)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \mathbf{e}_i^\top (\mathbf{I} + t\mathbf{K})^{-1} \mathbf{y} \right)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \mathbf{e}_i^\top \mathbf{y} \right)^2 \right)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \mathbf{e}_i^\top (\mathbf{I} + t\mathbf{K})^{-1} \mathbf{y} \mathbf{y}^\top (\mathbf{I} + t\mathbf{K})^{-1} \mathbf{e}_i \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \mathbf{e}_i^\top \mathbf{y} \mathbf{y}^\top \mathbf{e}_i \right)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbf{e}_i^\top (\mathbf{I} + t\mathbf{K})^{-1} (\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) (\mathbf{I} + t\mathbf{K})^{-1} \mathbf{e}_i}{\mathbf{e}_i^\top (\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) \mathbf{e}_i} \right) \\
&= \max_{\|\sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} \mathbf{e}_i\|_2 \neq 0} \left( \frac{\left( \left\| (\mathbf{I} + t\mathbf{K})^{-1} \sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} \mathbf{e}_i \right\|_2 \right)^2}{\left\| \sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} \mathbf{e}_i \right\|_2^2} \right) = \frac{1}{(1 + ts_{\min})^2},
\end{aligned}$$

where  $e^{-2ts_{\min}} \leq \frac{1}{(1 + ts_{\min})^2}$ , since  $e^x \geq 1 + x$ .

For part (d),

$$\max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( \left( \hat{f}_{\text{KRR}}(\mathbf{x}_i, 1/t) \right)^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2)} \right)$$

$$\begin{aligned}
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0}(y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( (e_i^\top (\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1}) \mathbf{y})^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( (e_i^\top \mathbf{y})^2 \right)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0}(y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( e_i^\top ((\mathbf{I} - \mathbf{I} + t\mathbf{K})^{-1}) \mathbf{y} \mathbf{y}^\top (\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1}) e_i \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( e_i^\top \mathbf{y} \mathbf{y}^\top e_i \right)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0}(y_i^2) \neq 0} \left( \frac{e_i^\top (\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1}) (\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) (\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1}) e_i}{e_i^\top (\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) e_i} \right) \\
&= \max_{\|\sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} e_i\|_2 \neq 0} \left( \frac{\left\| (\mathbf{I} - (\mathbf{I} + t\mathbf{K})^{-1}) \sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} e_i \right\|_2^2}{\left\| \sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} e_i \right\|_2^2} \right)^2 \\
&= \left( 1 - \frac{1}{(1 + ts_{\max})} \right)^2,
\end{aligned}$$

and

$$\begin{aligned}
&\max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0}(y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( (\hat{f}_{\text{KGF}}(\mathbf{x}_i, t) - \mathbf{y})^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} (y_i^2)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0}(y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( (e_i^\top (\mathbf{I} - \exp(-t\mathbf{K})) \mathbf{y})^2 \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( (e_i^\top \mathbf{y})^2 \right)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0}(y_i^2) \neq 0} \left( \frac{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( e_i^\top (\mathbf{I} - \exp(-t\mathbf{K})) \mathbf{y} \mathbf{y}^\top (\mathbf{I} - \exp(-t\mathbf{K})) e_i \right)}{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0} \left( e_i^\top \mathbf{y} \mathbf{y}^\top e_i \right)} \right) \\
&= \max_{\mathbb{E}_{\boldsymbol{\varepsilon}, \boldsymbol{\alpha}_0}(y_i^2) \neq 0} \left( \frac{e_i^\top (\mathbf{I} - \exp(-t\mathbf{K})) (\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) (\mathbf{I} - \exp(-t\mathbf{K})) e_i}{e_i^\top (\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}) e_i} \right) \\
&= \max_{\|\sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} e_i\|_2 \neq 0} \left( \frac{\left\| (\mathbf{I} - \exp(-t\mathbf{K})) \sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} e_i \right\|_2^2}{\left\| \sqrt{\boldsymbol{\Sigma}_\alpha \mathbf{K}^2 + \sigma_\varepsilon^2 \mathbf{I}} e_i \right\|_2^2} \right)^2 \\
&= (1 - e^{-ts_{\max}})^2,
\end{aligned}$$

where  $\left(1 - \frac{1}{1 + ts_{\max}}\right)^2 \leq (1 - e^{-ts_{\max}})^2$ , since  $e^x \geq 1 + x$ .  $\square$

*Proof of Proposition 4.3.* Unless otherwise stated, all expectations and covariances are with respect to the random variable  $\boldsymbol{\varepsilon}$ . We let  $\mathbf{K} = \mathbf{U} \mathbf{S} \mathbf{U}^\top$  denote the eigenvalue decomposition of  $\mathbf{K}$ .

Before starting the calculations, we show some intermediary results that will be needed:

$$\mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t)) - \boldsymbol{\alpha}_0 = (\mathbf{I} - \exp(-t\mathbf{K})) \mathbf{K}^{-1} \mathbb{E}(\mathbf{y}) - \boldsymbol{\alpha}_0$$



$$\begin{aligned}
&= (\mathbf{I} - \exp(-t\mathbf{K})) \boldsymbol{\alpha}_0 - \boldsymbol{\alpha}_0 = -\exp(-t\mathbf{K}) \boldsymbol{\alpha}_0. \\
\text{Cov}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t)) &= \mathbb{E} \left( \hat{\boldsymbol{\alpha}}_{\text{KGF}}(t) (\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t))^\top \right) - \mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t)) \mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t))^\top \\
&= (\mathbf{I} - \exp(-t\mathbf{K})) \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top (\mathbf{I} - \exp(-t\mathbf{K})) + \sigma_\varepsilon^2 \mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2 \\
&\quad - (\mathbf{I} - \exp(-t\mathbf{K})) \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top (\mathbf{I} - \exp(-t\mathbf{K})) \\
&= \sigma_\varepsilon^2 \mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2. \\
\mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t)) - \boldsymbol{\alpha}_0 &= (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbb{E}(\mathbf{y}) - \boldsymbol{\alpha}_0 = (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbf{K} \boldsymbol{\alpha}_0 - \boldsymbol{\alpha}_0 \\
&= \left( (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbf{K} - \mathbf{I} \right) \boldsymbol{\alpha}_0 = (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} (\mathbf{K} - \mathbf{K} - 1/t \cdot \mathbf{I}) \boldsymbol{\alpha}_0 \\
&= -1/t \cdot (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \boldsymbol{\alpha}_0. \\
\text{Cov}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t)) &= \mathbb{E} \left( \hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t) (\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t))^\top \right) \\
&\quad - \mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t)) \mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t))^\top \\
&= (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} (\mathbf{K} \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top \mathbf{K} + \sigma_\varepsilon^2 \mathbf{I}) (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \\
&\quad - (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbf{K} \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top \mathbf{K} (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \\
&= \sigma_\varepsilon^2 (\mathbf{K} + 1/t \cdot \mathbf{I})^{-2}. \\
\mathbb{E} \left( \hat{f}_{\text{KGF}}(\mathbf{x}^*, t) \right) - f_0(\mathbf{x}^*) &= \mathbf{k}(\mathbf{x}^*)^\top (\mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t)) - \boldsymbol{\alpha}_0) = \\
&\quad - \mathbf{k}(\mathbf{x}^*)^\top \exp(-t\mathbf{K}) \boldsymbol{\alpha}_0. \\
\mathbb{E} \left( \hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t) \right) - f_0(\mathbf{x}^*) &= \mathbf{k}(\mathbf{x}^*)^\top (\mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t)) - \boldsymbol{\alpha}_0) = \\
&\quad - 1/t \cdot \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \boldsymbol{\alpha}_0. \\
\text{Var}(\hat{f}_{\text{KGF}}(\mathbf{x}^*, t)) &= \mathbb{E} \left( \hat{f}_{\text{KGF}}(\mathbf{x}^*, t)^2 \right) - \mathbb{E} \left( \hat{f}_{\text{KGF}}(\mathbf{x}^*, t) \right)^2 \\
&= \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{I} - \exp(-t\mathbf{K})) \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top (\mathbf{I} - \exp(-t\mathbf{K})) \mathbf{k}(\mathbf{x}^*) \\
&\quad + \sigma_\varepsilon^2 \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2 \mathbf{k}(\mathbf{x}^*) \\
&\quad - \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{I} - \exp(-t\mathbf{K})) \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top (\mathbf{I} - \exp(-t\mathbf{K})) \mathbf{k}(\mathbf{x}^*) \\
&= \sigma_\varepsilon^2 \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2 \mathbf{k}(\mathbf{x}^*). \\
\text{Var}(\hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t)) &= \mathbb{E} \left( \hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t)^2 \right) - \mathbb{E} \left( \hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t) \right)^2 \\
&= \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} (\mathbf{K} \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top \mathbf{K} + \sigma_\varepsilon^2 \mathbf{I}) (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbf{k}(\mathbf{x}^*) \\
&\quad - \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbf{K} \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top \mathbf{K} (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbf{k}(\mathbf{x}^*) \\
&= \sigma_\varepsilon^2 \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-2} \mathbf{k}(\mathbf{x}^*).
\end{aligned}$$

