
Mind the spikes: Benign overfitting of kernels and neural networks in fixed dimension

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 The success of over-parameterized neural networks trained to near-zero training
2 error has caused great interest in the phenomenon of benign overfitting, where
3 estimators are statistically consistent even though they interpolate noisy training
4 data. While benign overfitting in fixed dimension has been established for some
5 learning methods, current literature suggests that for regression with typical kernel
6 methods and wide neural networks, benign overfitting requires a high-dimensional
7 setting where the dimension grows with the sample size. In this paper, we show
8 that the smoothness of the estimators, and not the dimension, is the key: benign
9 overfitting is possible if and only if the estimator’s derivatives are large enough.
10 We generalize existing inconsistency results to non-interpolating models and more
11 kernels to show that benign overfitting with moderate derivatives is impossible
12 in fixed dimension. Conversely, we show that benign overfitting is possible for
13 regression with a sequence of spiky-smooth kernels with large derivatives. Using
14 neural tangent kernels, we translate our results to wide neural networks. We prove
15 that while infinite-width networks do not overfit benignly with the ReLU activation,
16 this can be fixed by adding small high-frequency fluctuations to the activation
17 function. Our experiments verify that such neural networks, while overfitting, can
18 indeed generalize well even on low-dimensional data sets.

19 1 Introduction

20 While neural networks have shown great practical success, our theoretical understanding of their
21 generalization properties is still limited. A promising line of work considers the phenomenon of
22 benign overfitting, where researchers try to understand when and how models that interpolate noisy
23 training data can generalize (Zhang et al., 2021, Belkin et al., 2018, 2019). In the high-dimensional
24 regime, where the dimension grows with the number of sample points, consistency of minimum-norm
25 interpolants has been established for linear models and kernel regression (Hastie et al., 2022, Bartlett
26 et al., 2020, Liang and Rakhlin, 2020, Bartlett et al., 2021). In fixed dimension, minimum-norm
27 interpolation with standard kernels is inconsistent (Rakhlin and Zhai, 2019, Buchholz, 2022).

28 In this paper, we shed a differentiated light on benign overfitting with kernels and neural networks. We
29 argue that the dimension-dependent perspective does not capture the full picture of benign overfitting.
30 In particular, we show that harmless interpolation with kernel methods and neural networks is possible,
31 even in small fixed dimension, with adequately designed kernels and activation functions. The key is
32 to properly design estimators of the form ‘signal+spike’. While minimum-norm criteria have widely
33 been considered a useful inductive bias, we demonstrate that designing unusual norms can resolve the
34 shortcomings of standard norms. For wide neural networks, harmless interpolation can be realized by
35 adding tiny fluctuations to the activation function. Such networks do not require regularization and
36 can simply be trained to overfit (Figure 1).




Figure 1: **Spiky-smooth overfitting in 2 dimensions.** **a.** We plot the predicted function for ridgeless kernel regression with the Laplace kernel (blue) versus our spiky-smooth kernel (4) with Laplace components (orange) on \mathbb{S}^1 . The dashed black line shows the true regression function, black 'x' denote noisy training points. Further details can be found in Section 6.2. **b.** The predicted function of a trained 2-layer neural network with ReLU activation (blue) versus ReLU plus shifted high-frequency sin-function (8) (orange). Using the weights learned with the spiky-smooth activation function in a ReLU network (green) disentangles the spike component from the signal component. **c.** Training error (solid lines) and test error (dashed lines) over the course of training for b. evaluated on 10^4 test points. The dotted black line shows the optimal test error. The spiky-smooth activation function does not require regularization and can simply be trained to overfit.

37 On a technical level, we additionally prove that overfitting in kernel regression can only be consistent
 38 if the estimators have large derivatives. Using neural tangent kernels or neural network Gaussian
 39 process kernels, we can translate our results from kernel regression to the world of neural networks
 40 (Neal, 1996, Jacot et al., 2018). In particular, our results enable the design of activation functions that
 41 induce benign overfitting in fixed dimension: the spikes in kernels can be translated into infinitesimal
 42 fluctuations that can be added to an activation function to achieve harmless interpolation with neural
 43 networks. Such small high frequency oscillations can fit noisy observations without affecting the
 44 smooth component too much. Training finite neural networks with gradient descent shows that
 45 spiky-smooth activation functions can indeed achieve good generalization even when interpolating
 46 small, low-dimensional data sets (Figure 1 b,c).

47 Thanks to new technical contributions, our inconsistency results significantly extend existing ones.
 48 We use a novel noise concentration argument (Lemma D.6) to generalize existing inconsistency
 49 results on minimum-norm interpolants to the much more realistic regime of overfitting estimators with
 50 comparable Sobolev norm scaling, which includes training via gradient flow and gradient descent with
 51 “late stopping” as well as low levels of ridge regression. Moreover, a novel connection to eigenvalue
 52 concentration results for kernel matrices (Proposition 4) allows us to relax the smoothness assumption
 53 and to treat heteroscedastic noise in Theorem 5. Lastly, our Lemma E.1 translates inconsistency
 54 results from bounded open subsets of \mathbb{R}^d to the sphere \mathbb{S}^d , which leads to results for the neural tangent
 55 kernel and neural network Gaussian processes.

56 2 Setup and prerequisites

57 **General approach.** We consider a general regression problem on \mathbb{R}^d with an arbitrary, fixed dimension
 58 d and analyze kernel-based approaches to solve this problem: kernel ridge regression, kernel
 59 gradient flow and gradient descent, minimum-norm interpolation, and more generally, overfitting
 60 norm-bounded estimators. We then translate our results to neural networks via the neural network
 61 Gaussian process and the neural tangent kernel. Let us now introduce the formal framework.

62 **Notation.** We denote scalars by lowercase letters x , vectors by bold lowercase letters \mathbf{x} and matrices
 63 by bold uppercase letters \mathbf{X} . We denote the eigenvalues of \mathbf{A} as $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$ and the
 64 Moore-Penrose pseudo-inverse by \mathbf{A}^+ . We say that a probability distribution P has lower and upper
 65 bounded density if its density p satisfies $0 < c < p(\mathbf{x}) < C$ for suitable constants c, C and all \mathbf{x} on a
 66 given domain.

67 **Regression setup.** We consider a data set $D = ((\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)) \in (\mathbb{R}^d \times \mathbb{R})^n$ with i.i.d.
 68 samples $(\mathbf{x}_i, y_i) \sim P$, written as $D \sim P^n$, where P is a probability distribution on $\mathbb{R}^d \times \mathbb{R}$. We
 69 define $\mathbf{X} := (\mathbf{x}_1, \dots, \mathbf{x}_n)$ and $\mathbf{y} := (y_1, \dots, y_n)^\top \in \mathbb{R}^n$. Random variables $(\mathbf{x}, y) \sim P$ denote
 70 test points independent of D , and $P_{\mathbf{X}}$ denotes the probability distribution of \mathbf{x} . The (least squares)

71 *empirical risk* R_D and *population risk* R_P of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ are defined as

$$R_D(f) := \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2, \quad R_P(f) := \mathbb{E}_{\mathbf{x}, y} [(y - f(\mathbf{x}))^2].$$

We assume $\text{Var}(y|\mathbf{x}) < \infty$ for all \mathbf{x} . Then, R_P is minimized by the target function $f_P^*(\mathbf{x}) = \mathbb{E}[y|\mathbf{x}]$, and the *excess risk* of a function f is given by

$$R_P(f) - R_P(f_P^*) = \mathbb{E}_{\mathbf{x}} (f_P^*(\mathbf{x}) - f(\mathbf{x}))^2.$$

72 We call a data-dependent estimator f_D *consistent for* P if its excess risk converges to 0 in probability,
 73 that is, for all $\varepsilon > 0$, $\lim_{n \rightarrow \infty} P^n (D \in (\mathbb{R}^d \times \mathbb{R})^n \mid R_P(f_D) - R_P(f_P^*) \geq \varepsilon) = 0$. We call f_D
 74 *consistent in expectation for* P if $\lim_{n \rightarrow \infty} \mathbb{E}_D R_P(f_D) - R_P(f_P^*) = 0$. We call f_D *universally*
 75 *consistent* if it is consistent for all Borel probability measures P on $\mathbb{R}^d \times \mathbb{R}$.

Solutions by kernel regression. Recall that a kernel k induces a reproducing kernel Hilbert space \mathcal{H}_k , abbreviated RKHS (more details in Appendix B). For $f \in \mathcal{H}_k$, we consider the objective

$$\mathcal{L}_\rho(f) := \frac{1}{n} \sum_{i=1}^n (y_i - f(\mathbf{x}_i))^2 + \rho \|f\|_{\mathcal{H}_k}^2$$

with regularization parameter $\rho \geq 0$. Denote by $f_{t,\rho}$ the solution to this problem that is obtained by optimizing on \mathcal{L}_ρ in \mathcal{H}_k with gradient flow until time $t \in [0, \infty]$, using fixed a regularization constant $\rho > 0$, and initializing at $f = 0 \in \mathcal{H}_k$. We show in Appendix C.1 that it is given as

$$f_{t,\rho}(\mathbf{x}) := k(\mathbf{x}, \mathbf{X}) \left(\mathbf{I}_n - e^{-\frac{2}{n}t(k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n)} \right) (k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n)^{-1} \mathbf{y}, \quad (1)$$

76 where $k(\mathbf{x}, \mathbf{X})$ denotes the row vector $(k(\mathbf{x}, \mathbf{x}_i))_{i \in [n]}$ and $k(\mathbf{X}, \mathbf{X}) = (k(\mathbf{x}_i, \mathbf{x}_j))_{i,j \in [n]}$ the kernel
 77 matrix. $f_{t,\rho}$ elegantly subsumes several popular kernel regression estimators as special cases: (i)
 78 classical kernel ridge regression for $t \rightarrow \infty$, (ii) gradient flow on the unregularized objective for
 79 $\rho \searrow 0$, and (iii) kernel “ridgeless” regression $f_{\infty,0}(\mathbf{x}) = k(\mathbf{x}, \mathbf{X})k(\mathbf{X}, \mathbf{X})^+ \mathbf{y}$ in the joint limit of
 80 $\rho \rightarrow 0$ and $t \rightarrow \infty$. If $k(\mathbf{X}, \mathbf{X})$ is invertible, $f_{\infty,0}$ is the interpolating function $f \in \mathcal{H}_k$ with the
 81 smallest \mathcal{H}_k -norm.

82 **From kernels to neural networks: the neural tangent kernel (NTK) and the neural network**
 83 **Gaussian process (NNGP)**. Denote the output of a NN with parameters θ on input \mathbf{x} by $f_\theta(\mathbf{x})$.
 84 It is known that for suitable random initializations θ_0 , in the infinite-width limit the random initial
 85 function f_{θ_0} converges in distribution to a Gaussian Process with the so-called Neural Network
 86 Gaussian Process (NNGP) kernel (Neal, 1996, Lee et al., 2018, Matthews et al., 2018). In Bayesian
 87 inference, the posterior mean function is then of the form $f_{\infty,\rho}$. With minor modifications (Arora
 88 et al., 2019, Zhang et al., 2020), training infinitely wide NNs with gradient flow corresponds to
 89 learning the function $f_{t,0}$ with the neural tangent kernel (NTK) (Jacot et al., 2018, Lee et al., 2019).
 90 If only the last layer is trained, the NNGP kernel should be used instead (Daniely et al., 2016). For
 91 ReLU activation functions, the RKHS of the infinite-width NNGP and NTK on the sphere \mathbb{S}^d is
 92 typically a Sobolev space (Bietti and Bach, 2021, Chen and Xu, 2021), see Appendix B.4.

93 3 Related work

94 We here provide a short summary of related work. A more detailed account is provided in Appendix A.

95 **Kernel regression.** With appropriate regularization, kernel ridge regularization with typical universal
 96 kernels like the Gauss, Matérn, and Laplace kernels is universally consistent (Steinwart and Christ-
 97 mann, 2008, Chapter 9). Optimal rates in Sobolev RKHS can also be achieved using cross-validation
 98 of the regularization ρ (Steinwart et al., 2009) or early stopping rules (Yao et al., 2007, Raskutti et al.,
 99 2014, Wei et al., 2017). In the high-dimensional regime, the class of functions that is learnable with
 100 rotation-invariant kernels is quite limited (Donhauser et al., 2021, Ghorbani et al., 2021, Liang et al.,
 101 2020).

102 **Inconsistency results.** Besides Rakhlin and Zhai (2019) and Buchholz (2022), Beaglehole et al.
 103 (2022) derive inconsistency results for ridgeless kernel regression given assumptions on the spectral
 104 tail in the Fourier basis, and Li et al. (2023) show that polynomial convergence is impossible
 105 for common kernels including ReLU NTKs. Mallinar et al. (2022) conjecture inconsistency for

106 interpolation with ReLU NTKs based on their semi-rigorous result, which essentially assumes that
 107 the eigenfunctions can be replaced by structureless Gaussian random variables. Lai et al. (2023) show
 108 an inconsistency-type result for overfitting two-layer ReLU NNs with $d = 1$, but for fixed inputs
 109 \mathbf{X} . They also note that an earlier inconsistency result by Hu et al. (2021) relies on an unproven
 110 result. Mücke and Steinwart (2019) show that global minima of NNs can overfit both benignly and
 111 harmfully, but their result does not apply to gradient descent training. Overfitting with typical linear
 112 models around the interpolation peak is inconsistent (Ghosh and Belkin, 2022, Holzmüller, 2021).

113 **Classification.** For binary classification, benign overfitting is a more generic phenomenon than for
 114 regression (Muthukumar et al., 2021, Shamir, 2022), and consistency has been shown under linear
 115 separability assumptions (Montanari et al., 2019, Chatterji and Long, 2021, Frei et al., 2022), through
 116 complexity bounds for reference classes (Cao and Gu, 2019, Chen et al., 2019) or as long as the total
 117 variation distance of the class conditionals is sufficiently large and $f^*(\mathbf{x}) = \mathbb{E}[y|\mathbf{x}]$ lies in the RKHS
 118 with bounded norm (Liang and Recht, 2023). Chapter 8 of Steinwart and Christmann (2008) discusses
 119 how the overlap of the two classes may influence learning rates under positive regularization.