According to the bias covariance decomposition, we can write the risk of an estimator  $\hat{\boldsymbol{\theta}}$ , estimating  $\boldsymbol{\theta}_0$ , as

$$\text{Risk}(\hat{\boldsymbol{\theta}}; \boldsymbol{\theta}_0) = \mathbb{E} \left( \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2^2 \right) = \|\mathbb{E}(\hat{\boldsymbol{\theta}}) - \boldsymbol{\theta}_0\|_2^2 + \text{Tr}(\text{Cov}(\hat{\boldsymbol{\theta}})).$$

Thus,

$$\begin{aligned}
\text{Risk}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t); \boldsymbol{\alpha}_0) &= \|\mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t)) - \boldsymbol{\alpha}_0\|_2^2 + \text{Tr}(\text{Cov}(\hat{\boldsymbol{\alpha}}_{\text{KGF}})) \\
&= \|\exp(-t\mathbf{K})\boldsymbol{\alpha}_0\|_2^2 + \text{Tr}\left(\sigma_\varepsilon^2 \mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2\right) \\
&= \boldsymbol{\alpha}_0^\top \mathbf{U} e^{-t\mathbf{S}} \mathbf{U}^\top \mathbf{U} e^{-t\mathbf{S}} \mathbf{U}^\top \boldsymbol{\alpha}_0 + \sigma_\varepsilon^2 \text{Tr}\left(\mathbf{U} \left(\frac{1 - e^{-t\mathbf{S}}}{\mathbf{S}}\right)^2 \mathbf{U}^\top\right) \\
&= \boldsymbol{\alpha}_0^\top \mathbf{U} e^{-2t\mathbf{S}} \mathbf{U}^\top \boldsymbol{\alpha}_0 + \sigma_\varepsilon^2 \text{Tr}\left(\left(\frac{1 - e^{-t\mathbf{S}}}{\mathbf{S}}\right)^2\right) \\
&= \sum_{i=1}^n ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 e^{-2ts_i} + \sigma_\varepsilon^2 \sum_{i=1}^n \left(\frac{1 - e^{-ts_i}}{s_i}\right)^2,
\end{aligned}$$

where in the third equality, we have used the cyclic property of the trace.

Equivalently, for the risk of  $\hat{\boldsymbol{\alpha}}_{\text{KRR}}$ , for  $\lambda = 1/t$ ,

$$\begin{aligned}
\text{Risk}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t); \boldsymbol{\alpha}_0) &= \|\mathbb{E}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t)) - \boldsymbol{\alpha}_0\|_2^2 + \text{Tr}(\text{Cov}(\hat{\boldsymbol{\alpha}}_{\text{KRR}})) \\
&= \|\mathbf{I} - 1/t \cdot (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \boldsymbol{\alpha}_0\|_2^2 + \text{Tr}\left(\sigma_\varepsilon^2 (\mathbf{K} + 1/t \cdot \mathbf{I})^{-2}\right) \\
&= \boldsymbol{\alpha}_0^\top \mathbf{U} \frac{1/t}{\mathbf{S} + 1/t} \mathbf{U}^\top \mathbf{U} \frac{1/t}{\mathbf{S} + 1/t} \mathbf{U}^\top \boldsymbol{\alpha}_0 + \sigma_\varepsilon^2 \text{Tr}\left(\mathbf{U} \left(\frac{1}{\mathbf{S} + 1/t}\right)^2 \mathbf{U}^\top\right) \\
&= \boldsymbol{\alpha}_0^\top \mathbf{U} \frac{1}{(t \cdot \mathbf{S} + 1)^2} \mathbf{U}^\top \boldsymbol{\alpha}_0 + \sigma_\varepsilon^2 \text{Tr}\left(\frac{t^2}{(t \cdot \mathbf{S} + 1)^2}\right) \\
&= \sum_{i=1}^n ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 \frac{1}{(ts_i + 1)^2} + \sigma_\varepsilon^2 \sum_{i=1}^n \frac{t^2}{(ts_i + 1)^2}.
\end{aligned}$$

For  $x \geq 0$ , the following two inequalities are shown by [Ali et al. \(2019\)](#):

$$e^{-x} \leq 1/(x+1) \text{ and } 1 - e^{-x} \leq 1.2985 \cdot x/(x+1).$$

Using these, we obtain

$$\begin{aligned}
\text{Risk}(\hat{\boldsymbol{\alpha}}_{\text{KGF}}(t); \boldsymbol{\alpha}_0) &= \sum_{i=1}^n \left( ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 e^{-2ts_i} + \sigma_\varepsilon^2 \left(\frac{1 - e^{-ts_i}}{s_i}\right)^2 \right) \\
&\leq \sum_{i=1}^n \left( ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 \frac{1}{(ts_i + 1)^2} + \sigma_\varepsilon^2 \left(1.2985 \frac{ts_i}{ts_i + 1}\right)^2 \frac{1}{s_i^2} \right) \\
&= \sum_{i=1}^n \left( ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 \frac{1}{(ts_i + 1)^2} + \sigma_\varepsilon^2 1.2985^2 \frac{t^2}{(ts_i + 1)^2} \right) \\
&\leq 1.6862 \sum_{i=1}^n \left( ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 \frac{1}{(ts_i + 1)^2} + \sigma_\varepsilon^2 \frac{t^2}{(ts_i + 1)^2} \right) \\
&= 1.6862 \cdot \text{Risk}(\hat{\boldsymbol{\alpha}}_{\text{KRR}}(1/t); \boldsymbol{\alpha}_0),
\end{aligned}$$

which proves part (a).

The calculations for part (b) are basically identical to those for part (a). We obtain

$$\begin{aligned}\text{Risk}\left(\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t); \mathbf{f}_0(\mathbf{X})\right) &= \sum_{i=1}^n ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 s_i^2 e^{-2ts_i} + \sigma_\varepsilon^2 \sum_{i=1}^n (1 - e^{-ts_i})^2, \\ \text{Risk}\left(\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, 1/t); \mathbf{f}_0(\mathbf{X})\right) &= \sum_{i=1}^n ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 \frac{s_i^2}{(ts_i + 1)^2} + \sigma_\varepsilon^2 \sum_{i=1}^n \frac{t^2 s_i^2}{(ts_i + 1)^2},\end{aligned}$$

and thus

$$\begin{aligned}\text{Risk}\left(\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{X}, t); \mathbf{f}_0(\mathbf{X})\right) &= \sum_{i=1}^n \left( ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 s_i^2 e^{-2ts_i} + \sigma_\varepsilon^2 (1 - e^{-ts_i})^2 \right) \\ &\leq \sum_{i=1}^n \left( ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 \frac{s_i^2}{(ts_i + 1)^2} + \sigma_\varepsilon^2 \left( 1.2985 \frac{ts_i}{ts_i + 1} \right)^2 \right) \\ &\leq 1.2985^2 \sum_{i=1}^n \left( ((\mathbf{U}_{:,i})^\top \boldsymbol{\alpha}_0)^2 \frac{s_i^2}{(ts_i + 1)^2} + \sigma_\varepsilon^2 \frac{t^2 s_i^2}{(ts_i + 1)^2} \right) \\ &= 1.6862 \cdot \text{Risk}\left(\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{X}, 1/t); \mathbf{f}_0(\mathbf{X})\right),\end{aligned}$$

which proves part (b).

For the out-of-sample prediction risk, since  $\mathbf{k}(\mathbf{x}^*)$  and  $\mathbf{K}$  do not commute, calculations become slightly different. We now obtain

$$\begin{aligned}\text{Risk}\left(\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{x}^*, t); f_0(\mathbf{x}^*, t)\right) &= \left( \mathbb{E}(\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{x}^*, t)) - f_0(\mathbf{x}^*, t) \right)^2 + \text{Var}(\hat{\mathbf{f}}_{\text{KGF}}(\mathbf{x}^*, t)) \\ &= (-\mathbf{k}(\mathbf{x}^*)^\top \exp(-t\mathbf{K})\boldsymbol{\alpha}_0)^2 + \sigma_\varepsilon^2 \mathbf{k}(\mathbf{x}^*)^\top \mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2 \mathbf{k}(\mathbf{x}^*) \\ &= \boldsymbol{\alpha}_0^\top \exp(-t\mathbf{K}) \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \exp(-t\mathbf{K}) \boldsymbol{\alpha}_0 \\ &\quad + \sigma_\varepsilon^2 \text{Tr}\left(\mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2 \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top\right) \\ &= \text{Tr}\left(\exp(-t\mathbf{K}) \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top \exp(-t\mathbf{K}) \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top\right) \\ &\quad + \sigma_\varepsilon^2 \text{Tr}\left(\mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2 \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top\right),\end{aligned}$$

where we have used the fact that the trace of a scalar is the scalar itself, and the cyclic property of the trace.

Analogously, for the risk of  $\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{x}^*)$ ,

$$\begin{aligned}\text{Risk}\left(\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{x}^*, 1/t); f_0(\mathbf{x}^*)\right) &= \left( \mathbb{E}(\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{x}^*, 1/t)) - f_0(\mathbf{x}^*) \right)^2 + \text{Var}(\hat{\mathbf{f}}_{\text{KRR}}(\mathbf{x}^*, 1/t)) \\ &= \left( -1/t \cdot \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \boldsymbol{\alpha}_0 \right)^2 + \sigma_\varepsilon^2 \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-2} \mathbf{k}(\mathbf{x}^*)\end{aligned}$$

$$\begin{aligned}
&= 1/t^2 \cdot \boldsymbol{\alpha}_0^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \boldsymbol{\alpha}_0 \\
&\quad + \sigma_\varepsilon^2 \text{Tr} \left( (\mathbf{K} + 1/t \cdot \mathbf{I})^{-2} \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right) \\
&= \text{Tr} \left( 1/t^2 \cdot (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^\top (\mathbf{K} + 1/t \cdot \mathbf{I})^{-1} \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right) \\
&\quad + \sigma_\varepsilon^2 \text{Tr} \left( (\mathbf{K} + 1/t \cdot \mathbf{I})^{-2} \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right).
\end{aligned}$$

Taking expectation over  $\boldsymbol{\alpha}_0$ , for  $\boldsymbol{\alpha}_0 \sim (\mathbf{0}, \boldsymbol{\Sigma}_\alpha)$ , we obtain

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\alpha}_0} \left( \text{Risk} \left( \hat{f}_{\text{KGF}}(\mathbf{x}^*, t); f_0(\mathbf{x}^*) \right) \right) \\
&\quad = \text{Tr} \left( \left( \boldsymbol{\Sigma}_\alpha \exp(-t\mathbf{K})^2 + \sigma_\varepsilon^2 \mathbf{K}^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2 \right) \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right), \\
&\mathbb{E}_{\boldsymbol{\alpha}_0} \left( \text{Risk} \left( \hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t); f_0(\mathbf{x}^*) \right) \right) \\
&\quad = \text{Tr} \left( \left( \boldsymbol{\Sigma}_\alpha 1/t^2 \cdot (\mathbf{K} + 1/t \cdot \mathbf{I})^{-2} + \sigma_\varepsilon^2 (\mathbf{K} + 1/t \cdot \mathbf{I})^{-2} \right) \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right).
\end{aligned}$$

Finally comparing the risks, we obtain

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{\alpha}_0} \left( \text{Risk}(\hat{f}_{\text{KGF}}(\mathbf{x}^*, t); f_0(\mathbf{x}^*)) \right) \\
&\quad = \text{Tr} \left( \left( \boldsymbol{\Sigma}_\alpha \exp(-t\mathbf{K})^2 + \sigma_\varepsilon^2 t^2 (t\mathbf{K})^{-2} (\mathbf{I} - \exp(-t\mathbf{K}))^2 \right) \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right) \\
&\quad \leq \text{Tr} \left( \left( \boldsymbol{\Sigma}_\alpha (\mathbf{I} + t\mathbf{K})^{-2} + \sigma_\varepsilon^2 t^2 \left( 1.2985 (\mathbf{I} + t\mathbf{K})^{-1} \right)^2 \right) \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right) \\
&\quad = \text{Tr} \left( \left( \boldsymbol{\Sigma}_\alpha \cdot 1/t^2 \cdot (1/t \cdot \mathbf{I} + \mathbf{K})^{-2} \right. \right. \\
&\quad \quad \left. \left. + \sigma_\varepsilon^2 \cdot 1.2985^2 (1/t \cdot \mathbf{I} + \mathbf{K})^{-2} \right) \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right) \\
&\quad \leq 1.6862 \cdot \text{Tr} \left( \left( \boldsymbol{\Sigma}_\alpha \cdot 1/t^2 \cdot (1/t \cdot \mathbf{I} + \mathbf{K})^{-2} \right. \right. \\
&\quad \quad \left. \left. + \sigma_\varepsilon^2 (1/t \cdot \mathbf{I} + \mathbf{K})^{-2} \right) \mathbf{k}(\mathbf{x}^*) \mathbf{k}(\mathbf{x}^*)^\top \right) \\
&\quad = 1.6862 \cdot \mathbb{E}_{\boldsymbol{\alpha}_0} \left( \text{Risk}(\hat{f}_{\text{KRR}}(\mathbf{x}^*, 1/t); f_0(\mathbf{x}^*)) \right),
\end{aligned}$$

which proves part (c).  $\square$

*Proof of Proposition 5.1.* The proofs of the two parts are very similar, differing only in the details.

For part (a), when  $\mathbf{X}^\top \mathbf{X}$  is a diagonal matrix with elements  $\{s_{ii}\}_{i=1}^p$ ,

$$\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 = \mathbf{y}^\top \mathbf{y} - 2\mathbf{y}^\top \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\beta}^\top \mathbf{X}^\top \mathbf{X}\boldsymbol{\beta} = \sum_{i=1}^p (y_i^2 - 2(\mathbf{X}^\top \mathbf{y})_i \beta_i + s_{ii} \beta_i^2).$$

The constraint  $\|\boldsymbol{\beta}\|_\infty = \max_i |\beta_i| \leq c$  is equivalent to  $|\beta_i| \leq c$  for  $i = 1, 2, \dots, p$ . Thus,

$$\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \text{ s.t. } \|\boldsymbol{\beta}\|_\infty \leq c \iff \sum_{i=1}^p (y_i^2 - 2(\mathbf{X}^\top \mathbf{y})_i \beta_i + s_{ii} \beta_i^2) \text{ s.t. } |\beta_i| \leq c \forall i,$$

which decomposes element-wise. In the absence of the constraint,

$$\hat{\beta}_i = ((\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y})_i = (\mathbf{X}^\top \mathbf{y})_i / s_{ii}.$$

Assume  $(\mathbf{X}^\top \mathbf{y})_i / s_{ii} \geq 0$ . Then, the optimal value for  $\beta_i$  is  $(\mathbf{X}^\top \mathbf{y})_i / s_{ii}$ , unless  $(\mathbf{X}^\top \mathbf{y})_i / s_{ii} > c$ , then, due to convexity, the optimal value is  $c$ , i.e.  $\hat{\beta}_i = \min((\mathbf{X}^\top \mathbf{y})_i / s_{ii}, c)$ . Accounting also for the case of  $(\mathbf{X}^\top \mathbf{y})_i / s_{ii} < 0$ , we obtain

$$\hat{\beta}_i(c) = \text{sign}((\mathbf{X}^\top \mathbf{y})_i / s_{ii}) \cdot \min(|(\mathbf{X}^\top \mathbf{y})_i / s_{ii}|, c).$$

The sign gradient flow solution is calculated as follows:

$$\begin{aligned} \text{sign} \left( -\frac{\partial}{\partial \hat{\beta}(t)} \left( \|\mathbf{y} - \mathbf{X} \hat{\beta}(t)\|_2^2 \right) \right) &= \text{sign}(\mathbf{X}^\top \mathbf{y} - \mathbf{X}^\top \mathbf{X} \hat{\beta}(t)) \\ &= \text{sign} \left( \begin{bmatrix} (\mathbf{X}^\top \mathbf{y})_i - s_{ii} \hat{\beta}_i(t) \end{bmatrix} \right) = \begin{bmatrix} \text{sign}((\mathbf{X}^\top \mathbf{y})_i - s_{ii} \hat{\beta}_i(t)) \end{bmatrix}. \end{aligned}$$