120 4 Inconsistency of overfitting with common kernel estimators

121 We consider a regression problem on \mathbb{R}^d in arbitrary, fixed dimension d that is solved by kernel
 122 regression. In this section, we derive several new results, stating that overfitting estimators with
 123 moderate Sobolev norm are inconsistent, in a variety of settings. In the next section, we establish the
 124 other direction: overfitting estimators can be consistent when we adapt the norm that is minimized.

125 4.1 Beyond minimum-norm interpolants: general overfitting estimators with bounded norm

126 Existing generalization bounds often consider the perfect minimum norm interpolant. This is a rather
 127 theoretical construction; estimators obtained by training with gradient descent algorithms merely
 128 overfit and, in the best case, approximate interpolants with small norm. In this section, we extend
 129 existing bounds to arbitrary overfitting estimators whose norm does not grow faster than the minimum
 130 norm that would be required to interpolate the training data. Before we can state the theorem, we
 131 need to establish some technical assumptions.

132 **Assumptions on the data generating process.** The following assumptions (as in Buchholz (2022))
 133 allow for quite general domains and distributions. They are standard in nonparametric statistics.

134 (D1) Let P_X be a distribution on a bounded open Lipschitz domain $\Omega \subseteq \mathbb{R}^d$ with lower and
 135 upper bounded Lebesgue density. Consider data sets $D = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$, where
 136 $\mathbf{x}_i \sim P_X$ i.i.d. and $y_i = f^*(\mathbf{x}_i) + \varepsilon_i$, where ε_i is i.i.d. Gaussian noise with positive variance
 137 $\sigma^2 > 0$ and $f^* \in C_c^\infty(\Omega) \setminus \{0\}$ denotes a smooth function with compact support.

138 **Assumptions on the kernel.** Our assumption on the kernel is that its RKHS is equivalent to a
 139 Sobolev space. For integers $s \in \mathbb{N}$, the norm of a Sobolev space $H^s(\Omega)$ can be defined as

$$\|f\|_{H^s(\Omega)}^2 := \sum_{0 \leq |\alpha| \leq s} \|D^\alpha f\|_{L_2(\Omega)}^2,$$

140 where D^α denotes partial derivatives in multi-index notation for α . It measures the magnitude
 141 of derivatives up to some order s . For general $s > 0$, $H^s(\Omega)$ is (equivalent to) an RKHS if and
 142 only if $s > d/2$. For example, Laplace and Matérn kernels (Kanagawa et al., 2018, Example 2.6)
 143 have Sobolev RKHSs. The RKHS of the Gaussian kernel $\mathcal{H}^{\text{Gauss}}$ is contained in every Sobolev
 144 space, $\mathcal{H}^{\text{Gauss}} \subsetneq H^s$ for all $s \geq 0$ (Steinwart and Christmann, 2008, Corollary 4.36). Due to its
 145 smoothness, the Gaussian kernel is potentially even more prone to harmful overfitting than Sobolev
 146 kernels (Mallinar et al., 2022). We make the following assumption on the kernel:

147 (K) Let k be a positive definite kernel function whose RKHS \mathcal{H}_k is equivalent to the Sobolev
 148 space H^s for $s \in (\frac{d}{2}, \frac{3d}{4}]$.

149 Now we are ready to state the main result of this section:

150 **Theorem 1 (Overfitting estimators with small norms are inconsistent).** *Let assumptions (D1) and*
 151 *(K) hold. Let $c_{\text{fit}} \in (0, 1]$ and $C_{\text{norm}} > 0$. Then, there exist $c > 0$ and $n_0 \in \mathbb{N}$ such that the following*
 152 *holds for all $n \geq n_0$ with probability $1 - O(1/n)$ over the draw of the data set D with n samples:*
 153 *Every function $f \in \mathcal{H}_k$ that satisfies the following two conditions*

- 154 (O) $\frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2 \leq (1 - c_{\text{fit}}) \cdot \sigma^2$ (training error of f is below Bayes risk)
 155 (N) $\|f\|_{\mathcal{H}_k} \leq C_{\text{norm}} \|f_{\infty,0}\|_{\mathcal{H}_k}$ (norm comparable to minimum-norm interpolant (1)),

has an excess risk that satisfies

$$R_P(f) - R_P(f^*) \geq c > 0. \quad (2)$$

156 In words: In fixed dimension d , every differentiable function f that overfits the training data and is
 157 not much “spikier” than the minimum RKHS-norm interpolant is inconsistent!

158 **Proof idea.** Our proof follows a similar approach as Rakhlin and Zhai (2019), Buchholz (2022), and
 159 also holds for kernels with adaptive bandwidths. For small bandwidths, $\|f_{\infty,0}\|_{L_2(P_X)}$ is too small,
 160 because $f_{\infty,0}$ decays to 0 between the training points, which shows that purely ‘spiky’ estimators are
 161 inconsistent. For all other bandwidths, interpolating $\Theta(n)$ many noisy labels y_i incurs $\Theta(1)$ error
 162 in an area of volume $\Omega(1/n)$ around $\Theta(n)$ data points with high probability, which accumulates
 163 to a total error $\Omega(1)$. Our observation is that the same logic holds when overfitting by a constant
 164 fraction. Formally, we show that f^* and f must then be separated by a constant on a constant fraction
 165 of training points, with high probability, by using the fact that a constant fraction of the total noise
 166 cannot concentrate on less than $\Theta(n)$ noise variables, with high probability (Lemma D.6). The full
 167 proof can be found in Appendix D. \square

168 Assumption (O) is necessary in Theorem 1, because optimally regularized kernel ridge regression
 169 fulfills all other assumptions of Theorem 1 while achieving consistency with minimax optimal
 170 convergence rates (see Section 3). The necessity of Assumption (N) is demonstrated by Section 5.

171 The following proposition establishes that Theorem 1 covers the entire overfitting regime of the
 172 popular (regularized) gradient flow estimators $f_{t,\rho}$ for all times $t \in [0, \infty]$ and any regularization
 173 $\rho \geq 0$. The proof in Appendix C.2 also covers gradient descent.

174 **Proposition 2 (Popular estimators fulfill the norm bound (N)).** *Let $t \in [0, \infty]$ and let $\rho \geq 0$
 175 arbitrary. Then $f_{t,\rho}$ as defined in (1) fulfills Assumption (N) with $C_{\text{norm}} = 1$.*

176 4.2 Inconsistency of overfitting with neural kernels

177 We would now like to apply the above results to neural kernels, which would allow us to translate
 178 our inconsistency results from the kernel domain to neural networks. However, to achieve this, we
 179 need to take one more technical hurdle: the equivalence results for NTKs and NNGPs only hold for
 180 probability distributions on the sphere \mathbb{S}^d (detailed summary in Appendix B.4). Lemma E.1 provides
 181 the missing technical link: It establishes a smooth correspondence between the respective kernels,
 182 Sobolev spaces, and probability distributions. The inconsistency of overfitting with (deep) ReLU
 183 NTKs and NNGP kernels then immediately follows from adapting Theorem 1 via Lemma E.1.

184 **Theorem 3 (Overfitting with neural network kernels in fixed dimension is inconsistent).** *Let
 185 $c \in (0, 1)$, and let P be a probability distribution with lower and upper bounded Lebesgue density on
 186 an arbitrary spherical cap $T := \{\mathbf{x} \in \mathbb{S}^d \mid x_{d+1} < v\} \subseteq \mathbb{S}^d$, $v \in (-1, 1)$. Let k either be*

- 187 (i) *the fully-connected ReLU NTK with 0-initialized biases of any fixed depth $L \geq 2$, and $d \geq 2$, or*
 188 (ii) *the fully-connected ReLU NNGP kernel without biases of any fixed depth $L \geq 3$, and $d \geq 6$.*

189 *Then, if $f_{t,\rho}$ fulfills Assumption (O) with probability at least c over the draw of the data set D , $f_{t,\rho}$ is
 190 inconsistent for P .*

191 Theorem 3 also holds for more general estimators as in Theorem 1, cf. the proof in Appendix E.

192 Mallinar et al. (2022) already observed empirically that overfitting common network architectures
 193 yields suboptimal generalization performance on large data sets in fixed dimension. Theorem 3 now
 194 provides a rigorous proof for this phenomenon since sufficiently wide trained neural networks and the
 195 corresponding NTKs have a similar generalization behavior (e.g. (Arora et al., 2019, Theorem 3.2)).

196 4.3 Relaxing smoothness and noise assumptions via spectral concentration bounds

197 In this section, we consider a different approach to derive lower bounds for the generalization error
 198 of overfitting kernel regression: through concentration results for the eigenvalues of kernel matrices.
 199 On a high level, we obtain similar results as in the last section. The novelty of this section is on the
 200 technical side, and we suggest that non-technical readers skip this section in their first reading.

201 We define the convolution kernel of a given kernel k as $k_*(\mathbf{x}, \mathbf{x}') := \int k(\mathbf{x}, \mathbf{x}'')k(\mathbf{x}'', \mathbf{x}') dP_X(\mathbf{x}'')$,
 202 which is possible whenever $k(\mathbf{x}, \cdot) \in L_2(P_X)$ for all \mathbf{x} . The latter condition is satisfied for bounded
 203 kernels. Our starting point is the following new lower bound:

Proposition 4 (Spectral lower bound). *Assume that the kernel matrix $k(\mathbf{X}, \mathbf{X})$ is almost surely positive definite, and that $\text{Var}(y|\mathbf{x}) \geq \sigma^2$ for P_X -almost all \mathbf{x} . Then, the expected excess risk satisfies*

$$\mathbb{E}_D R_P(f_{t,\rho}) - R_P^* \geq \frac{\sigma^2}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{X}} \frac{\lambda_i(k_*(\mathbf{X}, \mathbf{X})/n) (1 - e^{-2t(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + \rho)})^2}{(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + \rho)^2}. \quad (3)$$

204 Using concentration inequalities for kernel matrices and the relation between the integral operators of
 205 k and k_* , it can be seen that for $t = \infty$ and $\rho = 0$, every term in the sum in Eq. (3) should converge
 206 to 1 as $n \rightarrow \infty$. However, since the number of terms in the sum increases with n and the convergence
 207 may not be uniform, this is not sufficient to show inconsistency in expectation. Instead, relative
 208 concentration bounds that are even stronger than the ones by Valdivia (2018) would be required to
 209 show inconsistency in expectation. However, by combining multiple weaker bounds and further
 210 arguments on kernel equivalences, we can still show inconsistency in expectation for a class of
 211 dot-product kernels on the sphere, including certain NTK and NNGP kernels (Appendix B.4):

212 **Theorem 5 (Inconsistency for Sobolev dot-product kernels on the sphere).** *Let k be a dot-product*
 213 *kernel on \mathbb{S}^d , i.e., a kernel of the form $k(\mathbf{x}, \mathbf{x}') = \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle)$, such that its RKHS \mathcal{H}_k is equivalent*
 214 *to a Sobolev space $H^s(\mathbb{S}^d)$, $s > d/2$. Moreover, let P be a distribution on $\mathbb{S}^d \times \mathbb{R}$ such that*
 215 *P_X has a lower and upper bounded density w.r.t. the uniform distribution $\mathcal{U}(\mathbb{S}^d)$, and such that*
 216 *$\text{Var}(y|\mathbf{x}) \geq \sigma^2 > 0$ for P_X -almost all $\mathbf{x} \in \mathbb{S}^d$. Then, for every $C > 0$, there exists $c > 0$ independent*
 217 *of σ^2 such that for all $n \geq 1$, $t \in (C^{-1}n^{2s/d}, \infty]$, and $\rho \in [0, Cn^{-2s/d})$, the expected excess risk*
 218 *satisfies*

$$\mathbb{E}_D R_P(f_{t,\rho}) - R_P^* \geq c\sigma^2 > 0.$$

219 The assumptions of Theorem 5 and Theorem 3 differ in several ways. Theorem 5 applies to arbitrarily
 220 high smoothness s and therefore to ReLU NTKs and NNGPs in arbitrary dimension d . Moreover, it
 221 applies to distributions on the whole sphere and allows more general noise distributions. On the flip
 222 side, it only shows inconsistency in expectation, which we believe could be extended to inconsistency
 223 for Gaussian noise. Moreover, it only applies to functions of the form $f_{t,\rho}$ but provides an explicit
 224 bound on t and ρ to get inconsistency. For $t = \infty$, the bound $\rho = O(n^{-2s/d})$ appears to be tight, as
 225 larger ρ yield consistency for comparable Sobolev kernels on \mathbb{R}^d (Steinwart et al., 2009, Corollary 3).

226 The spectral lower bounds in Theorem F.2 show that our approach can directly benefit from developing
 227 better kernel matrix concentration inequalities. Conversely, the investigation of consistent kernel
 228 interpolation might provide information about where such concentration inequalities do not hold.

229 5 Consistency via spiky-smooth estimators – even in fixed dimension


230 In Section 4, we have seen that when common kernel estimators overfit, they are inconsistent for
 231 many kernels and a wide variety of distributions. We now design consistent interpolating kernel
 232 estimators. The key is to violate Assumption (N) and allow for quickly exploding derivatives.

233 5.1 Almost universal consistency of spiky-smooth ridgeless kernel regression

234 In high dimensional regimes (where the dimension d is supposed to grow with the number of data
 235 points), benign overfitting of linear and kernel regression has been understood by an additive decom-
 236 position of the minimum-norm interpolant into a smooth regularized component that is responsible
 237 for good generalization, and a spiky component that interpolates the noisy data points while not
 238 harming generalization (Bartlett et al., 2021). This inspires us to enforce such a decomposition in
 239 arbitrary fixed dimension by adding a sharp kernel spike $\rho \check{k}_{\gamma_n}$ to a common kernel \check{k} . In this way, we
 240 can still generate any Sobolev RKHS (see Appendix G.2).

Definition 6 (Spiky-smooth kernel). Let \tilde{k} denote any universal kernel function on \mathbb{R}^d . We call it the smooth component. Consider a second, translation invariant kernel \check{k}_γ of the form $k_\gamma(\mathbf{x}, \mathbf{y}) = q(\frac{\mathbf{x}-\mathbf{y}}{\gamma})$, for some function $q : \mathbb{R}^d \rightarrow \mathbb{R}$. We call it the spiky component. Then we define the ρ -regularized spiky-smooth kernel with spike bandwidth γ as

$$k_{\rho,\gamma}(\mathbf{x}, \mathbf{y}) = \tilde{k}(\mathbf{x}, \mathbf{y}) + \rho \cdot \check{k}_\gamma(\mathbf{x}, \mathbf{y}), \quad \mathbf{x}, \mathbf{y} \in \mathbb{R}^d. \quad (4)$$



We now show that the minimum-norm interpolant of the spiky-smooth kernel sequence with properly chosen $\rho_n, \gamma_n \rightarrow 0$ is consistent for a large class of distributions, on a space with fixed (possibly small) dimension d . We establish our result under the following assumption (as in Mücke and Steinwart (2019)), which is weaker than our previous Assumption (D1).

(D2) There exists a constant $\beta_X > 0$ and a continuous function $\phi : [0, \infty) \rightarrow [0, 1]$ with $\phi(0) = 0$ such that the data generating probability distribution satisfies $P_X(B_t(\mathbf{x})) \leq \phi(t) = O(t^{\beta_X})$ for all $\mathbf{x} \in \Omega$ and all $t \geq 0$ (here $B_t(\mathbf{x})$ denotes the Euclidean ball of radius t around \mathbf{x}).

Theorem 7 (Consistency of spiky-smooth ridgeless kernel regression). Assume that the training set D consists of n i.i.d. pairs $(\mathbf{x}, y) \sim P$ such that the marginal P_X fulfills (D2) and $\mathbb{E}y^2 < \infty$. Let the kernel components satisfy:

- \tilde{k} is a universal kernel, and $\rho_n \rightarrow 0$ and $n\rho_n^4 \rightarrow \infty$.
- \check{k}_{γ_n} denotes the Laplace kernel with a sequence of positive bandwidths (γ_n) fulfilling $\gamma_n = O(n^{-(2+\alpha)/\beta_X} / \log(n))$, where $\alpha > 0$ arbitrary.

Then the minimum-norm interpolant of the ρ_n -regularized spiky-smooth kernel sequence $k_n := k_{\rho_n, \gamma_n}$ is consistent for P .

Proof idea. With sharp spikes $\gamma \rightarrow 0$, it holds that $\check{k}_\gamma(\mathbf{X}, \mathbf{X}) \approx \mathbf{I}_n$, with high probability. Hence, ridgeless kernel regression with the spiky-smooth kernel interpolates the training set while approximating kernel ridge regression with the smooth component \tilde{k} and regularization ρ . \square

The theorem even holds under much weaker assumptions on the decay behavior of the spike component \check{k}_{γ_n} , including Gaussian and Matérn kernels. The full version of the theorem and its proof can be found in Appendix G. It also applies to kernels and distributions on the sphere \mathbb{S}^d .

5.2 From spiky-smooth kernels to spiky-smooth activation functions

So far, our discussion revolved around the properties of kernels and whether they lead to estimators that are consistent. We now turn our attention to the neural network side. The big question is whether it is possible to specifically design activation functions that enable benign overfitting in fixed, possibly small dimension. We will see that the answer is yes: similarly to adding sharp spikes to a kernel, we add tiny fluctuations to the activation function. Concretely, we exploit (Simon et al., 2022, Theorem 3.1). It states that any dot-product kernel on the sphere that is a dot-product kernel in every dimension d can be written as an NNGP kernel or an NTK of two-layer fully-connected networks with a specifically chosen activation function. Further details can be found in Appendix H.

Theorem 8 (Connecting kernels and activation functions (Simon et al., 2022)). Let $\kappa : [-1, 1] \rightarrow \mathbb{R}$ be a function such that $k_d : \mathbb{S}^d \times \mathbb{S}^d \rightarrow \mathbb{R}, k_d(\mathbf{x}, \mathbf{x}') = \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle)$ is a kernel for every $d \geq 1$. Then, there exist $b_i \geq 0$ with $\sum_{i=0}^{\infty} b_i < \infty$ such that $\kappa(t) = \sum_{i=0}^{\infty} b_i t^i$, and for any choice of signs $(s_i)_{i \in \mathbb{N}} \subseteq \{-1, +1\}$, the kernel k_d can be realized as the NNGP or NTK of a two-layer fully-connected network with activation function

$$\phi_{NNGP}^{k_d}(x) = \sum_{i=0}^{\infty} s_i (b_i)^{1/2} h_i(x), \quad \phi_{NTK}^{k_d}(x) = \sum_{i=0}^{\infty} s_i \left(\frac{b_i}{i+1} \right)^{1/2} h_i(x). \quad (5)$$

Here, h_i denotes the i -th Probabilist's Hermite polynomial normalized such that $\|h_i\|_{L_2(\mathcal{N}(0,1))} = 1$.

The following proposition justifies the approach of adding spikes $\rho^{1/2} \phi^{\check{k}_\gamma}$ to an activation function to enable harmless interpolation with wide neural networks. Here we state the result for the case of the NTK; an analogous result holds for induced NNGP activation functions.




Figure 3: **a.**, **b.** Gaussian NTK activation components $\phi_{NTK}^{\tilde{k}_\gamma}$ defined via (5) induced by the Gaussian kernel with varying bandwidth $\gamma \in [0.2, 0.1, 0.05]$ (the darker, the smaller γ) for **a.** bi-alternating signs $s_i = +1$ iff $\lfloor i/2 \rfloor$ even, and **b.** randomly iid chosen signs $s_i \sim \mathcal{U}(\{-1, +1\})$. **c.** Coefficients of the Hermite series of a Gaussian NTK activation component with varying bandwidth γ . Observe peaks at $2/\gamma$. For reliable approximations of activation functions use a truncation $\geq 4/\gamma$. The sum of squares of the coefficients follows Eq. (6). Figure I.8 visualizes NNGP activation components.

Proposition 9 (Additive decomposition of spiky-smooth activation functions). Fix $\tilde{\gamma}, \rho > 0$ arbitrary. Let $k = \tilde{k} + \rho \tilde{k}_\gamma$ denote the spiky-smooth kernel where \tilde{k} and \tilde{k}_γ are Gaussian kernels of bandwidth $\tilde{\gamma}$ and γ , respectively. Assume that we choose signs $\{s_i\}_{i \in \mathbb{N}}$ and then the activation functions ϕ_{NTK}^k , $\phi_{NTK}^{\tilde{k}}$ and $\phi_{NTK}^{\tilde{k}_\gamma}$ as in Theorem 8. Then, for $\gamma > 0$ small enough, it holds that

$$\|\phi_{NTK}^k - (\phi_{NTK}^{\tilde{k}} + \sqrt{\rho} \cdot \phi_{NTK}^{\tilde{k}_\gamma})\|_{L_2(\mathcal{N}(0,1))}^2 \leq 2^{1/2} \rho \gamma^{3/2} \exp\left(-\frac{1}{\gamma}\right) + \frac{4\pi(1 + \tilde{\gamma})\gamma}{\tilde{\gamma}}.$$

Proof idea. When the spikes are sharp enough (γ small enough), the smooth and the spiky component of the activation function are approximately orthogonal in $L_2(\mathcal{N}(0,1))$ (Figure 3c), so that the spiky-smooth activation function can be approximately additively decomposed into the smooth activation component $\phi^{\tilde{k}}$ and the spike component $\phi^{\tilde{k}_\gamma}$ responsible for interpolation. \square

To motivate why the added spike functions $\rho^{1/2} \phi^{\tilde{k}_\gamma}$ should have small amplitudes, observe that Gaussian activation components $\phi^{\tilde{k}_\gamma}$ satisfy

$$\|\phi_{NNGP}^{\tilde{k}_\gamma}\|_{L_2(\mathcal{N}(0,1))}^2 = 1, \quad \|\phi_{NTK}^{\tilde{k}_\gamma}\|_{L_2(\mathcal{N}(0,1))}^2 = \frac{\gamma}{2} \left(1 - \exp\left(-\frac{2}{\gamma}\right)\right). \quad (6)$$

Hence, the average amplitude of NNGP spike activation components $\rho^{1/2} \phi^{\tilde{k}_\gamma}$ does not depend on γ , while the average amplitude of NTK spike components decays to 0 with $\gamma \rightarrow 0$. Since consistency requires the quasi-regularization $\rho \rightarrow 0$, the spiky component of induced NTK as well as NNGP activation functions should vanish for large data sets $n \rightarrow \infty$ to achieve consistency.

6 Experiments

Now we explore how appropriate spiky-smooth activation functions might look like and whether they indeed enable harmless interpolation for trained networks of finite width on finite data sets. Further experimental results are reported in Appendix I.

6.1 What do common activation functions lack in order to achieve harmless interpolation?

To understand which properties we have to introduce into activation functions to enable harmless interpolation, we plot NTK spike components $\phi^{\tilde{k}_\gamma}$ induced by the Gaussian kernel (Figure 3a,b) as well as their Hermite series coefficients (Figure 3c). Remarkably, the spike components $\phi^{\tilde{k}_\gamma}$ approximately correspond to a shifted, high-frequency sin-curve, when choosing the signs s_i in (5) to alternate every second i , that is $s_i = +1$ iff $\lfloor i/2 \rfloor$ even (Figure 3a). We empirically determine (Appendix I.6) that the NNGP activation functions are well approximated by the fluctuation function

$$\omega_{NNGP}(x; \gamma) := \sqrt{2} \cdot \sin\left(\sqrt{2/\gamma} \cdot x + \pi/4\right) = \sin\left(\sqrt{2/\gamma} \cdot x\right) + \cos\left(\sqrt{2/\gamma} \cdot x\right), \quad (7)$$

where the last equation follows from the trigonometric addition theorem. For small bandwidths γ , the NTK activation functions are increasingly well approximated by

$$\omega_{NTK}(x; \gamma) := \sqrt{\gamma} \cdot \sin\left(\sqrt{2/\gamma} \cdot x + \pi/4\right) = \sqrt{\gamma/2} \left(\sin\left(\sqrt{2/\gamma} \cdot x\right) + \cos\left(\sqrt{2/\gamma} \cdot x\right)\right). \quad (8)$$

285 With decreasing bandwidth $\gamma \rightarrow 0$ the frequency increases, while the amplitude decreases for the
 286 NTK and remains constant for the NNGP (see Eq. (6)). Plotting equivalent spike components $\phi^{\tilde{k}\gamma}$
 287 with different choices of the signs s_i (Figure 3b and Appendix I.5) suggests that harmless interpolation
 288 requires activation functions that contain **small high-frequency oscillations** or that **explode at large**
 289 $|x|$, which only affects few neurons. The Hermite series expansion of suitable activation functions
 290 should contain **non-negligible weight spread across high-order coefficients** (Figure 3c). While
 291 Simon et al. (2022) already truncate the Hermite series of induced activation functions at order 5,
 292 Figure 3c shows that an accurate approximation of spiky-smooth activation functions requires the
 293 truncation index to be larger than $2/\gamma$. Only a careful implementation allows us to capture the
 294 high-order fluctuations in the Hermite series of the spiky activation functions. Our implementation
 295 can be found in the supplementary material.

296 6.2 Training neural networks to achieve harmless interpolation in low dimension

297 In Figure 1, we plot the results of (a) ridgeless kernel regression and (b) trained 2-layer neural
 298 networks with standard choices of kernels and activation functions (blue) as well as our spiky-smooth
 299 alternatives (orange). We trained on 15 points sampled i.i.d. from $x = (x_1, x_2) \sim \mathcal{U}(\mathbb{S}^1)$ and
 300 $y = x_1 + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, 0.25)$. The figure shows that both the Laplace kernel and standard
 301 ReLU networks interpolate the training data too smoothly in low dimension, and do not generalize
 302 well. However, our spiky-smooth kernel and neural networks with spiky-smooth activation functions
 303 achieve close to optimal generalization while interpolating the training data with sharp spikes.
 304 We achieve this by using the adjusted activation function with high-frequency oscillations $x \mapsto$
 305 $\text{ReLU}(x) + \omega_{\text{NTK}}(x; \frac{1}{5000})$ as defined in Eq. (8). With this choice, we avoid activation functions
 306 with exploding behavior, which would induce exploding gradients. Other choices of amplitude and
 307 frequency in Eq. (8) perform worse. Over the course of training (Figure 1c), the standard ReLU
 308 network exhibits harmful overfitting, whereas the NN with a spiky-smooth activation function quickly
 309 interpolates the training set with nearly optimal generalization. Training details and hyperparameter
 310 choices can be found in Appendix I.1. Although the high-frequency oscillations perturb the gradients,
 311 the NN with spiky smooth activation has a stable training trajectory using gradient descent with a large
 312 learning rate of 0.4 or stochastic gradient descent with a learning rate of 0.04. Since our activation
 313 function is the sum of two terms, we can additively decompose the network into its ReLU-component
 314 and its ω_{NTK} -component. Figure 1b and Appendix I.2 demonstrate that our interpretation of the
 315 ω_{NTK} -component as 'spiky' is accurate: The oscillations in the hidden neurons induced by ω_{NTK}
 316 interfere constructively to interpolate the noise in the training points and regress to 0 between training
 317 points. This entails immediate access to the signal component of the trained neural network in form
 318 of its ReLU-component.

319 7 Conclusion

320 Conceptually, our work shows that inconsistency of overfitting is quite a generic phenomenon for
 321 regression in fixed dimension. However, particular spiky-smooth estimators enable benign overfitting,
 322 even in fixed dimension. We translate the spikes that lead to benign overfitting in kernel regression
 323 into infinitesimal fluctuations that can be added to activation functions to consistently interpolate with
 324 wide neural networks. Our experiments verify that neural networks with spiky-smooth activation
 325 functions can exhibit benign overfitting even on small, low-dimensional data sets.

326 Technically, our inconsistency results cover many distributions, Sobolev spaces of arbitrary order, and
 327 arbitrary RKHS-norm-bounded overfitting estimators. Lemma E.1 serves as a generic tool to extend
 328 generalization bounds to the sphere \mathbb{S}^d , allowing us to cover (deep) ReLU NTKs and ReLU NNGPs.

329 **Future work.** While our experiments serve as a promising proof of concept, it remains unclear how
 330 to design activation functions that enable harmless interpolation of more complex neural network
 331 architectures and data sets. As another interesting insight, our consistent kernel sequence shows that
 332 although kernels may have equivalent RKHS (see Appendix G.2), their generalization error can differ
 333 arbitrarily much; the constants of the equivalence matter and the narrative that depth does not matter
 334 in the NTK regime as in Bietti and Bach (2021) is too simplified. More promisingly, analyses that
 335 extend our analysis in the infinite-width limit to a joint scaling of width and depth could help us to
 336 understand the influence of depth (Fort et al., 2020, Li et al., 2021, Seleznova and Kutyniok, 2022).