Since element  $i$  in the vector only depends on  $\hat{\beta}$  through  $\hat{\beta}_i$ ,

$$\begin{aligned} \frac{\partial \hat{\beta}(t)}{\partial t} &= \text{sign} \left( -\frac{\partial}{\partial \hat{\beta}(t)} \left( \|\mathbf{y} - \mathbf{X} \hat{\beta}(t)\|_2^2 \right) \right) \\ \iff \frac{\partial \hat{\beta}_i(t)}{\partial t} &= \text{sign}((\mathbf{X}^\top \mathbf{y})_i - s_{ii} \hat{\beta}_i(t)) \\ &= \begin{cases} 1 & \text{if } (\mathbf{X}^\top \mathbf{y})_i > s_{ii} \hat{\beta}_i(t) \\ -1 & \text{if } (\mathbf{X}^\top \mathbf{y})_i < s_{ii} \hat{\beta}_i(t), \quad i = 1, 2, \dots, p. \\ 0 & \text{if } (\mathbf{X}^\top \mathbf{y})_i = s_{ii} \hat{\beta}_i(t) \end{cases} \end{aligned}$$

Thus, with  $\hat{\beta}_i(0) = 0$ ,  $\hat{\beta}_i(t) = \pm t$ , depending on the sign of  $(\mathbf{X}^\top \mathbf{y})_i / s_{ii}$ . Once  $\hat{\beta}_i = (\mathbf{X}^\top \mathbf{y})_i / s_{ii}$  (and  $\text{sign}(s_{ii} \hat{\beta}_i - (\mathbf{X}^\top \mathbf{y})_i) = 0$ ), then  $\hat{\beta}_i = (\mathbf{X}^\top \mathbf{y})_i / s_{ii}$ . That is,

$$\hat{\beta}_i(t) = \text{sign}((\mathbf{X}^\top \mathbf{y})_i / s_{ii}) \cdot \min(t, |(\mathbf{X}^\top \mathbf{y})_i / s_{ii}|).$$

For part (b), when  $\mathbf{K}$  is a diagonal matrix with elements  $\{k_{ii}\}_{i=1}^n$ ,

$$\|\mathbf{y} - \mathbf{K} \boldsymbol{\alpha}\|_{\mathbf{K}^{-1}}^2 = \left\| (\sqrt{\mathbf{K}})^{-1} \mathbf{y} - \sqrt{\mathbf{K}} \boldsymbol{\alpha} \right\|_2^2 = \sum_{i=1}^n \left( \frac{y_i}{\sqrt{k_{ii}}} - \sqrt{k_{ii}} \alpha_i \right)^2.$$

The constraint  $\|\boldsymbol{\alpha}\|_\infty = \max_i |\alpha_i| \leq c$  is equivalent to  $|\alpha_i| \leq c$ , for  $i = 1, 2, \dots, n$ . Thus,

$$\|\mathbf{y} - \mathbf{K} \boldsymbol{\alpha}\|_{\mathbf{K}^{-1}}^2 \text{ s.t. } \|\boldsymbol{\alpha}\|_\infty \leq c \iff \sum_{i=1}^n \left( \frac{y_i}{\sqrt{k_{ii}}} - \sqrt{k_{ii}} \alpha_i \right)^2 \text{ s.t. } |\alpha_i| \leq c \quad \forall i,$$

which decomposes element-wise. Assume  $y_i/k_{ii} \geq 0$ . Then, the optimal value for  $\alpha_i$  is  $y_i/k_{ii}$ , unless  $y_i/k_{ii} > c$ , then, due to convexity, the optimal value is  $c$ , i.e.  $\hat{\alpha}_i = \min(y_i/k_{ii}, c)$ . Accounting also for the case of  $y_i/k_{ii} < 0$ , we obtain

$$\hat{\alpha}_i(c) = \text{sign}(y_i/k_{ii}) \cdot \min(|y_i/k_{ii}|, c).$$

The sign gradient flow solution is calculated as follows:

$$\begin{aligned} \text{sign} \left( -\frac{\partial}{\partial \hat{\alpha}(t)} \left( \|\mathbf{y} - \mathbf{K} \hat{\alpha}(t)\|_{\mathbf{K}^{-1}}^2 \right) \right) &= \text{sign}(\mathbf{y} - \mathbf{K} \hat{\alpha}(t)) \\ &= \text{sign} \left( \begin{bmatrix} y_1 - k_{11} \hat{\alpha}_1(t) \\ \vdots \\ y_n - k_{nn} \hat{\alpha}_n(t) \end{bmatrix} \right) = \begin{bmatrix} \text{sign}(y_1 - k_{11} \hat{\alpha}_1(t)) \\ \vdots \\ \text{sign}(y_n - k_{nn} \hat{\alpha}_n(t)) \end{bmatrix}. \end{aligned}$$

Since element  $i$  in the vector only depends on  $\hat{\alpha}$  through  $\hat{\alpha}_i$ ,

$$\begin{aligned} \frac{\partial \hat{\alpha}(t)}{\partial t} &= \text{sign} \left( -\frac{\partial}{\partial \hat{\alpha}(t)} \left( \|\mathbf{y} - \mathbf{K} \hat{\alpha}(t)\|_{\mathbf{K}^{-1}}^2 \right) \right) \\ \iff \frac{\partial \hat{\alpha}_i(t)}{\partial t} &= \text{sign}(y_i - k_{ii} \hat{\alpha}_i(t)) = \begin{cases} 1 & \text{if } y_i > k_{ii} \hat{\alpha}_i(t) \\ -1 & \text{if } y_i < k_{ii} \hat{\alpha}_i(t), \quad i = 1, 2, \dots, n. \\ 0 & \text{if } y_i = k_{ii} \hat{\alpha}_i(t) \end{cases} \end{aligned}$$

Thus, with  $\hat{\alpha}_i(0) = 0$ ,  $\hat{\alpha}_i(t) = \pm t$ , depending on the sign of  $y_i/k_{ii}$ . Once  $\hat{\alpha}_i = y_i/k_{ii}$  (and  $\text{sign}(y_i - k_{ii} \hat{\alpha}_i) = 0$ ), then  $\hat{\alpha}_i = y_i/k_{ii}$ . That is,

$$\hat{\alpha}_i(t) = \text{sign}(y_i/k_{ii}) \cdot \min(t, |y_i/k_{ii}|). \quad \square$$

*Proof of Proposition 5.2.* We first calculate the gradients of  $\|\mathbf{y} - \Phi \beta\|_2^2$ ,  $\|\mathbf{y} - \Phi \beta\|_1$  and  $\|\mathbf{y} - \Phi \beta\|_\infty$ , respectively:

$$\frac{\partial}{\partial \beta} \left( \|\mathbf{y} - \Phi \beta\|_2^2 \right) = \frac{\partial}{\partial \beta} \left( (\mathbf{y} - \Phi \beta)^\top (\mathbf{y} - \Phi \beta) \right) = -\Phi^\top (\mathbf{y} - \Phi \beta).$$

Since  $\|\mathbf{v}\|_1 = \text{sign}(\mathbf{v})^\top \mathbf{v}$ ,

$$\frac{\partial}{\partial \beta} (\|\mathbf{y} - \Phi \beta\|_1) = \frac{\partial}{\partial \beta} \left( \text{sign}(\mathbf{y} - \Phi \beta)^\top (\mathbf{y} - \Phi \beta) \right) = -\Phi^\top \text{sign}(\mathbf{y} - \Phi \beta).$$

Let  $\mathbb{I}_\infty(\mathbf{v})$  be a vector denoting the sign of the largest (absolute) value in  $\mathbf{v}$ , such that

$$\mathbb{I}_\infty(\mathbf{v})_d = \begin{cases} \text{sign}(v_d) & \text{if } d = \arg \max_{d'} |v_{d'}| \\ 0 & \text{else.} \end{cases}$$

Then  $\|\mathbf{v}\|_\infty = \mathbb{I}_\infty(\mathbf{v})^\top \mathbf{v}$ , and

$$\frac{\partial}{\partial \beta} (\|\mathbf{y} - \Phi \beta\|_\infty) = \frac{\partial}{\partial \beta} \left( \mathbb{I}_\infty(\mathbf{y} - \Phi \beta)^\top (\mathbf{y} - \Phi \beta) \right) = -\Phi^\top \mathbb{I}_\infty(\mathbf{y} - \Phi \beta).$$

The three update rules for gradient descent in  $\beta$  are thus

$$\begin{aligned}\hat{\beta}_{k+1} &= \hat{\beta}_k + \eta \cdot \Phi^\top (\mathbf{y} - \Phi \hat{\beta}_k), & \text{for } \|\mathbf{y} - \Phi \hat{\beta}_k\|_2^2 \\ \hat{\beta}_{k+1} &= \hat{\beta}_k + \eta \cdot \Phi^\top \text{sign}(\mathbf{y} - \Phi \hat{\beta}_k), & \text{for } \|\mathbf{y} - \Phi \hat{\beta}_k\|_1 \\ \hat{\beta}_{k+1} &= \hat{\beta}_k + \eta \cdot \Phi^\top \mathbb{I}_\infty(\mathbf{y} - \Phi \hat{\beta}_k), & \text{for } \|\mathbf{y} - \Phi \hat{\beta}_k\|_\infty.\end{aligned}\quad (23)$$

Since the gradient of  $\|\mathbf{y} - \mathbf{K}\alpha\|_{\mathbf{K}^{-1}}^2$  is  $-(\mathbf{y} - \mathbf{K}\alpha)$ , the update rules for gradient descent, sign gradient descent and coordinate descent are, respectively

$$\begin{aligned}\hat{\alpha}_{k+1} &= \hat{\alpha}_k + \eta \cdot (\mathbf{y} - \mathbf{K}\hat{\alpha}_k) \\ \hat{\alpha}_{k+1} &= \hat{\alpha}_k + \eta \cdot \text{sign}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k) \\ \hat{\alpha}_{k+1} &= \hat{\alpha}_k + \eta \cdot \mathbb{I}_\infty(\mathbf{y} - \mathbf{K}\hat{\alpha}_k).\end{aligned}\quad (24)$$

By multiplying each term in Equation 24 by  $\Phi^\top$ , and using  $\beta = \Phi^\top \alpha$  and  $\mathbf{K} = \Phi \Phi^\top$ , and thus  $\mathbf{K}\alpha = \Phi \beta$ , we obtain Equation 23.  $\square$

*Proof of Proposition 5.3.* For KGD, let  $L_2(\hat{\alpha}) := \|\mathbf{y} - \mathbf{K}\hat{\alpha}\|_2^2$ . Then  $L_2'(\hat{\alpha}) = -\mathbf{K}(\mathbf{y} - \mathbf{K}\hat{\alpha})$ ,  $L_2''(\hat{\alpha}) = \mathbf{K}^2$  and higher order derivatives are zero. With  $\Delta\hat{\alpha} = \eta \cdot (\mathbf{y} - \mathbf{K}\hat{\alpha})$ , according to Taylor's theorem

$$\begin{aligned}L_2(\hat{\alpha}_{k+1}) - L_2(\hat{\alpha}_k) &= L_2(\hat{\alpha}_k + \Delta\hat{\alpha}_k) - L_2(\hat{\alpha}_k) \\ &= L_2(\hat{\alpha}_k) + L_2'(\hat{\alpha}_k)^\top \Delta\hat{\alpha}_k + \frac{1}{2} \Delta\hat{\alpha}_k^\top L_2''(\hat{\alpha}_k) \Delta\hat{\alpha}_k - L_2(\hat{\alpha}_k) \\ &= -\eta \cdot (\mathbf{y} - \mathbf{K}\hat{\alpha}_k)^\top \mathbf{K}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k) + \frac{\eta^2}{2} (\mathbf{y} - \mathbf{K}\hat{\alpha}_k)^\top \mathbf{K}^2 (\mathbf{y} - \mathbf{K}\hat{\alpha}_k) \\ &\leq -\eta \cdot (\mathbf{y} - \mathbf{K}\hat{\alpha}_k)^\top \mathbf{K}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k) + \frac{\eta^2}{2} s_{\max}(\mathbf{K}) (\mathbf{y} - \mathbf{K}\hat{\alpha}_k)^\top \mathbf{K}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k).\end{aligned}$$

Factoring out  $\eta \cdot (\mathbf{y} - \mathbf{K}\hat{\alpha}_k)^\top \mathbf{K}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k)$ , and using that  $(\mathbf{y} - \mathbf{K}\hat{\alpha}_k)^\top \mathbf{K}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k) \geq 0$ , since  $\mathbf{K}$  is positive semi-definite, we see that  $L_2(\hat{\alpha}_{k+1}) - L_2(\hat{\alpha}_k) \leq 0$  if  $\eta \leq \frac{2}{s_{\max}(\mathbf{K})}$ .

For KSGD,  $L_1(\hat{\alpha}) := \|\mathbf{y} - \mathbf{K}\hat{\alpha}\|_1 = \text{sign}(\mathbf{y} - \mathbf{K}\hat{\alpha})^\top (\mathbf{y} - \mathbf{K}\hat{\alpha})$ . Then  $L_1'(\hat{\alpha}) = -\mathbf{K} \cdot \text{sign}(\mathbf{y} - \mathbf{K}\hat{\alpha})$  and higher order derivatives are zero. With  $\Delta\hat{\alpha} = \eta \cdot \text{sign}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k)$ , according to Taylor's theorem, for  $\eta \leq \min_{1 \leq i \leq n} (\|\mathbf{y} - \mathbf{K}\hat{\alpha}_k\|_i)$

$$\begin{aligned}L_1(\hat{\alpha}_{k+1}) - L_1(\hat{\alpha}_k) &= L_1(\hat{\alpha}_k + \Delta\hat{\alpha}_k) - L_1(\hat{\alpha}_k) = L_1(\hat{\alpha}_k) + L_1'(\hat{\alpha}_k)^\top \Delta\hat{\alpha}_k \\ &= -\eta \cdot \text{sign}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k)^\top \mathbf{K} \cdot \text{sign}(\mathbf{y} - \mathbf{K}\hat{\alpha}_k) \leq 0.\end{aligned}$$

For KCD, let  $\mathbb{I}_\infty(\mathbf{v})$  be a vector denoting the sign of the largest (absolute) value in  $\mathbf{v}$ , such that

$$\mathbb{I}_\infty(\mathbf{v})_d = \begin{cases} \text{sign}(v_d) & \text{if } d = \arg \max_{d'} |v_{d'}| \\ 0 & \text{else,} \end{cases}$$

so that  $\|\mathbf{v}\|_\infty = \mathbb{I}_\infty(\mathbf{v})^\top \mathbf{v}$ . Furthermore, let  $L_\infty(\hat{\boldsymbol{\alpha}}) := \|\mathbf{y} - \mathbf{K}\hat{\boldsymbol{\alpha}}\|_\infty = \mathbb{I}_\infty(\mathbf{y} - \mathbf{K}\hat{\boldsymbol{\alpha}})^\top (\mathbf{y} - \mathbf{K}\hat{\boldsymbol{\alpha}})$ . Then  $L'_\infty(\hat{\boldsymbol{\alpha}}) = -\mathbf{K} \cdot \mathbb{I}_\infty(\mathbf{y} - \mathbf{K}\hat{\boldsymbol{\alpha}})$  and higher order derivatives are zero. With  $\Delta\hat{\boldsymbol{\alpha}} = \eta \cdot \mathbb{I}_\infty(\mathbf{y} - \mathbf{K}\hat{\boldsymbol{\alpha}})$ , according to Taylor's theorem, for  $\eta \leq \min_{1 \leq i \leq n} (|\mathbf{y} - \mathbf{K}\hat{\boldsymbol{\alpha}}|_i)$

$$\begin{aligned} L_\infty(\hat{\boldsymbol{\alpha}}_{k+1}) - L_\infty(\hat{\boldsymbol{\alpha}}_k) &= L_\infty(\hat{\boldsymbol{\alpha}}_k + \Delta\hat{\boldsymbol{\alpha}}_k) - L_\infty(\hat{\boldsymbol{\alpha}}_k) = L_\infty(\hat{\boldsymbol{\alpha}}_k) + L'(\hat{\boldsymbol{\alpha}}_k)^\top \Delta\hat{\boldsymbol{\alpha}} \\ &= -\eta \cdot \mathbb{I}_\infty(\mathbf{y} - \mathbf{K}\hat{\boldsymbol{\alpha}})^\top \mathbf{K} \cdot \mathbb{I}_\infty(\mathbf{y} - \mathbf{K}\hat{\boldsymbol{\alpha}}) \leq 0. \end{aligned} \quad \square$$

*Proof of Proposition A.1.* Let  $\mathbf{f}^+ := [\hat{\mathbf{f}}^\top, \hat{\mathbf{f}}^{*\top}]^\top$ ,  $\mathbf{y}^+ := [\mathbf{y}^\top, \tilde{\mathbf{y}}^\top]^\top$ , and let

$$\hat{\mathbf{I}} := [\mathbf{I}_{n \times n} \quad \mathbf{0}_{n \times n^*}] \in \mathbb{R}^{n \times (n+n^*)}$$

denote the training data selection matrix, so that  $\mathbf{f} = \hat{\mathbf{I}}\mathbf{f}^+$  and  $\mathbf{K} = \hat{\mathbf{I}} \cdot [\mathbf{K}^\top, \mathbf{K}^{*\top}]^\top = \hat{\mathbf{I}}\mathbf{K}^{**}\hat{\mathbf{I}}^\top$ .