337 **References**

- 338 R.A. Adams and J.J.F. Fournier. *Sobolev Spaces*. Elsevier Science, 2003.
- 339 Michael Aerni, Marco Milanta, Konstantin Donhauser, and Fanny Yang. Strong inductive biases
340 provably prevent harmless interpolation. In *International Conference on Learning Representations*
341 (*ICLR*), 2023.
- 342 Mikhail S Agranovich. *Sobolev spaces, their generalizations and elliptic problems in smooth and*
343 *Lipschitz domains*. Springer, 2015.
- 344 Ludovic Arnould, Claire Boyer, and Erwan Scornet. Is interpolation benign for random forest
345 regression? In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2023.
- 346 Sanjeev Arora, Simon S. Du, Wei Hu, Zhiyuan Li, Russ R. Salakhutdinov, and Ruosong Wang.
347 On exact computation with an infinitely wide neural net. In *Advances in Neural Information*
348 *Processing Systems (NeurIPS)*, 2019.
- 349 Yaroslav Averbaynov and Alain Celisse. Early stopping and polynomial smoothing in regression with
350 reproducing kernels. *arXiv:2007.06827*, 2020.
- 351 Peter L. Bartlett, Philip M. Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear
352 regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- 353 Peter L. Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint.
354 *Acta Numerica*, 30:87–201, 2021.
- 355 Daniel Beaglehole, Mikhail Belkin, and Parthe Pandit. Kernel ridgeless regression is inconsistent in
356 low dimensions. *arXiv:2205.13525*, 2022.
- 357 Mikhail Belkin, Siyuan Ma, and Soumik Mandal. To understand deep learning we need to understand
358 kernel learning. In *International Conference on Machine Learning (ICML)*, 2018.
- 359 Mikhail Belkin, Alexander Rakhlin, and Alexandre B. Tsybakov. Does data interpolation contra-
360 dict statistical optimality? In *International Conference on Artificial Intelligence and Statistics*
361 (*AISTATS*), 2019.
- 362 Jordan Bell. The singular value decomposition of compact operators on Hilbert spaces, 2014.
- 363 Alberto Bietti and Francis Bach. Deep equals shallow for ReLU networks in kernel regimes. In
364 *International Conference on Learning Representations (ICLR)*, 2021.
- 365 Simon Buchholz. Kernel interpolation in sobolev spaces is not consistent in low dimensions. In
366 *Conference on Learning Theory (COLT)*, 2022.
- 367 Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep
368 neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- 369 Niladri S. Chatterji and Philip M. Long. Finite-sample analysis of interpolating linear classifiers in
370 the overparameterized regime. *Journal of Machine Learning Research (JMLR)*, 2021.
- 371 Lin Chen and Sheng Xu. Deep neural tangent kernel and laplace kernel have the same rkhs. In
372 *International Conference on Learning Representations (ICLR)*, 2021.
- 373 Zixiang Chen, Yuan Cao, Difan Zou, and Quanquan Gu. How much over-parameterization is
374 sufficient to learn deep ReLU networks? *arXiv:1911.12360*, 2019.
- 375 Fan Chung and Linyuan Lu. Concentration inequalities and martingale inequalities: a survey. *Internet*
376 *Mathematics*, 3(1):79 – 127, 2006.
- 377 Amit Daniely, Roy Frostig, and Yoram Singer. Toward deeper understanding of neural networks:
378 The power of initialization and a dual view on expressivity. In *Advances in Neural Information*
379 *Processing Systems (NeurIPS)*, 2016.
- 380 Ronald A DeVore and Robert C Sharpley. Besov spaces on domains in \mathbb{R}^d . *Transactions of the*
381 *American Mathematical Society*, 335(2):843–864, 1993.

- 382 Luc Devroye, Laszlo Györfi, and Adam Krzyżak. The Hilbert kernel regression estimate. *Journal of*
383 *Multivariate Analysis*, 65(2):209–227, 1998.
- 384 Eleonora Di Nezza, Giampiero Palatucci, and Enrico Valdinoci. Hitchhiker’s guide to the fractional
385 Sobolev spaces. *Bulletin des Sciences Mathématiques*, 136(5):521–573, 2012.
- 386 Konstantin Donhauser, Mingqi Wu, and Fanny Yang. How rotational invariance of common kernels
387 prevents generalization in high dimensions. In *International Conference on Machine Learning*
388 *(ICML)*, 2021.
- 389 Konstantin Donhauser, Nicolò Ruggeri, Stefan Stojanovic, and Fanny Yang. Fast rates for noisy
390 interpolation require rethinking the effect of inductive bias. In *International Conference on*
391 *Machine Learning (ICML)*, 2022.
- 392 Lutz Duembgen. Bounding standard gaussian tail probabilities. *arXiv:1012.2063*, 2010.
- 393 Stanislav Fort, Gintare Karolina Dziugaite, Mansheej Paul, Sepideh Kharaghani, Daniel M Roy,
394 and Surya Ganguli. Deep learning versus kernel learning: an empirical study of loss landscape
395 geometry and the time evolution of the neural tangent kernel. In *Advances in Neural Information*
396 *Processing Systems (NeurIPS)*, 2020.
- 397 Spencer Frei, Niladri S Chatterji, and Peter Bartlett. Benign overfitting without linearity: Neural
398 network classifiers trained by gradient descent for noisy linear data. In *Conference on Learning*
399 *Theory (COLT)*, 2022.
- 400 Semjon Gerschgorin. Über die Abgrenzung der Eigenwerte einer Matrix. *Izv. Akad. Nauk. USSR Otd.*
401 *Fiz.-Mat. Nauk*, 1931.
- 402 Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers
403 neural networks in high dimension. *The Annals of Statistics*, 49(2):1029 – 1054, 2021.
- 404 Nikhil Ghosh and Mikhail Belkin. A universal trade-off between the model size, test loss, and training
405 loss of linear predictors. *arXiv:2207.11621*, 2022.
- 406 Tilmann Gneiting. Strictly and non-strictly positive definite functions on spheres. *Bernoulli*, 19(4):
407 1327–1349, 2013.
- 408 Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J. Tibshirani. Surprises in high-
409 dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949 – 986,
410 2022.
- 411 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing
412 human-level performance on imagenet classification. In *IEEE international conference on computer*
413 *vision (ICCV)*, 2015.
- 414 David Holzmüller. On the universality of the double descent peak in ridgeless regression. In
415 *International Conference on Learning Representations (ICLR)*, 2021.
- 416 David Holzmüller and Ingo Steinwart. Training two-layer relu networks with gradient descent is
417 inconsistent. *Journal of Machine Learning Research (JMLR)*, 23(181):1–82, 2022.
- 418 Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2013.
- 419 Tianyang Hu, Wenjia Wang, Cong Lin, and Guang Cheng. Regularization matters: A nonparametric
420 perspective on overparametrized neural network. In *International Conference on Artificial*
421 *Intelligence and Statistics (AISTATS)*, 2021.
- 422 Simon Hubbert, Quoc Le Gia, and Tanya Morton. *Spherical Radial Basis Functions, Theory and*
423 *Applications*. Springer International Publishing, 2015.
- 424 Simon Hubbert, Emilio Porcu, Chris Oates, and Mark Girolami. Sobolev spaces, kernels and
425 discrepancies over hyperspheres. *arXiv:2211.09196*, 2022.

- 426 Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural Tangent Kernel: Convergence and gener-
427 alization in neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*,
428 2018.
- 429 Ziwei Ji, Justin D. Li, and Matus Telgarsky. Early-stopped neural networks are consistent. In
430 *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- 431 Fredrik Johansson et al. *mpmath: a Python library for arbitrary-precision floating-point arithmetic*
432 *(version 1.3.0)*, 2023. <https://mpmath.org/>.
- 433 Motonobu Kanagawa, Philipp Hennig, Dino Sejdinovic, and Bharath K. Sriperumbudur. Gaussian
434 processes and kernel methods: A review on connections and equivalences. *arXiv:1805.08845v1*,
435 2018.
- 436 Jianfa Lai, Manyun Xu, Rui Chen, and Qian Lin. Generalization ability of wide neural networks on
437 \mathbb{R} . *arXiv:2302.05933*, 2023.
- 438 Béatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection.
439 *The Annals of Statistics*, 28(5):1302 – 1338, 2000.
- 440 Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha
441 Sohl-Dickstein. Deep neural networks as Gaussian processes. In *International Conference on*
442 *Learning Representations (ICLR)*, 2018.
- 443 Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-
444 Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models
445 under gradient descent. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- 446 Mufan Li, Mihai Nica, and Dan Roy. The future is log-gaussian: Resnets and their infinite-depth-and-
447 width limit at initialization. In *Advances in Neural Information Processing Systems (NeurIPS)*,
448 2021.
- 449 Yicheng Li, Haobo Zhang, and Qian Lin. Kernel interpolation generalizes poorly. *arXiv:2303.15809*,
450 2023.
- 451 Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel “ridgeless” regression can generalize.
452 *Annals of Statistics*, 48(3):1329–1347, 2020.
- 453 Tengyuan Liang and Benjamin Recht. Interpolating classifiers make few mistakes. *Journal of*
454 *Machine Learning Research (JMLR)*, 24(20):1–27, 2023.
- 455 Tengyuan Liang, Alexander Rakhlin, and Xiyu Zhai. On the multiple descent of minimum-norm
456 interpolants and restricted lower isometry of kernels. In *Conference on Learning Theory (COLT)*,
457 2020.
- 458 Jacques Louis Lions and Enrico Magenes. *Non-homogeneous boundary value problems and applica-*
459 *tions: Vol. 1*. Springer Science & Business Media, 2012.
- 460 Neil Rohit Mallinar, James B Simon, Amirhesam Abedsoltan, Parthe Pandit, Misha Belkin, and
461 Preetum Nakkiran. Benign, tempered, or catastrophic: Toward a refined taxonomy of overfitting.
462 In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- 463 Alexander G. de G. Matthews, Jiri Hron, Mark Rowland, Richard E. Turner, and Zoubin Ghahramani.
464 Gaussian process behaviour in wide deep neural networks. In *International Conference on Learning*
465 *Representations (ICLR)*, 2018.
- 466 Andrew D. Mcrae, Santhosh Karnik, Mark Davenport, and Vidya K. Muthukumar. Harmless
467 interpolation in regression and classification with structured features. In *International Conference*
468 *on Artificial Intelligence and Statistics (AISTATS)*, 2022.
- 469 Robert E. Megginson. *An Introduction to Banach Space Theory*. Springer-Verlag, New York, 1998.
- 470 Leon Mirsky. On the trace of matrix products. *Mathematische Nachrichten*, 20(3-6):171–174, 1959.

- 471 Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of
472 max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime.
473 *arXiv:1911.01544*, 2019.
- 474 Vidya Muthukumar, Adhyayan Narang, Vignesh Subramanian, Mikhail Belkin, Daniel Hsu, and Anant
475 Sahai. Classification vs regression in overparameterized regimes: Does the loss function matter?
476 *JMLR*, 22(1):10104–10172, 2021.
- 477 Nicole Mücke and Ingo Steinwart. Global minima of DNNs: The plenty pantry. *arXiv:1905.10686*,
478 2019.
- 479 Radford M. Neal. *Priors for Infinite Networks*. Springer New York, 1996.
- 480 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
481 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward
482 Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,
483 Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep
484 learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- 485 Iosif Pinelis. Exact lower and upper bounds on the incomplete gamma function. *Mathematical
486 Inequalities & Applications*, 23(4):1261–1278, 2020.
- 487 Alexander Rakhlin and Xiyu Zhai. Consistency of interpolation with Laplace kernels is a high-
488 dimensional phenomenon. In *Conference on Learning Theory (COLT)*, 2019.
- 489 Akshay Rangamani, Lorenzo Rosasco, and Tomaso Poggio. For interpolating kernel machines,
490 minimizing the norm of the erm solution maximizes stability. *Analysis and Applications*, 21:193 –
491 215, 2023.
- 492 Garvesh Raskutti, Martin J. Wainwright, and Bin Yu. Early stopping and non-parametric regression:
493 An optimal data-dependent stopping rule. *Journal of Machine Learning Research (JMLR)*, 15(11):
494 335–366, 2014.
- 495 Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*.
496 MIT Press, 2005.
- 497 Hans Richter. Zur Abschätzung von Matrizennormen. *Mathematische Nachrichten*, 18(1-6):178–187,
498 1958.
- 499 Robert Schaback. Superconvergence of kernel-based interpolation. *Journal of Approximation Theory*,
500 235:1–19, 2018.
- 501 Cornelia Schneider and Nadine Große. Sobolev spaces on Riemannian manifolds with bounded
502 geometry: General coordinates traces. *Mathematische Nachrichten*, 286(16), 2013.
- 503 Mariia Seleznova and Gitta Kutyniok. Neural tangent kernel beyond the infinite-width limit: Effects
504 of depth and initialization. In *International Conference on Machine Learning (ICML)*, 2022.
- 505 Ohad Shamir. The implicit bias of benign overfitting. In *Conference on Learning Theory (COLT)*,
506 2022.
- 507 James Benjamin Simon, Sajant Anand, and Mike Deweese. Reverse engineering the neural tangent
508 kernel. In *International Conference on Machine Learning (ICML)*, 2022.
- 509 Ingo Steinwart. Consistency of support vector machines and other regularized kernel machines. *IEEE
510 Transactions on Information Theory*, 2001.
- 511 Ingo Steinwart and Andreas Christmann. *Support Vector Machines*. Springer New York, 2008.
- 512 Ingo Steinwart, Don R. Hush, and Clint Scovel. Optimal rates for regularized least squares regression.
513 In *Conference on Learning Theory (COLT)*, 2009.
- 514 Daniel W. Stroock et al. *Essentials of integration theory for analysis*. Springer, 2011.