Differentiating Equation 17a with respect  $\mathbf{f}^+$  and setting the gradient to  $\mathbf{0}$ , we obtain

$$\begin{aligned} \mathbf{0} &= \frac{\partial}{\partial \hat{\mathbf{f}}^+} \left( \frac{1}{2} \|\mathbf{y} - \hat{\mathbf{I}}\mathbf{f}^+\|_2^2 + \frac{\lambda}{2} \|\mathbf{f}^+\|_{(\mathbf{K}^{**})^{-1}}^2 \right) \\ &= \hat{\mathbf{I}}^\top (\hat{\mathbf{I}}\hat{\mathbf{f}}^+ - \mathbf{y}) + \lambda(\mathbf{K}^{**})^{-1} \hat{\mathbf{f}}^+ \\ \iff \hat{\mathbf{f}}^+ &= \left( \hat{\mathbf{I}}^\top \hat{\mathbf{I}} + \lambda(\mathbf{K}^{**})^{-1} \right)^{-1} \hat{\mathbf{I}}^\top \mathbf{y}. \end{aligned}$$

According to the matrix inversion lemma,

$$(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{C}\mathbf{A}^{-1} = \mathbf{D}^{-1}\mathbf{C}(\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1}.$$

For  $\mathbf{A} = \mathbf{I}$ ,  $\mathbf{B} = \hat{\mathbf{I}}$ ,  $\mathbf{C} = \hat{\mathbf{I}}^\top$  and  $\mathbf{D} = \lambda(\mathbf{K}^{**})^{-1}$ , we obtain

$$\begin{aligned} \left( \lambda(\mathbf{K}^{**})^{-1} + \hat{\mathbf{I}}^\top \hat{\mathbf{I}} \right)^{-1} \hat{\mathbf{I}}^\top &= 1/\lambda \underbrace{\mathbf{K}^{**}\hat{\mathbf{I}}^\top}_{=[\mathbf{K}^\top, \mathbf{K}^{*\top}]^\top} \left( \mathbf{I} + 1/\lambda \underbrace{\hat{\mathbf{I}}\mathbf{K}^{**}\hat{\mathbf{I}}^\top}_{=\mathbf{K}} \right)^{-1} \\ &= \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\lambda\mathbf{I} + \mathbf{K})^{-1}, \end{aligned}$$

and thus

$$\hat{\mathbf{f}}^+ = \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\mathbf{K} + \lambda\mathbf{I})^{-1} \mathbf{y}.$$

Before differentiating Equation 17c with respect to  $\mathbf{f}^+$ , we first note that according to the definition of  $\tilde{\mathbf{y}}$ ,

$$\mathbf{y}^+ - \mathbf{f}^+ = \begin{bmatrix} \mathbf{y} - \hat{\mathbf{I}}\mathbf{f}^+ \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{n \times n} \\ \mathbf{0}_{n^* \times n} \end{bmatrix} (\mathbf{y} - \hat{\mathbf{I}}\mathbf{f}^+) = \hat{\mathbf{I}}^\top \mathbf{y} - \hat{\mathbf{I}}^\top \hat{\mathbf{I}}\mathbf{f}^+.$$



Now, differentiating Equation 17c with respect  $\mathbf{f}^+$  and setting the gradient to  $\mathbf{0}$  we obtain

$$\begin{aligned}
 \mathbf{0} &= \frac{\partial}{\partial \hat{\mathbf{f}}^+} \left( \frac{1}{2} \|\mathbf{y}^+ - \hat{\mathbf{f}}^+\|_{K^{**}}^2 + \frac{\lambda}{2} \|\hat{\mathbf{f}}^+\|_2^2 \right) \\
 &= -K^{**} \mathbf{y}^+ + K^{**} \hat{\mathbf{f}}^+ + \lambda \hat{\mathbf{f}}^+ \\
 &= -K^{**} (\mathbf{y}^+ - \hat{\mathbf{f}}^+) + \lambda \hat{\mathbf{f}}^+ = -K^{**} (\hat{\mathbf{I}}^\top \mathbf{y} - \hat{\mathbf{I}}^\top \hat{\mathbf{I}} \hat{\mathbf{f}}^+) + \lambda \hat{\mathbf{f}}^+ \\
 &= -K^{**} \hat{\mathbf{I}}^\top \mathbf{y} + (K^{**} \hat{\mathbf{I}}^\top \hat{\mathbf{I}} + \lambda I) \hat{\mathbf{f}}^+ \\
 \iff \hat{\mathbf{f}}^+ &= \left( \begin{bmatrix} K \\ K^* \end{bmatrix} \hat{\mathbf{I}} + \lambda I \right)^{-1} \begin{bmatrix} K \\ K^* \end{bmatrix} \mathbf{y},
 \end{aligned}$$

Using the matrix inversion lemma,

$$(D - CA^{-1}B)^{-1}CA^{-1} = D^{-1}C(A - BD^{-1}C)^{-1},$$

with  $A = I$ ,  $B = \hat{\mathbf{I}}$ ,  $C = [K^\top, K^{*\top}]^\top$  and  $D = \lambda I$ , we obtain

$$\begin{aligned}
 \left( \lambda I + \begin{bmatrix} K \\ K^* \end{bmatrix} \hat{\mathbf{I}} \right)^{-1} \begin{bmatrix} K \\ K^* \end{bmatrix} &= 1/\lambda \begin{bmatrix} K \\ K^* \end{bmatrix} \left( I + 1/\lambda \underbrace{\hat{\mathbf{I}} \begin{bmatrix} K \\ K^* \end{bmatrix}}_{=K} \right)^{-1} \\
 &= \begin{bmatrix} K \\ K^* \end{bmatrix} (\lambda I + K)^{-1},
 \end{aligned}$$

and thus

$$\hat{\mathbf{f}}^+ = \begin{bmatrix} K \\ K^* \end{bmatrix} (K + \lambda I)^{-1} \mathbf{y}. \quad \square$$

*Proof of Equation 19.* With  $\hat{\boldsymbol{\eta}}(t) = \hat{\mathbf{f}}(t) - \mathbf{y}$  we obtain  $\frac{d\hat{\boldsymbol{\eta}}(t)}{dt} = \frac{d\hat{\mathbf{f}}(t)}{dt}$  and the first part of Equation 18 can be written as

$$\frac{d\hat{\boldsymbol{\eta}}(t)}{dt} = -K\hat{\boldsymbol{\eta}}(t) \iff \hat{\boldsymbol{\eta}}(t) = \exp(-tK) \hat{\boldsymbol{\eta}}_0.$$

Now,

$$\hat{\mathbf{f}}(0) = \mathbf{0} \implies \hat{\boldsymbol{\eta}}_0 = \hat{\boldsymbol{\eta}}(0) = -\mathbf{y} \implies \hat{\boldsymbol{\eta}}(t) = -\exp(-tK)\mathbf{y}$$

Solving for  $\hat{\mathbf{f}}(t)$ , we obtain

$$\hat{\mathbf{f}}(t) = (I - \exp(-tK)) \mathbf{y}.$$

Finally,

$$\begin{aligned}
 \frac{d\hat{\mathbf{f}}^*(t)}{dt} &= K^* (\mathbf{y} - \hat{\mathbf{f}}(t)) = K^* \exp(-tK) \mathbf{y} \\
 \implies \hat{\mathbf{f}}^*(t) &= \mathbf{c} - K^* K^{-1} \exp(-tK) \mathbf{y}.
 \end{aligned}$$

Now,

$$\hat{\mathbf{f}}^*(0) = \mathbf{0} \implies \mathbf{c} = K^* K^{-1} \mathbf{y} \implies \hat{\mathbf{f}}^*(t) = K^* K^{-1} (I - \exp(-tK)) \mathbf{y}. \quad \square$$

*Proof of Equation 19 with momentum and Nesterov accelerated gradient.* Analogous to the case of  $\hat{\mathbf{a}}$  in the proof of the Remark 4.2, for momentum and Nesterov accelerated gradient, Equation 18 generalizes into

$$(1 - \gamma) \cdot \left[ \frac{d\hat{\mathbf{f}}(t)}{dt} \right] = \begin{bmatrix} \mathbf{K} \\ \mathbf{K}^* \end{bmatrix} (\mathbf{y} - \hat{\mathbf{f}}(t)).$$

Solving the differential equations in the same way as for Equation 19, we obtain

$$\begin{bmatrix} \hat{\mathbf{f}}(t) \\ \hat{\mathbf{f}}^*(t) \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{K}^* \mathbf{K}^{-1} \end{bmatrix} \left( \mathbf{I} - \exp \left( -\frac{t}{1-\gamma} \mathbf{K} \right) \right) \mathbf{y}. \quad \square$$

*Proof of Proposition A.2.* When  $\mathbf{K}^{**}$  is a diagonal matrix with elements  $\{k_{ii}\}_{i=1}^{n+n^*}$ ,  $k_{ii} > 0$ ,

$$\|\mathbf{y}^+ - \mathbf{f}^+\|_{\mathbf{K}^{**}}^2 = \sum_{i=1}^{n+n^*} \left( k_{ii} (y_i^+ - f_i^+) \right)^2.$$

The constraint  $\|\mathbf{f}^+\|_\infty = \max_i |f_i^+| \leq c$  is equivalent to  $|f_i^+| \leq c$  for  $i = 1, 2, \dots, n + n^*$ .

Thus,

$$\|\mathbf{y}^+ - \mathbf{f}^+\|_{\mathbf{K}^{**}}^2 \text{ s.t. } \|\mathbf{f}^+\|_\infty \leq c \iff \sum_{i=1}^{n+n^*} \left( k_{ii} (y_i^+ - f_i^+) \right)^2 \text{ s.t. } |f_i^+| \leq c \forall i,$$

which decomposes element-wise to

$$\hat{f}_i^+ = \arg \min_{f_i^+ \in \mathbb{R}} \left( k_{ii} (y_i^+ - f_i^+) \right)^2 \text{ s.t. } |f_i^+| \leq c. \quad (25)$$

We solve Equation 25 separately for the two cases  $i \leq n$  and  $i > n$ . For  $i > n$ , we have defined  $y_i^+$  to be a copy of  $f_i^+$ , which means that the reconstruction error is always 0 and Equation 25 reduces to

$$\begin{aligned} \hat{f}_i^+ &= \arg \min_{f_i^+ \in \mathbb{R}} \left( k_{ii} (y_i^+ - f_i^+) \right)^2 \text{ s.t. } |f_i^+| \leq c \\ &= \arg \min_{f_i^+ \in \mathbb{R}} \left( k_{ii} (y_i^+ - f_i^+) \right)^2 + \lambda_i |f_i^+| = \lambda_i |f_i^+|, \end{aligned} \quad (26)$$

for  $\lambda_i = \max(0, k_{ii}(y_i^+ - c))$ , where the first equality is due to Lagrangian duality, and the second is due to the reconstruction error being 0.

First assume  $\lambda_i = 0$ , in which case Equation 26 reduces to the ill-posed problem  $\hat{f}_i^+ = \arg \min_{f_i^+ \in \mathbb{R}} 0$ . However, in this case, we can use Equation 3 with  $\lambda = 0$ . Since a diagonal  $\mathbf{K}^{**}$  implies  $\mathbf{K}^* = \mathbf{0}$ , we obtain, for  $i > n$ ,  $\hat{f}_i^+ = 0$ . If,

on the other hand,  $\lambda_i > 0$ , then  $\hat{f}_i^+ = \arg \min_{f_i^+ \in \mathbb{R}} \lambda_i |f_i^+| = 0$ . Thus, for  $i > n$ , regardless of  $\lambda_i$ ,  $\hat{f}_i^+ = 0$ .

For  $i \leq n$ , we have  $y_i^+ = y_i$ . Assume  $y_i \geq 0$ . Then, the optimal value for  $f_i^+$  is  $y_i$ , unless  $y_i > c$ , then, due to convexity, the optimal value is  $c$ , i.e.  $\hat{f}_i^+ = \min(y_i, c)$ . Accounting also for the case of  $y_i < 0$ , we obtain

$$\hat{f}_i^+(c) = \begin{cases} \text{sign}(y_i) \cdot \min(|y_i|, c) & \text{if } i \leq n \\ 0 & \text{if } i > n. \end{cases}$$

The sign gradient flow solution is calculated as follows:

$$\begin{aligned} \text{sign} \left( -\frac{\partial}{\partial \hat{\mathbf{f}}^+(t)} \left( \left\| \mathbf{y}^+ - \hat{\mathbf{f}}^+(t) \right\|_{\mathbf{K}^{**}}^2 \right) \right) &= \text{sign} \left( \mathbf{K}^{**}(\mathbf{y}^+ - \hat{\mathbf{f}}^+(t)) \right) \\ &= \text{sign} \left( \begin{bmatrix} k_{ii}(y_i^+ - \hat{f}_i^+(t)) \\ \mathbf{0} \end{bmatrix} \right) = \text{sign} \left( \begin{bmatrix} [k_{ii}(y_i - \hat{f}_i(t))]_{i=1}^n \\ \mathbf{0} \end{bmatrix} \right) \\ &= \begin{bmatrix} [\text{sign}(k_{ii}(y_i - \hat{f}_i(t)))]_{i=1}^n \\ \mathbf{0} \end{bmatrix}, \end{aligned}$$

where  $[\text{sign}(k_{ii}(y_i - \hat{f}_i(t)))]_{i=1}^n \in \mathbb{R}^n$  and  $\mathbf{0} \in \mathbb{R}^{n^*}$ . Since element  $i$  in the vector depends on  $\hat{\mathbf{f}}^+$  only through  $\hat{f}_i^+$ , and since  $k_{ii} > 0$ ,

$$\begin{aligned} \frac{\partial \hat{\mathbf{f}}^+(t)}{\partial t} &= \text{sign} \left( -\frac{\partial}{\partial \hat{\mathbf{f}}^+(t)} \left( \left\| \mathbf{y}^+ - \hat{\mathbf{f}}^+(t) \right\|_{\mathbf{K}^{**}}^2 \right) \right) \\ \iff \frac{\partial \hat{f}_i^+(t)}{\partial t} &= \begin{cases} \text{sign}(k_{ii}(y_i - \hat{f}_i(t))) & \text{if } i \leq n \\ 0 & \text{if } i > n \end{cases} \\ &= \begin{cases} 1 & \text{if } i \leq n \text{ and } y_i > \hat{f}_i^+(t) \\ -1 & \text{if } i \leq n \text{ and } y_i < \hat{f}_i^+(t) \\ 0 & \text{if } i > n \text{ or } y_i = \hat{f}_i^+(t). \end{cases} \end{aligned}$$

Thus, since  $\hat{f}_i^+(0) = 0$ ,  $\hat{f}_i^+(t) = 0$ , for  $i > n$ . For  $i \leq n$ ,  $\hat{f}_i^+(t) = \pm t$ , depending on the sign of  $y_i$ . Once  $\hat{f}_i^+ = y_i$  (and  $\text{sign}(k_{ii}(y_i - \hat{f}_i^+) = 0$ ), then  $\hat{f}_i^+ = y$ :

$$\hat{f}_i^+(t) = \begin{cases} \text{sign}(y_i) \cdot \min(t, |y_i|) & \text{if } i \leq n \\ 0 & \text{if } i > n. \end{cases} \quad \square$$