- 515 Sattar Vakili, Michael Bromberg, Jezabel Garcia, Da-shan Shiu, and Alberto Bernacchia. Uniform
516 generalization bounds for overparameterized neural networks. *arXiv:2109.06099*, 2021.
- 517 Ernesto Araya Valdivia. Relative concentration bounds for the spectrum of kernel matrices.
518 *arXiv:1812.02108*, 2018.
- 519 Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*.
520 Cambridge University Press, 2018.
- 521 Guorong Wang, Yimin Wei, Sanzheng Qiao, Peng Lin, and Yuzhuo Chen. *Generalized Inverses:
522 Theory and Computations*. Springer, 2018.
- 523 Yutong Wang and Clayton Scott. Consistent interpolating ensembles via the manifold-Hilbert kernel.
524 In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- 525 Yuting Wei, Fanny Yang, and Martin J Wainwright. Early stopping for kernel boosting algorithms:
526 A general analysis with localized complexities. In *Advances in Neural Information Processing
527 Systems (NeurIPS)*, 2017.
- 528 Holger Wendland. *Scattered Data Approximation*. Cambridge University Press, 2005.
- 529 Xingyu Xu and Yuantao Gu. Benign overfitting of non-smooth neural networks beyond lazy training.
530 In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2023.
- 531 Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning.
532 *Constructive Approximation*, 26:289–315, 2007.
- 533 Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep
534 learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115,
535 2021.
- 536 Yaoyu Zhang, Zhi-Qin John Xu, Tao Luo, and Zheng Ma. A type of generalization error induced by
537 initialization in deep neural networks. In *Proceedings of The First Mathematical and Scientific
538 Machine Learning Conference*, 2020.

539 Appendices

540 Appendix Contents.

541	A Detailed related work	16
542	B Kernels and Sobolev spaces on the sphere	18
543	B.1 Background on Sobolev spaces	18
544	B.2 General kernel theory and notation	18
545	B.3 Dot-product kernels on the sphere	19
546	B.4 Neural kernels	20
547	C Gradient flow and gradient descent with kernels	21
548	C.1 Derivation of gradient flow and gradient descent	21
549	C.2 Gradient flow and gradient descent initialized at 0 have monotonically growing \mathcal{H} -norm	22
550	D Proof of Theorem 1	23
551	D.1 Auxiliary results for the proof of Theorem 1	27
552	E Translating between \mathbb{R}^d and \mathbb{S}^d	31
553	F Spectral lower bound	34
554	F.1 General lower bounds	34
555	F.2 Equivalences of norms and eigenvalues	37
556	F.3 Kernel matrix eigenvalue bounds	39
557	F.4 Spectral lower bound for dot-product kernels on the sphere	41
558	G Proof of Theorem 7	43
559	G.1 Auxiliary results for the proof of Theorem 7	46
560	G.2 RKHS norm bounds	47
561	H Spiky-smooth activation functions induced by Gaussian components	48
562	I Additional experimental results	51
563	I.1 Experimental details of Figure 1	51
564	I.2 Disentangling signal from noise in neural networks with spiky-smooth activation	
565	functions	51
566	I.3 Repeating the finite-sample experiments	53
567	I.4 Spiky-smooth activation functions	53
568	I.5 Isolated spike activation functions	55
569	I.6 Additive decomposition and sin-fit	57
570	I.7 Spiky-smooth kernel hyper-parameter selection	58

571 **A Detailed related work**

572 Motivated by Zhang et al. (2021) and Belkin et al. (2018), an abundance of papers have tried to
 573 grasp when and how benign overfitting occurs in different settings. Rigorous understanding is mainly
 574 restricted to linear (Bartlett et al., 2020), feature (Hastie et al., 2022) and kernel regression (Liang and
 575 Rakhlin, 2020) under restrictive distributional assumptions. In the well-specified linear setting under
 576 additional assumptions, the minimum-norm interpolant is consistent if and only if $k \ll n \ll d$, the
 577 top- k eigendirections of the covariate covariance matrix align with the signal, followed by sufficiently
 578 many ‘quasi-isotropic’ directions with eigenvalues of similar magnitude (Bartlett et al., 2020).

579 **Kernel methods.** The analysis of kernel methods is more nuanced and depends on the interplay
 580 between the chosen kernel, the choice of regularization and the data distribution. L_2 -generalization
 581 error bounds can be derived in the eigenbasis of the kernel’s integral operator (Mcrae et al., 2022),
 582 where upper bounds of the form $\sqrt{\mathbf{y}^\top k(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}}/n$ promise good generalization when the regres-
 583 sion function f^* is aligned with the dominant eigendirections of the kernel, or in other words, when
 584 $\|f^*\|_{\mathcal{H}}$ is small. Most recent work focuses on high-dimensional limits, where the data dimensionality
 585 $d \rightarrow \infty$. For $d \rightarrow \infty$, the Hilbert space and its norm change, so that consistency results that demand
 586 bounded Hilbert norm of rotation-invariant kernels do not even include simple functions like sparse
 587 products (Donhauser et al., 2021, Lemma 2.1). In the regime $d^{l+\delta} \leq n \leq d^{l+1-\delta}$, rotation-invariant
 588 (neural) kernel methods (Ghorbani et al., 2021, Donhauser et al., 2021) can in fact only learn the
 589 polynomial parts up to order l of the regression function f^* , and fully-connected NTKs do so. Liang
 590 et al. (2020) uncover a related multiple descent phenomenon in kernel regression, where the risk
 591 vanishes for most $n \rightarrow \infty$, but peaks at $n = d^i$ for all $i \in \mathbb{N}$. The slower d grows, the slower the
 592 optimal rate $n^{-\frac{1}{2i+1}}$ between the peaks. Note, however, that these bounds are only upper bounds, and
 593 whether they are optimal remains an open question to the best of our knowledge. Another recent line
 594 of work analyzes how different inductive biases, measured in $\|\cdot\|_p$ -norm minimization, $p \in [1, 2]$,
 595 (Donhauser et al., 2022) or in the filter size of convolutional kernels (Aerni et al., 2023), affects the
 596 generalization properties of minimum-norm interpolants. While the risk on noiseless training samples
 597 (bias) decreases with decreasing p or small filter size, the sensitivity to noise in the training data
 598 (variance) increases. Hence only ‘weak inductive biases’, that is large p or large filter sizes, enable
 599 harmless interpolation. Our results suggest that to achieve harmless interpolation in fixed dimension
 600 one has to construct and minimize more unusual norms than $\|\cdot\|_p$ -norms.

601 **Regularised kernel regression achieves optimal rates.** With appropriate regularization, kernel
 602 ridge regularization with typical universal kernels like the Gauss, Matérn, and Laplace kernels is
 603 universally consistent (Steinwart and Christmann, 2008, Chapter 9). Steinwart et al. (2009, Corollary
 604 6) even implies minimax optimal nonparametric rates for clipped kernel ridge regression with Sobolev
 605 kernels and $f^* \in H^\beta$ where $d/2 < \beta \leq s$ for the choice $\rho_n = n^{-2s/(2\beta+d)}$. Although f^* is not
 606 necessarily in the RKHS, KRR is adaptive and can still achieve optimal learning rates. Lower
 607 smoothness β of f^* as well as higher smoothness of the kernel should be met with faster decay of
 608 ρ_n . Optimal rates in Sobolev RKHS can also be achieved using cross-validation of the regularization
 609 ρ (Steinwart et al., 2009), early stopping rules based on empirical localized Rademacher (Raskutti
 610 et al., 2014) or Gaussian complexity (Wei et al., 2017) or smoothing of the empirical risk via kernel
 611 matrix eigenvalues (Averyanov and Celisse, 2020).

612 **Lower bounds for kernel regression.** Besides Rakhlin and Zhai (2019) and Buchholz (2022),
 613 Beaglehole et al. (2022) derive inconsistency results for kernel ridgeless regression given assumptions
 614 on the spectral tail in the Fourier basis. Mallinar et al. (2022) provide a characterization of kernel
 615 ridge regression into benign, tempered and catastrophic overfitting using a *heuristic* approximation of
 616 the risk via the kernel’s eigenspectrum, essentially assuming that the eigenfunctions can be replaced
 617 by structureless Gaussian random variables. A general lower bound for ridgeless linear regression
 618 Holzmüller (2021) predicts bad generalization near the “interpolation threshold”, where the dimension
 619 of the feature space is close to n , also known as the *double descent* phenomenon. In this regime,
 620 Ghosh and Belkin (2022) also consider overfitting by a fraction beyond the noise level and derive a
 621 lower bound for linear models.

622 **Benign overfitting in fixed dimension.** Only few works have established consistency results for
 623 interpolating models in fixed dimension. The first statistical guarantees for Nadaraya-Watson kernel

624 smoothing with singular kernels were given by Devroye et al. (1998). Optimal non-asymptotic results
625 have only been established more recently. Belkin et al. (2019) show that Nadaraya-Watson kernel
626 smoothing achieves minimax optimal convergence rates for $a \in (0, d/2)$ under smoothness assump-
627 tions on f^* , when using singular kernels such as truncated Hilbert kernels $K(u) = \|u\|_2^a \mathbb{1}_{\|u\| \leq 1}$,
628 which do not induce RKHS that only contain weakly differentiable functions (as our results do). By
629 thresholding the kernel they can adjust the amount of overfitting without affecting the generalization
630 bound. To the best of our knowledge, rigorously proving or disproving analogous bounds for kernel
631 ridge regression remains an open question. Arnould et al. (2023) show that median random forests are
632 able to interpolate consistently in fixed dimension because of an averaging effect introduced through
633 feature randomization. They conjecture consistent interpolation for Breiman random forests based on
634 numerical experiments.

635 **Classification.** For binary classification tasks, benign overfitting is a more generic phenomenon
636 than for regression tasks (Muthukumar et al., 2021, Shamir, 2022). Consistency has been shown under
637 linear separability assumptions (Montanari et al., 2019, Chatterji and Long, 2021, Frei et al., 2022)
638 and through complexity bounds with respect to reference classes like the ‘Neural Tangent Random
639 Feature’ model (Cao and Gu, 2019, Chen et al., 2019). Most recently, Liang and Recht (2023) have
640 shown that the 0-1-generalization error of minimum RKHS-norm interpolants \hat{f}_0 is upper bounded
641 by $\frac{\|\hat{f}_0\|_{\mathcal{H}}^2}{n}$ and analogously that kernel ridge regression \hat{f}_ρ generalizes as $\frac{\mathbf{y}^\top (k(\mathbf{X}, \mathbf{X}) + \rho \mathbf{I})^{-1} \mathbf{y}}{n}$, where
642 the numerator upper bounds $\|\hat{f}_\rho\|_{\mathcal{H}}^2$. Their bounds imply consistency as long as the total variation
643 distance between the class conditionals is sufficiently large and the regression function has bounded
644 RKHS-norm, and their Lemma 7 shows that the upper bound is rate optimal. Under a noise condition
645 on the regression function $f^*(\mathbf{x}) = \mathbb{E}[y|\mathbf{x}]$ for binary classification and bounded $\|f^*\|_{\mathcal{H}}$, our results
646 together with Liang and Recht (2023) reiterate the distinction between benign overfitting for binary
647 classification and inconsistent overfitting for least squares regression for a large class of distributions
648 in kernel regression over Sobolev RKHS. Chapter 8 of Steinwart and Christmann (2008) discusses
649 how the overlap of the two classes may influence learning rates under positive regularization. Using
650 Nadaraya-Watson kernel smoothing, Wang and Scott (2022) offer the first consistency result for a
651 simple interpolating ensemble method with data-independent base classifiers.

652 **Connection to neural networks.** It is known that neural networks can behave like kernel methods
653 in certain infinite-width limits. For example, the function represented by a randomly initialized NN
654 behaves like a Gaussian process with the NN Gaussian process (NNGP) kernel, which depends on
655 details such as the activation function and depth of the NN (Neal, 1996, Lee et al., 2018, Matthews
656 et al., 2018). Hence, Bayesian inference in infinitely wide NNs is GP regression, whose posterior
657 predictive mean function is of the form $f_{\infty, \rho}$, where ρ depends on the assumed noise variance.
658 Moreover, gradient flow training of certain infinitely wide NNs is similar to gradient flow training
659 with the so-called *neural tangent kernel* (NTK) (Jacot et al., 2018, Lee et al., 2019, Arora et al.,
660 2019), and the correspondence can be made exact using small modifications to the NN to remove
661 the stochastic effect of the random initial function (Arora et al., 2019, Zhang et al., 2020). In other
662 words, certain infinitely wide NNs trained with gradient flow learn functions of the form $f_{t,0}$.

663 When considering the sphere $\Omega = \mathbb{S}^d$, the NTK and NNGP kernels of fully-connected NNs are
664 dot-product kernels, i.e., $k(\mathbf{x}, \mathbf{x}') = \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle)$ for some function $\kappa : [-1, 1] \rightarrow \mathbb{R}$. Moreover, from
665 Bietti and Bach (2021) and Chen and Xu (2021) it follows that the RKHSs of typical NTK and NNGP
666 kernels for the ReLU activation function are equivalent to the Sobolev spaces $H^{(d+1)/2}(\mathbb{S}^d)$ and
667 $H^{(d+3)/2}(\mathbb{S}^d)$, respectively, cf Appendix B.4.

668 Regarding consistency, Ji et al. (2021) use the NTK correspondence to show that early-stopped
669 wide NNs for classification are universally consistent under some assumptions. On the other hand,
670 Holzmüller and Steinwart (2022) show that zero-initialized biases can prevent certain two-layer
671 ReLU NNs from being universally consistent. Lai et al. (2023) show an inconsistency-type result
672 for overfitting two-layer ReLU NNs with $d = 1$, but for fixed inputs \mathbf{X} . They also note that an
673 earlier inconsistency result by Hu et al. (2021) relies on an unproven result. Li et al. (2023) show
674 that consistency with polynomial convergence rates is impossible for minimum-norm interpolants of
675 common kernels including ReLU NTKs. Mallinar et al. (2022) conjecture tempered overfitting and
676 therefore inconsistency for interpolation with ReLU NTKs based on their semi-rigorous result and the
677 results of Bietti and Bach (2021) and Chen and Xu (2021). Xu and Gu (2023) establish consistency
678 of overfitting wide 2-layer neural networks beyond the NTK regime for binary classification in very

679 high dimension $d = \Omega(n^2)$ and for a quite restricted class of distributions (the mean difference μ of
 680 the class conditionals needs to fulfill $\mu = \Omega((d/n)^{1/4} \log^{1/4}(md/n))$ and $\mu = O((d/n)^{1/2})$).

681 **B Kernels and Sobolev spaces on the sphere**

682 **B.1 Background on Sobolev spaces**

683 We say that two Hilbert spaces $\mathcal{H}_1, \mathcal{H}_2$ are equivalent, written as $\mathcal{H}_1 \cong \mathcal{H}_2$, if they are equal as sets
 684 and the corresponding norms $\|\cdot\|_{\mathcal{H}_1}$ and $\|\cdot\|_{\mathcal{H}_2}$ are equivalent.

685 Let Ω be an open set with C^∞ boundary. In this paper, we will mainly consider ℓ_2 -balls for Ω .
 686 There are multiple equivalent ways to define a (fractional) Sobolev space $H^s(\Omega)$, $s \in \mathbb{R}_{\geq 0}$, these are
 687 equivalent in the sense that the resulting Hilbert spaces will be equivalent. For example, $H^s(\Omega)$ can
 688 be defined through restrictions of functions from $H^s(\mathbb{R}^d)$, through interpolation spaces, or through
 689 Sobolev-Slobodetski norms (see e.g. Chapter 5 and 14 in Agranovich, 2015 and Chapters 7–10
 690 in Lions and Magenes, 2012). Some requirements on Ω can be relaxed, for example to Lipschitz
 691 domains, by using more general extension operators (e.g. DeVore and Sharpley, 1993). Since
 692 our results are based on equivalent norms and not specific norms, we do not care which of these
 693 definitions is used. Further background on Sobolev spaces can be found in Adams and Fournier
 694 (2003), Wendland (2005) and Di Nezza et al. (2012).

695 **B.2 General kernel theory and notation**

696 There is a one-to-one correspondence between kernel functions k and the corresponding reproducing
 697 kernel Hilbert spaces (RKHS) \mathcal{H}_k . Mercer's theorem (Steinwart and Christmann, 2008, Theorem
 698 4.49) states that for compact Ω , continuous k and a Borel probability measure P_X on Ω whose
 699 support is Ω , the integral operator $T_{k, P_X} : L_2(P_X) \rightarrow L_2(P_X)$ given by

$$T_{k, P_X} f(\mathbf{x}) = \int_{\Omega} f(\mathbf{x}') k(\mathbf{x}, \mathbf{x}') dP_X(\mathbf{x}'),$$

700 can be decomposed into an orthonormal basis $(e_i)_{i \in I}$ of $L_2(P_X)$ and corresponding eigenvalues
 701 $(\lambda_i)_{i \in I} \geq 0$, $\lambda_i \searrow 0$, such that

$$T_{k, P_X} f = \sum_{i \in I} \lambda_i \langle f, e_i \rangle e_i, \quad f \in L_2(P_X).$$

702 We write $\lambda_i(T_{k, P_X}) := \lambda_i$. Moreover, $k(\mathbf{x}, \mathbf{x}') = \sum_{i \in I} \lambda_i e_i(\mathbf{x}) e_i(\mathbf{x}')$ converges absolutely and
 703 uniformly, and the RKHS is given by

$$\mathcal{H}_k = \left\{ \sum_{i \in I} a_i \sqrt{\lambda_i} e_i \mid \sum_{i \in I} a_i^2 < \infty \right\}. \quad (\text{B.1})$$

704 The corresponding inner product between $f = \sum_{i \in I} a_i \sqrt{\lambda_i} e_i \in \mathcal{H}$ and $g = \sum_{i \in I} b_i \sqrt{\lambda_i} e_i \in \mathcal{H}$ can
 705 then be written as

$$\langle f, g \rangle_{\mathcal{H}} = \sum_{i \in I} a_i b_i. \quad (\text{B.2})$$

We use asymptotic notation O, Ω, Θ for integers n in the following way: We write

$$\begin{aligned} f(n) = O(g(n)) &\Leftrightarrow \exists C > 0 \forall n : f(n) \leq Cg(n) \\ f(n) = \Omega(g(n)) &\Leftrightarrow g(n) = O(f(n)) \\ f(n) = \Theta(g(n)) &\Leftrightarrow f(n) = O(g(n)) \text{ and } g(n) = O(f(n)). \end{aligned}$$

Above, we require that the inequality $f(n) \leq Cg(n)$ holds for all n and not only for $n \geq n_0$. This implies that if $f(n) = \Omega(g(n))$, then f must be nonzero whenever g is nonzero. This is an important detail when arguing about equivalence of RKHSs, since it allows the following statement: If we have two kernels k, \tilde{k} with Mercer representations

$$k(\mathbf{x}, \mathbf{x}') = \sum_{i \in I} \lambda_i e_i(\mathbf{x}) e_i(\mathbf{x}')$$

$$\tilde{k}(\mathbf{x}, \mathbf{x}') = \sum_{i \in I} \tilde{\lambda}_i e_i(\mathbf{x}) e_i(\mathbf{x}')$$

706 with identical eigenfunctions e_i and eigenvalues satisfying $\lambda_i = \Theta(\tilde{\lambda}_i)$, then the associated RKHSs
707 are equivalent by (B.1) and (B.2).

708 B.3 Dot-product kernels on the sphere

709 A kernel of the form $k(\mathbf{x}, \mathbf{x}') = \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle)$ for some function κ is called *dot-product kernel*. Dot-
710 product kernels are rotationally invariant. Especially, NTKs and NNGPs of fully-connected NNs
711 restricted to the sphere \mathbb{S}^d are dot-product kernels. Moreover, kernels like the Laplace, Matérn, and
712 Gaussian kernels that only depend on the distance between their inputs are also dot-product kernels
713 when restricted to the sphere \mathbb{S}^d . Therefore, in this section, we will assume that $k : \mathbb{S}^d \times \mathbb{S}^d \rightarrow \mathbb{R}$ is a
714 dot-product kernel.

We can then leverage some convenient results from the theory of dot-product kernels on the sphere,
which are summarized in more detail by Hubbert et al. (2022). For example, a Mercer representation
of k is given by

$$k(\mathbf{x}, \mathbf{x}') = \sum_{l=0}^{\infty} \mu_l \sum_{i=1}^{N_{l,d}} Y_{l,i}(\mathbf{x}) Y_{l,i}(\mathbf{x}'),$$

715 where $\{Y_{l,1}, \dots, Y_{l,N_{l,d}}\}$ is a real orthonormal basis for the space of spherical harmonics of degree
716 l within $L_2(\mathbb{S}^d)$. Especially, the integral operator $T_{k, \mathcal{U}(\mathbb{S}^d)}$ for the uniform distribution $\mathcal{U}(\mathbb{S}^d)$ has
717 eigenvalues μ_l with multiplicity $N_{l,d}$ and unnormalized eigenfunctions $Y_{l,i}$. The RKHS of k is then
718 given by

$$\mathcal{H}_k = \left\{ \sum_{l=0}^{\infty} \sqrt{\mu_l} \sum_{i=1}^{N_{l,d}} a_{l,i} Y_{l,i} \mid \sum_{l=0}^{\infty} \sum_{i=1}^{N_{l,d}} a_{l,i}^2 < \infty \right\}.$$

719 Since the index l starts from zero, we will denote decay asymptotics for l in the form $\Theta((l+1)^{-q})$
720 and not $\Theta(l^{-q})$, cf. our definition of Θ notation in Appendix B.2.

721 **Lemma B.1** (Sobolev dot-product kernels on the sphere). *For a dot-product kernel k on \mathbb{S}^d as above,*
722 *the RKHS \mathcal{H}_k is equivalent to the Sobolev space $H^s(\mathbb{S}^d)$, $s > d/2$, if and only if $\mu_l = \Theta((l+1)^{-2s})$.*
723 *In this case, we have*

$$\lambda_i(T_{k, \mathcal{U}(\mathbb{S}^d)}) = \Theta(i^{-2s/d}).$$

724 *Proof. Step 0: Equivalence.* If $\mu_l = \Theta((l+1)^{-2s})$, it is stated in Section 3 in Hubbert et al. (2022)
725 that $\mathcal{H}_k \cong H^s(\mathbb{S}^d)$. On the other hand, if $\mu_l \neq \Theta((l+1)^{-2s})$, it is easy to see that \mathcal{H}_k is not equivalent
726 to the RKHS of a kernel with $\mu_l = \Theta((l+1)^{-2s})$. It remains to show $\lambda_i(T_{k, \mathcal{U}(\mathbb{S}^d)}) = \Theta(i^{-2s/d})$.

Step 1: Ordering the eigenvalues. Consider a permutation $\pi : \mathbb{N}_0 \rightarrow \mathbb{N}_0$ such that

$$\mu_{\pi(0)} \geq \mu_{\pi(1)} \geq \dots$$

We can then define the partial sums

$$S_l := \sum_{i=0}^l N_{\pi(i), d}.$$

727 For $S_{l-1} < i \leq S_l$, we then have $\lambda_i(T_{k, \mathcal{U}(\mathbb{S}^d)}) = \mu_{\pi(l)}$.

Step 2: Show $\pi(i) = \Theta(i)$. Let $c, C > 0$ such that $c(l+1)^{-2s} \leq \mu_l \leq C(l+1)^{-2s}$ for all $l \in \mathbb{N}_0$.
For indices $i, j \in \mathbb{N}_0$, we have the implications

$$\begin{aligned} i > j &\Rightarrow c(\pi(i)+1)^{-2s} \leq \mu_{\pi(i)} \leq \mu_{\pi(j)} \leq C(\pi(j)+1)^{-2s} \\ &\Rightarrow \pi(i)+1 \geq \left(\frac{c}{C}\right)^{1/(2s)} (\pi(j)+1). \end{aligned}$$

Therefore, for $i \geq 1$ and $j \geq 0$,

$$\pi(i)+1 \geq \left(\frac{c}{C}\right)^{1/(2s)} \max_{i' < i} (\pi(i')+1) \geq \left(\frac{c}{C}\right)^{1/(2s)} ((i-1)+1) \geq \Omega(i+1),$$

$$\pi(j) + 1 \leq \left(\frac{C}{c}\right)^{1/(2s)} \min_{j' > j} (\pi(j') + 1) \leq \left(\frac{C}{c}\right)^{1/(2s)} ((j+1) + 1) \leq O(j+1).$$

728 We can thus conclude that $\pi(i) + 1 = \Theta(i+1)$.

Step 3: Individual Eigenvalue decay. As explained in Section 2.1 in Hubbert et al. (2022), we have $N_{l,d} = \Theta((l+1)^{d-1})$. Therefore,

$$S_l = \sum_{i=0}^l \Theta((\pi(i) + 1)^{d-1}) = \sum_{i=0}^l \Theta((i+1)^{d-1}) = \Theta((l+1)^d).$$

Now, let $i \geq 1$ and let $l \in \mathbb{N}_0$ such that $S_{l-1} < i \leq S_l$. We have $i \geq \Omega(l^d)$, and $i \leq O((l+1)^d)$, which implies $i = \Theta((l+1)^d)$ since $i \geq 1$. Therefore,

$$\lambda_i = \mu_{\pi(l)} = \Theta((\pi(l) + 1)^{-2s}) = \Theta((l+1)^{-2s}) = \Theta(i^{-2s/d}). \quad \square$$

729 B.4 Neural kernels

730 Several NTK and NNGP kernels have RKHSs that are equivalent to Sobolev spaces on \mathbb{S}^d . In the
731 following cases, we can deduct this from known results:

- 732 • Consider fully-connected NNs with $L \geq 3$ layers without biases and the activation function
733 $\varphi(x) = \max\{0, x\}^m$, $m \in \mathbb{N}_0$. Especially, the case $m = 1$ corresponds to the ReLU
734 activation. Vakili et al. (2021) generalize the result by Bietti and Bach (2021) from $m = 1$
735 to $m \geq 1$, showing that the NTK-RKHS is equivalent to $H^s(\mathbb{S}^d)$ for $s = (d + 2m - 1)/2$
736 and the NNGP-RKHS is equivalent to $H^s(\mathbb{S}^d)$ for $s = (d + 2m + 1)/2$. For $m = 0$,
737 Bietti and Bach (2021) essentially show that the NNGP-RKHS is equivalent to $H^s(\mathbb{S}^d)$ for
738 $s = (d + 2^{2-L})/2$. However, all of the aforementioned result have the problem that the
739 main theorem by Bietti and Bach (2021) allows for the possibility that finitely many μ_l are
740 zero, which can change the RKHS. Using our Lemma B.2 below, it follows that all μ_l are in
741 fact nonzero for NNGPs and NTKs since they are kernels in every dimension d using the
742 same function κ independent of the dimension. Hence, the equivalences to Sobolev spaces
743 stated before are correct.
- 744 • Chen and Xu (2021) prove that the RKHS of the NTK corresponding to fully-connected
745 ReLU NNs with zero-initialized biases and $L \geq 2$ (as opposed to no biases and $L \geq 3$
746 above) layers is equivalent to the RKHS of the Laplace kernel on the sphere. Since the
747 Laplace kernel is a Matérn kernel of order $\nu = 1/2$ (see e.g. Section 4.2 in Rasmussen and
748 Williams (2005)), we can use Proposition 5.2 of Hubbert et al. (2022) to obtain equivalence
749 to $H^s(\mathbb{S}^d)$ with $s = (d+1)/2$. Alternatively, we can obtain the RKHS of the Laplace kernel
750 from Bietti and Bach (2021) combined with Lemma B.2.

751 Bietti and Bach (2021) also show that under an integrability condition on the derivatives, C^∞
752 activations induce NTK and NNGP kernels whose RKHSs are smaller than every Sobolev space.