## Funding

The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

## References

- Ali, A., Kolter, J. Z., and Tibshirani, R. J. (2019). A continuous-time view of early stopping for least squares regression. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1370–1378. PMLR.
- Ali, M., Prasad, R., Xiang, Y., and Yaseen, Z. M. (2020). Complete ensemble empirical mode decomposition hybridized with random forest and kernel ridge regression model for monthly rainfall forecasts. *Journal of Hydrology*, 584:124647.
- Allerbo, O., Jonasson, J., and Jörnsten, R. (2023). Elastic gradient descent, an iterative optimization method approximating the solution paths of the elastic net. *Journal of Machine Learning Research*, 24(277):1–53. [MR4664714](#)
- Balles, L. and Hennig, P. (2018). Dissecting Adam: The sign, magnitude and variance of stochastic gradients. In *International Conference on Machine Learning*, pages 404–413. PMLR.
- Bondell, H. D. and Reich, B. J. (2008). Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR. *Biometrics*, 64(1):115–123. [MR2422825](#)
- Chen, H. and Leclair, J. (2021). Optimizing etching process recipe based on kernel ridge regression. *Journal of Manufacturing Processes*, 61:454–460.
- Chen, Z., Hu, J., Qiu, X., and Jiang, W. (2022). Kernel ridge regression-based TV regularization for motion correction of dynamic MRI. *Signal Processing*, 197:108559. [MR4577146](#)
- De Brabanter, K., Pelckmans, K., De Brabanter, J., Debruyne, M., Suykens, J. A., Hubert, M., and De Moor, B. (2009). Robustness of kernel based regression: a comparison of iterative weighting schemes. In *Artificial Neural Networks–ICANN 2009: 19th International Conference, Limassol, Cyprus, September 14–17, 2009, Proceedings, Part I 19*, pages 100–110. Springer.
- Debruyne, M., Christmann, A., Hubert, M., and Suykens, J. A. (2010). Robustness of reweighted least squares kernel based regression. *Journal of Multivariate Analysis*, 101(2):447–463. [MR2564353](#)
- Demontis, A., Biggio, B., Fumera, G., Giacinto, G., Roli, F., et al. (2017). Infinity-norm support vector machines against adversarial label contamination. In *CEUR Workshop Proceedings*, volume 1816, pages 106–115. CEUR-WS.
- Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al. (2004). Least angle regression. *Annals of Statistics*, 32(2):407–499. [MR2060166](#)
- Fan, P., Deng, R., Qiu, J., Zhao, Z., and Wu, S. (2021). Well logging curve reconstruction based on kernel ridge regression. *Arabian Journal of Geosciences*, 14(16):1–10.
- Feng, Y., Lv, S.-G., Hang, H., and Suykens, J. A. (2016). Kernelized elastic net regularization: generalization bounds, and sparse recovery. *Neural Computation*, 28(3):525–562. [MR3867099](#)
- Friedman, J. and Popescu, B. E. (2004). Gradient directed regularization. *Unpublished manuscript*.
- Guigue, V., Rakotomamonjy, A., and Canu, S. (2005). Kernel basis pursuit. In

- Machine Learning: ECML 2005*, pages 146–157. Springer. [MR2249863](#)
- Hamidieh, K. (2018). A data-driven statistical model for predicting the critical temperature of a superconductor. *Computational Materials Science*, 154:346–354.
- Hastie, T., Taylor, J., Tibshirani, R., Walther, G., et al. (2007). Forward stage-wise regression and the monotone lasso. *Electronic Journal of Statistics*, 1:1–29. [MR2312144](#)
- Hinton, G., Srivastava, N., and Swersky, K. (2012). Neural networks for machine learning lecture 6a overview of mini-batch gradient descent. *Unpublished result*, <http://www.cs.toronto.edu/~hinton/coursera/lecture6/lec6.pdf>.
- Hwang, C. and Shim, J. (2005). A simple quantile regression via support vector machine. In *Advances in Natural Computation: First International Conference, ICNC 2005, Changsha, China, August 27-29, 2005, Proceedings, Part I*, pages 512–520. Springer.
- Hwang, S., Kim, D., Jeong, M. K., and Yum, B.-J. (2015). Robust kernel-based regression with bounded influence for outliers. *Journal of the Operational Research Society*, 66(8):1385–1398.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint [arXiv:1412.6980](#)*.
- Krige, D. G. (1951). A statistical approach to some basic mine valuation problems on the Witwatersrand. *Journal of the Southern African Institute of Mining and Metallurgy*, 52(6):119–139.
- Le, Y., Jin, S., Zhang, H., Shi, W., and Yao, H. (2021). Fingerprinting indoor positioning method based on kernel ridge regression with feature reduction. *Wireless Communications and Mobile Computing*, 2021:6631585.
- Li, Y., Liu, Y., and Zhu, J. (2007). Quantile regression in reproducing kernel Hilbert spaces. *Journal of the American Statistical Association*, 102(477):255–268. [MR2293307](#)
- Mason, L., Baxter, J., Bartlett, P., and Frean, M. (1999). Boosting algorithms as gradient descent. *Advances in Neural Information Processing Systems*, 12.
- Matheron, G. (1963). Principles of geostatistics. *Economic Geology*, 58(8):1246–1266.
- Mercer, J. (1909). XVI. functions of positive and negative type, and their connection the theory of integral equations. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 209(441-458):415–446.
- Nesterov, Y. (1983). A method for unconstrained convex minimization problem with the rate of convergence  $\mathcal{O}(1/k^2)$ . In *Dokl. Akad. Nauk SSSR*, volume 269, pages 543–547. [MR0701288](#)
- Pace, R. K. and Barry, R. (1997). Sparse spatial autoregressions. *Statistics & Probability Letters*, 33(3):291–297.
- Polyak, B. T. (1964). Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17. [MR0169403](#)
- Richards, D., Dobriban, E., and Rebeschini, P. (2021). Comparing classes of

- estimators: When does gradient descent beat ridge regression in linear models? *arXiv preprint [arXiv:2108.11872](#)*.
- Rockafellar, R. T. (1976). Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898. [MR0410483](#)
- Roth, V. (2004). The generalized LASSO. *IEEE Transactions on Neural Networks*, 15(1):16–28.
- Russu, P., Demontis, A., Biggio, B., Fumera, G., and Roli, F. (2016). Secure kernel machines against evasion attacks. In *Proceedings of the 2016 ACM Workshop on Artificial Intelligence and Security*, pages 59–69.
- Safari, M. J. S. and Rahimzadeh Arashloo, S. (2021). Kernel ridge regression model for sediment transport in open channel flow. *Neural Computing and Applications*, 33(17):11255–11271.
- Sathishkumar, V. E., Shin, C., and Cho, Y. (2021). Efficient energy consumption prediction model for a data analytic-enabled industry building in a smart city. *Building Research & Information*, 49(1):127–143.
- Shahsavari, A., Jamei, M., and Karbasi, M. (2021). Experimental evaluation and development of predictive models for rheological behavior of aqueous Fe<sub>3</sub>O<sub>4</sub> ferrofluid in the presence of an external magnetic field by introducing a novel grid optimization based-Kernel ridge regression supported by sensitivity analysis. *Powder Technology*, 393:1–11.
- Singh Alvarado, J., Goffinet, J., Michael, V., Liberti, W., Hatfield, J., Gardner, T., Pearson, J., and Mooney, R. (2021). Neural dynamics underlying birdsong practice and performance. *Nature*, 599(7886):635–639.
- Takeuchi, I., Le, Q. V., Sears, T. D., and Smola, A. J. (2006). Nonparametric quantile estimation. *Journal of Machine Learning Research*, 7(45):1231–1264. [MR2274404](#)
- Tang, Q., Gu, Y., and Wang, B. (2024). fastkqr: A fast algorithm for kernel quantile regression. *arXiv preprint [arXiv:2408.05393](#)*.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288. [MR1379242](#)
- Tibshirani, R. J. (2015). A general framework for fast stagewise algorithms. *Journal of Machine Learning Research*, 16(1):2543–2588. [MR3450516](#)
- Wibowo, A. (2009). Robust kernel ridge regression based on M-estimation. *Computational Mathematics and Modeling*, 20(4):438–446. [MR2660712](#)
- Wood, S. N., Li, Z., Shaddick, G., and Augustin, N. H. (2017). Generalized additive models for gigadata: Modeling the UK black smoke network daily data. *Journal of the American Statistical Association*, 112(519):1199–1210. [MR3735370](#)
- Wu, Y., Prezhdo, N., and Chu, W. (2021). Increasing efficiency of nonadiabatic molecular dynamics by Hamiltonian interpolation with kernel ridge regression. *The Journal of Physical Chemistry A*, 125(41):9191–9200.
- Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67. [MR2212574](#)

- Zahrt, A. F., Henle, J. J., Rose, B. T., Wang, Y., Darrow, W. T., and Denmark, S. E. (2019). Prediction of higher-selectivity catalysts by computer-driven workflow and machine learning. *Science*, 363(6424):eaau5631.
- Zheng, S. (2022). Fast quantile regression in reproducing kernel Hilbert space. *Journal of the Korean Statistical Society*, 51(2):568–588. [MR4422521](#)
- Zhou, Y., Jin, R., and Hoi, S. C.-H. (2010). Exclusive lasso for multi-task feature selection. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 988–995.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476):1418–1429. [MR2279469](#)