Lemma B.2 (Guaranteeing non-zero eigenvalues). *Let $\kappa : [-1, 1] \rightarrow \mathbb{R}$, let $d \geq 1$, and let*

$$\begin{aligned} k_d : \mathbb{S}^d \times \mathbb{S}^d, k_d(\mathbf{x}, \mathbf{x}') &:= \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle) \\ k_{d+2} : \mathbb{S}^{d+2} \times \mathbb{S}^{d+2}, k_{d+2}(\mathbf{x}, \mathbf{x}') &:= \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle). \end{aligned}$$

753 *Suppose that k_{d+2} is a kernel. Then, k_d is a kernel. Moreover, if the corresponding eigenvalues μ_l
754 satisfy $\mu_l > 0$ for infinitely many l , then they satisfy $\mu_l > 0$ for all $l \in \mathbb{N}_0$.*

755 *Proof.* The fact that k_d is a kernel follows directly from the inclusion $\Phi_{d+2} \subseteq \Phi_d$ mentioned in
756 Gneiting (2013). For $D \in \{d, d+2\}$, let $\mu_{l,d}$ be the sequence of eigenvalues μ_l associated with k_D .
757 Then, as mentioned for example by Hubbert et al. (2022), the Schoenberg coefficients $b_{l,d}$ satisfy

$$b_{l,d} = \frac{\Gamma\left(\frac{d+1}{2}\right) N_{m,d} \mu_{l,d}}{2\pi^{(d+1)/2}}.$$

758 Especially, the Schoenberg coefficients $b_{l,d}$ have the same sign as the eigenvalues $\mu_{l,d}$. We use

$$0 \leq b_{l,d+2} = \begin{cases} b_{l,d} - \frac{1}{2}b_{l+2,d} & , l = 0 \text{ and } d = 1 \\ \frac{1}{2}(l+1)(b_{l,d} - b_{l+2,d}) & , l \geq 1 \text{ and } d = 1 \\ \frac{(l+d-1)(l+d)}{d(2l+d-1)}b_{l,d} - \frac{(l+1)(l+2)}{d(2l+d+3)}b_{l+2,d} & , d \geq 2, \end{cases}$$

759 where the inequality follows from the fact that k_{d+2} is a kernel and the equality is the statement of
 760 Corollary 3 by Gneiting (2013). In any of the three cases, $b_{l+2,d} > 0$ implies $b_{l,d} > 0$. Hence, if
 761 $b_{l,d} > 0$ for infinitely many l , then $b_{l,d} > 0$ for all l , which implies $\mu_{l,d} > 0$ for all l . \square

762 C Gradient flow and gradient descent with kernels

763 C.1 Derivation of gradient flow and gradient descent

Here, we derive expressions for gradient flow and gradient descent in the RKHS for the regularized loss

$$L(f) := \frac{1}{n} \sum_{i=1}^n (y_i - f(\mathbf{x}_i))^2 + \rho \|f\|_{\mathcal{H}_k}^2 = \frac{1}{n} \sum_{i=1}^n (y_i - \langle k(\mathbf{x}_i, \cdot), f \rangle_{\mathcal{H}_k})^2 + \rho \langle f, f \rangle_{\mathcal{H}_k}.$$

764 Note that we will take derivatives in the RKHS with respect to f , which is different from taking
 765 derivatives w.r.t. the coefficients \mathbf{c} in a model $f(\mathbf{x}) = \mathbf{c}^\top k(\mathbf{X}, \mathbf{x})$.

In the RKHS-Norm, the Fréchet derivative of L is

$$\frac{\partial L(f)}{f} = \frac{2}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i) \langle k(\mathbf{x}_i, \cdot), \cdot \rangle_{\mathcal{H}_k} + 2\rho \langle f, \cdot \rangle_{\mathcal{H}_k},$$

which is represented in \mathcal{H}_k by

$$L'(f) = \frac{2}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i) k(\mathbf{x}_i, \cdot) + 2\rho f.$$

Now assume that $f = \sum_{i=1}^n a_i k(\mathbf{x}_i, \cdot) = \mathbf{a}^\top k(\mathbf{X}, \cdot)$. Then,

$$\begin{aligned} L'(f) &= \frac{2}{n} \sum_{i=1}^n (\mathbf{a}^\top k(\mathbf{X}, \mathbf{x}_i) - y_i) k(\mathbf{x}_i, \cdot) + 2\rho \mathbf{a}^\top k(\mathbf{X}, \cdot) \\ &= \frac{2}{n} (\mathbf{a}^\top k(\mathbf{X}, \mathbf{X}) k(\mathbf{X}, \cdot) - \mathbf{y}^\top k(\mathbf{X}, \cdot) + \rho n \mathbf{a}^\top k(\mathbf{X}, \cdot)) \\ &= \frac{2}{n} ((k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n) \mathbf{a} - \mathbf{y})^\top k(\mathbf{X}, \cdot). \end{aligned}$$

Especially, under gradient flow of f , the coefficients \mathbf{a} follow the dynamics

$$\dot{\mathbf{a}}(t) = \frac{2}{n} (\mathbf{y} - (k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n) \mathbf{a}(t)),$$

which is solved for $\mathbf{a}(0) = \mathbf{0}$ by

$$\mathbf{a}(t) = \left(\mathbf{I}_n - e^{-\frac{2}{n} t (k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n)} \right) (k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n)^{-1} \mathbf{y},$$

766 which is the closed form expression (1) of $f_{t,\rho}$.

For gradient descent, assuming that $f_{t,\rho}^{\text{GD}} = \mathbf{c}_{t,\rho}^\top k(\mathbf{X}, \cdot)$, we have

$$\begin{aligned} f_{t+1,\rho}^{\text{GD}} &= f_{t,\rho}^{\text{GD}} - \eta_t L'(f_{t,\rho}^{\text{GD}}) = \mathbf{c}_{t,\rho}^\top k(\mathbf{X}, \cdot) - \eta_t \frac{2}{n} ((k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n) \mathbf{c}_{t,\rho} - \mathbf{y})^\top k(\mathbf{X}, \cdot) \\ &= \left(\mathbf{c}_{t,\rho} + \eta_t \frac{2}{n} (\mathbf{y} - (k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n) \mathbf{c}_{t,\rho}) \right)^\top k(\mathbf{X}, \cdot) \end{aligned}$$

If $f_{0,\rho}^{\text{GD}} \equiv 0$, the coefficients evolve as $\mathbf{c}_0 = \mathbf{0}$ and

$$\mathbf{c}_{t+1,\rho} = \mathbf{c}_{t,\rho} + \eta_t \frac{2}{n} (\mathbf{y} - (k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n) \mathbf{c}_{t,\rho}).$$

767 For an analysis of gradient descent for kernel regression with $\rho = 0$, we refer to, e.g., Yao et al.
 768 (2007).

769 **C.2 Gradient flow and gradient descent initialized at 0 have monotonically growing \mathcal{H} -norm**

770 In the following proposition we show that under gradient flow and gradient descent with suffi-
 771 ciently small learning rates initialized at 0, the RKHS norm grows monotonically with time t . This
 772 immediately implies that Assumption (N) with $C_{\text{norm}} = 1$ holds for all estimators $f_{t,\rho}$ from (1).

773 **Proposition C.1.**

- 774 (i) For any $t \in [0, \infty]$ and any $\rho \geq 0$, $f_{t,\rho}$ from (1) fulfills Assumption (N) with $C_{\text{norm}} = 1$.
 775 (ii) For any $t \in \mathbb{N}_0 \cup \{\infty\}$ and any $\rho \geq 0$, with sufficiently small fixed learning rate $0 \leq \eta \leq$
 776 $\frac{1}{2(\rho + \lambda_{\max}(k(\mathbf{X}, \mathbf{X}))/n)}$, $f_{t,\rho}^{\text{GD}}$ fulfills Assumption (N) with $C_{\text{norm}} = 1$.

777 *Proof. Proof of (i):*

We write $f_{t,\rho}(\mathbf{x}) = k(\mathbf{x}, \mathbf{X})\mathbf{c}_{t,\rho}$, where $\mathbf{c}_{t,\rho} := A_{t,\rho}(\mathbf{X})\mathbf{y}$. We now show that the RKHS-norm of $f_{t,\rho}$ grows monotonically in t , by using the eigendecomposition $k(\mathbf{X}, \mathbf{X}) = \mathbf{U}\Lambda\mathbf{U}^\top$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^{n \times n}$ is diagonal and $\mathbf{U} \in \mathbb{R}^{n \times n}$ is orthonormal, and writing $\tilde{\mathbf{y}} := \mathbf{U}^\top \mathbf{y}$. Then it holds that

$$\begin{aligned} \|f_{t,\rho}\|_{\mathcal{H}}^2 &= (\mathbf{c}_{t,\rho})^\top k(\mathbf{X}, \mathbf{X})\mathbf{c}_{t,\rho} = \tilde{\mathbf{y}}^\top (\Lambda + \rho n \mathbf{I}_n)^{-1} \left(\mathbf{I}_n - \exp\left(-\frac{2t}{n}(\Lambda + \rho n \mathbf{I}_n)\right) \right) \Lambda \\ &\quad \left(\mathbf{I}_n - \exp\left(-\frac{2t}{n}(\Lambda + \rho n \mathbf{I}_n)\right) \right) (\Lambda + \rho n \mathbf{I}_n)^{-1} \tilde{\mathbf{y}} \\ &= \sum_{k=1, \lambda_k + \rho n > 0}^n \tilde{y}_k^2 \underbrace{\frac{\lambda_k}{(\lambda_k + \rho n)^2}}_{\leq 1/\lambda_k} \underbrace{\left(1 - \exp\left(-\frac{2t}{n}(\lambda_k + \rho n)\right)\right)}_{\leq 1} \\ &\leq \sum_{k=1, \lambda_k > 0}^n \tilde{y}_k^2 \frac{1}{\lambda_k} = \|f_{\infty,0}\|_{\mathcal{H}}^2. \end{aligned}$$

778 **Proof of (ii):**

779 Expanding the iteration in the definition of $\mathbf{c}_{t,\rho}$ yields

$$\mathbf{c}_{t+1,\rho} = \sum_{i=0}^t \prod_{j=0}^{t-i-1} \left(\mathbf{I} - \frac{2\eta_{t-j}}{n} (k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}) \right) \frac{2\eta_i}{n} \mathbf{y}.$$

We again use the eigendecomposition $k(\mathbf{X}, \mathbf{X}) = \mathbf{U}\Lambda\mathbf{U}^\top$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^{n \times n}$ is diagonal and $\mathbf{U} \in \mathbb{R}^{n \times n}$ is orthonormal, and write $\tilde{\mathbf{y}} := \mathbf{U}^\top \mathbf{y}$. Then, using sufficiently small learning rates $0 \leq \eta_t \leq \frac{1}{2(\rho + \lambda_{\max}(k(\mathbf{X}, \mathbf{X}))/n)}$ in all time steps $t \in \mathbb{N}$, it holds that

$$\begin{aligned} &\|f_{t,\rho}^{\text{GD}}\|_{\mathcal{H}}^2 \\ &= (\mathbf{c}_{t,\rho})^\top k(\mathbf{X}, \mathbf{X})\mathbf{c}_{t,\rho} \\ &= \tilde{\mathbf{y}}^\top \left(\sum_{i=0}^t \frac{2\eta_i}{n} \prod_{j=0}^{t-i-1} \left((1 - 2\eta_{t-j}\rho) \mathbf{I} - \frac{2\eta_{t-j}}{n} \Lambda \right) \right) \Lambda \left(\sum_{i=0}^t \frac{2\eta_i}{n} \prod_{j=0}^{t-i-1} \left((1 - 2\eta_{t-j}\rho) \mathbf{I} - \frac{2\eta_{t-j}}{n} \Lambda \right) \right) \tilde{\mathbf{y}} \\ &= \sum_{k=1}^n \underbrace{\tilde{y}_k^2 \lambda_k}_{\geq 0} \left(\sum_{i=0}^t \frac{2\eta_i}{n} \prod_{j=0}^{t-i-1} \underbrace{\left(1 - 2\eta_{t-j}(\rho + \lambda_k/n)\right)}_{\in [0,1]} \right)^2. \end{aligned} \tag{C.1}$$

The last display shows that $\|f_{t,\rho}^{\text{GD}}\|_{\mathcal{H}}^2$ grows monotonically in t , strictly monotonically if $\eta_t \in (0, \frac{1}{2(\rho + \lambda_{\max}(k(\mathbf{X}, \mathbf{X}))/n)})$ holds for all t . It also shows that if $\rho' \geq \rho$ then $\|f_{t,\rho'}^{\text{GD}}\|_{\mathcal{H}} \leq \|f_{t,\rho}^{\text{GD}}\|_{\mathcal{H}}$ for any $t \in \mathbb{N} \cup \{\infty\}$. To see that $\|f_{t,\rho}^{\text{GD}}\|_{\mathcal{H}}^2 \leq \|f_{\infty,0}\|_{\mathcal{H}}^2$ for all $t \in \mathbb{N} \cup \{\infty\}$ and all $\rho \geq 0$, observe that with fixed learning rates $\eta_t = \eta \in (0, \frac{1}{2(\rho + \lambda_{\max}(k(\mathbf{X}, \mathbf{X}))/n)}) \subseteq (0, \frac{1}{2\lambda_{\max}(k(\mathbf{X}, \mathbf{X}))/n})$, for all $t \in \mathbb{N} \cup \{\infty\}$ it holds that

$$\sum_{i=0}^t \frac{2\eta_i}{n} \prod_{j=0}^{t-i-1} (1 - 2\eta_{t-j} \lambda_k/n) = \frac{2\eta}{n} \sum_{i=0}^t (1 - 2\eta \lambda_k/n)^{t-i}$$

$$= \frac{2\eta}{n} \sum_{i=0}^t (1 - 2\eta\lambda_k/n)^i = \frac{2\eta}{n} \frac{1 - (1 - 2\eta\lambda_k/n)^{t+1}}{2\eta\lambda_k/n} \leq \frac{1}{\lambda_k}.$$

780 Since it suffices to consider the case $\rho \rightarrow 0$, using the above derivation in (C.1) yields $\|f_{t,\rho}^{\text{GD}}\|_{\mathcal{H}}^2 \leq$
781 $\|f_{\infty,0}\|_{\mathcal{H}}^2$ for all $t \in \mathbb{N}$, which concludes the proof. \square

782 D Proof of Theorem 1

783 Our goal in this section is to prove Theorem D.1, which can be seen as a generalization of Theorem 1
784 to varying bandwidths. To be able to speak of bandwidths, we need to consider translation-invariant
785 kernels. Although Theorem 1 is formulated for general kernels with Sobolev RKHS, it follows from
786 Theorem D.1 since we can always find, for a fixed bandwidth, a translation-invariant kernel with
787 equivalent RKHS, such that only the constant C_{norm} changes in the theorem statement.

To generate the RKHS H^s , Buchholz (2022) uses the translation-invariant kernel $k^B(\mathbf{x}, \mathbf{y}) =$
 $u^B(\mathbf{x} - \mathbf{y})$ defined via its Fourier transform $\hat{u}^B(\boldsymbol{\xi}) = (1 + |\boldsymbol{\xi}|^2)^{-s}$. Adapting the bandwidth, the
kernel is then normalized in the usual L_1 -sense,

$$k_\gamma^B(\mathbf{x}, \mathbf{y}) = \gamma^{-d} u^B((\mathbf{x} - \mathbf{y})/\gamma). \quad (\text{D.1})$$

788 **Theorem D.1 (Inconsistency of overfitting estimators).** *Let assumptions (D1) and (K) hold. Let*
789 $c_{\text{fit}} \in (0, 1]$ *and* $C_{\text{norm}} > 0$. *Then, there exist* $c > 0$ *and* $n_0 \in \mathbb{N}$ *such that the following holds for*
790 *all* $n \geq n_0$ *with probability* $1 - O(1/n)$ *over the draw of the data set* D *with* n *samples: For every*
791 *function* $f \in \mathcal{H}_k$ *with*

- 792 (O) $\frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2 \leq (1 - c_{\text{fit}}) \cdot \sigma^2$ (training error below Bayes risk) and
793 (N) $\|f\|_{\mathcal{H}_k} \leq C_{\text{norm}} \|f_{\infty,0}\|_{\mathcal{H}_k}$ (norm comparable to minimum-norm interpolant (1)),

the excess risk satisfies

$$R_P(f) - R_P(f^*) \geq c > 0. \quad (\text{D.2})$$

794 If k_γ denotes a L_1 -normalized translation-invariant kernel with bandwidth $\gamma > 0$, i.e. there exists a
795 $q : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $k_\gamma(x, y) = \gamma^{-d} q(\frac{x-y}{\gamma})$, then inequality (D.2) holds with c independent of the
796 sequence of bandwidths $(\gamma_n)_{n \in \mathbb{N}} \subseteq (0, 1)$, as long as f_D fulfills (N) for the sequence $(\mathcal{H}_{\gamma_n})_{n \in \mathbb{N}}$ with
797 constant $C_{\text{norm}} > 0$.

Proof. By assumption, the RKHS norm $\|\cdot\|_{\mathcal{H}_k}$ induced by the kernel k (or k_γ if we allow bandwidth
adaptation) is equivalent to the RKHS norm $\|\cdot\|_{\mathcal{H}_\gamma}$ induced by a kernel of the form (D.1) with an
arbitrary but fixed choice of bandwidth $\gamma \in (0, 1)$, which means that there exists a constant $C_\gamma > 0$
such that $\frac{1}{C_\gamma} \|f\|_{\mathcal{H}_\gamma} \leq \|f\|_{\mathcal{H}_k} \leq C_\gamma \|f\|_{\mathcal{H}_\gamma}$ for all $f \in \mathcal{H}_k$. Hence the minimum-norm interpolant
 $g_{D,\gamma}$ in \mathcal{H}_γ satisfies

$$\|f_D\|_{\mathcal{H}_\gamma} \leq C_\gamma \|f_D\|_{\mathcal{H}_k} \leq C_\gamma C_{\text{norm}} \|g_D\|_{\mathcal{H}_k} \leq C_\gamma C_{\text{norm}} \|g_{D,\gamma}\|_{\mathcal{H}_k} \leq C_\gamma^2 C_{\text{norm}} \|g_{D,\gamma}\|_{\mathcal{H}_\gamma},$$

798 where $\|g_D\|_{\mathcal{H}_k} \leq \|g_{D,\gamma}\|_{\mathcal{H}_k}$ because, g_D is the minimum-norm interpolant in \mathcal{H}_k .

799 Now consider the RKHS norm $\|\cdot\|_{\tilde{\mathcal{H}}_\gamma}$ of a translation-invariant kernel k_γ . Then the functions
800 $\{h_p(x) = e^{ip \cdot x}\}_{p \in \mathbb{R}^d}$ are eigenfunctions of the kernel's integral operator, so that the RKHS norm
801 can be written as (Rakhlin and Zhai, 2019)

$$\|f\|_{\tilde{\mathcal{H}}_\gamma}^2 = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \frac{|\hat{f}(\omega)|^2}{\hat{q}(\omega)} d\omega,$$

802 where \hat{f} denotes the Fourier transform of f .

By assumption we know that there exists a $C_{\gamma_0} > 0$ such that $\frac{1}{C_{\gamma_0}} \|f\|_{\mathcal{H}_{\gamma_0}} \leq \|f\|_{\tilde{\mathcal{H}}_{\gamma_0}} \leq C_{\gamma_0} \|f\|_{\mathcal{H}_{\gamma_0}}$
holds for some fixed bandwidth $\gamma_0 > 0$, then substituting by $\tilde{\omega} = \frac{\gamma}{\gamma_0} \omega$ yields

$$\|f\|_{\tilde{\mathcal{H}}_\gamma} = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \frac{|\hat{f}(\omega)|^2}{\hat{q}_1(\gamma\omega)} d\omega = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \frac{|\hat{f}(\frac{\gamma_0}{\gamma} \tilde{\omega})|^2}{\hat{q}_1(\gamma_0 \tilde{\omega})} \left(\frac{\gamma_0}{\gamma}\right)^d d\tilde{\omega} = \|\tilde{f}\|_{\tilde{\mathcal{H}}_{\gamma_0}}$$

$$\leq C_{\gamma_0} \|\tilde{f}\|_{\mathcal{H}_{\gamma_0}} = \frac{C_{\gamma_0}}{(2\pi)^d} \int_{\mathbb{R}^d} \frac{|\hat{f}(\frac{\gamma_0}{\gamma} \tilde{\omega})|^2}{\hat{q}_2(\gamma_0 \tilde{\omega})} \left(\frac{\gamma_0}{\gamma}\right)^d d\tilde{\omega} = \frac{C_{\gamma_0}}{(2\pi)^d} \int_{\mathbb{R}^d} \frac{|\hat{f}(\omega)|^2}{\hat{q}_2(\gamma\omega)} d\omega = C_{\gamma_0} \|f\|_{\mathcal{H}_{\gamma}}.$$

803 In the same way we get $\frac{1}{C_{\gamma_0}} \|f\|_{\mathcal{H}_{\gamma}} \leq \|f\|_{\tilde{\mathcal{H}}_{\gamma}}$ for arbitrary $\gamma \in (0, 1)$. This shows that the constant
 804 C_{γ_0} , that quantifies the equivalence between $\|\cdot\|_{\mathcal{H}_{\gamma}}$ and $\|\cdot\|_{\tilde{\mathcal{H}}_{\gamma}}$ does not depend on the bandwidth γ .
 805 Finally Proposition D.4, Proposition D.2 and Remark D.3 together yield the result. \square

806 The following proposition generalizes the inconsistency result for large bandwidths, Proposition 4
 807 in Buchholz (2022), beyond interpolating estimators to estimators that overfit at least an arbitrary
 808 constant fraction beyond the Bayes risk and whose RKHS norm is at most a constant factor larger than
 809 the RKHS norm of the minimum-norm interpolant. Compared to Rakhlin and Zhai (2019), Buchholz
 810 gets a statement in probability over the draw of a training set D and less restrictive assumptions on
 811 the domain Ω and dimension d .

812 **Proposition D.2 (Inconsistency for large bandwidths).** *Let $c_{\text{fit}} \in (0, 1]$ and $C_{\text{norm}} > 0$. Let the*
 813 *data set $D = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ be drawn i.i.d. from a distribution P that fulfills Assumption*
 814 *(D1), let $g_{D, \gamma}$ be the minimum-norm interpolant in $\mathcal{H} := \mathcal{H}_{\gamma}$ with respect to the kernel (D.1) for a*
 815 *bandwidth $\gamma > 0$. Then, for every $A > 0$, there exist $c > 0$ and $n_0 \in \mathbb{N}$ such that the following holds*
 816 *for all $n \geq n_0$ with probability $1 - O(1/n)$ over the draw of the data set D with n samples:*

For every function $f \in \mathcal{H}$ that fulfills Assumption (O) with c_{fit} and Assumption (N) with C_{norm} the excess risk satisfies

$$\mathbb{E}_{\mathbf{x}}(f(\mathbf{x}) - f^*(\mathbf{x}))^2 \geq c > 0,$$

817 *where c depends neither on n nor on $1 > \gamma > An^{-1/d} > 0$.*

818 **Remark D.3.** Proposition D.2 holds for any kernel that fulfills Assumption (K). The reason is that
 819 any kernel k that fulfills assumption (K) and the kernel defined in (D.1) have the same RKHS and
 820 equivalent norms. Therefore every function $f \in \mathcal{H}_k = \mathcal{H}_{\gamma}$ (equality as sets) that fulfills Assumptions
 821 (O) and (N) for the kernel k also fulfills Assumptions (O) and (N) with an adapted constant C_{norm}
 822 for the kernel (D.1). \blacktriangleleft

823 **Proof. Step 1: Generalizing the procedure in Buchholz (2022).**

824 We write $[n] = \{1, \dots, n\}$ and follow the proof of Proposition 4 in Buchholz (2022). Define
 825 $u(\mathbf{x}) = f(\mathbf{x}) - f^*(\mathbf{x})$. We need to show that with probability at least $1 - O(n^{-1})$ over the draw of
 826 D it holds that $\|u\|_{L^2(P_X)} \geq c > 0$, where c depends neither on n nor on γ .

For this purpose we show that with probability at least $1 - 3n^{-1}$ over the draw of D there exist a
 constants $c'', \kappa'' > 0$ depending only on c_{fit} and a subset $\mathcal{P}'' \subseteq [n]$ with $|\mathcal{P}''| \geq \lfloor \kappa'' \cdot n \rfloor$ such that

$$|f(\mathbf{x}_i) - f^*(\mathbf{x}_i)| \geq c'' > 0 \text{ holds for all } i \in \mathcal{P}'' \tag{D.3}$$

827 Then via Lemma 7 in Buchholz (2022) as well as Lemma D.7 we can choose a large subset $\mathcal{P}''' \subseteq [n]$
 828 of the training point indices with $|\mathcal{P}'''| \geq n - |\mathcal{P}''|/2$, such that the \mathbf{x}_i for $i \in \mathcal{P}'''$ are well-separated
 829 in the sense that $\min_{\{i, j \in \mathcal{P}''', i \neq j\}} \|\mathbf{x}_i - \mathbf{x}_j\| \geq d_{\text{min}}$ with $d_{\text{min}} := c''' n^{-1/d}$, where c''' depends on
 830 c_{fit}, d , the upper bound on the Lebesgue density C_u and on the smoothness of the RKHS s . Then the
 831 intersection $\mathcal{P}'' \cap \mathcal{P}'''$ contains at least $\frac{|\mathcal{P}''|}{2}$ points. Now we can replace \mathcal{P}' in the proof of Proposition
 832 4 for $s \in \mathbb{N}$ in Buchholz (2022) by the intersection $\mathcal{P}'' \cap \mathcal{P}'''$. The rest of the proof applies without
 833 modification, where (42) holds by our assumption $\|f\|_{\mathcal{H}} \leq C_{\mathcal{H}} \|g_D\|_{\mathcal{H}}$. Our modifications do not
 834 affect Buchholz' arguments for the extension to $s \notin \mathbb{N}$.

835 **Step 2: The existence of \mathcal{P}'' .**

Given a choice of $\kappa'', c'' > 0$, consider the event (over the draw of D)

$$\begin{aligned} E &:= \{\exists \mathcal{P}'' \subseteq [n] \text{ with } |\mathcal{P}''| \geq \lfloor \kappa'' \cdot n \rfloor \text{ that fulfills (D.3)}\} \\ &= \{\exists \tilde{\mathcal{P}} \subseteq [n] \text{ with } |\tilde{\mathcal{P}}| \geq \lceil (1 - \kappa'')n \rceil \text{ such that } |f^*(\mathbf{x}_i) - f(\mathbf{x}_i)| < c'' \quad \forall i \in \tilde{\mathcal{P}}\}. \end{aligned}$$

836 With the proper choices of c'' and κ'' independent of n and f , we will show $P(E) \leq 3n^{-1}$. We
 837 will find a small $c'' > 0$ such that if f^* and f are closer than c'' on too many training points $\tilde{\mathcal{P}}$

838 and f overfits by at least the fraction c_{fit} , the noise variables ε_i on the complement $\tilde{\mathcal{P}}^c$ would have
 839 to be unreasonably large, contradicting the event E_{6i} defined below, and implying (D.3) with high
 840 probability. We will use the notation $\|\mathbf{f}\|_{\tilde{\mathcal{P}}}^2 := \sum_{i \in \tilde{\mathcal{P}}} f(\mathbf{x}_i)^2$ and $\|\mathbf{y}\|_{\tilde{\mathcal{P}}}^2 := \sum_{i \in \tilde{\mathcal{P}}} y_i^2$.

841 **Step 2b: Noise bounds.**

842 Lemma D.6 (i) states that there exists a $\kappa'' > 0$ small enough such that the event (over the draw of
 843 D)

$$E_{6i} := \{\forall \mathcal{P}_1 \subseteq [n] \text{ with } |\mathcal{P}_1| \leq \lfloor \kappa'' \cdot n \rfloor \text{ it holds that } \frac{1}{n} \|\mathbf{f}^* - \mathbf{y}\|_{\mathcal{P}_1}^2 = \frac{1}{n} \sum_{i \in \mathcal{P}_1} \varepsilon_i^2 < \frac{c_{\text{fit}}}{4} \sigma^2\},$$

844 fulfills, for n large enough, $P(E_{6i}) \geq 1 - n^{-1}$.

845 Lemma D.6 (ii) implies that there exists a $c_{\text{lower}} > 0$ such that the event (over the draw of D)

$$E_{6ii} := \{\forall \mathcal{P}_2 \text{ with } |\mathcal{P}_2| \geq \lfloor (1 - \kappa'')n \rfloor \text{ it holds that } \frac{1}{n} \|\mathbf{f}^* - \mathbf{y}\|_{\mathcal{P}_2}^2 \geq c_{\text{lower}} \cdot \sigma^2\},$$

846 fulfills, for n large enough, $P(E_{6ii}) \geq 1 - n^{-1}$.

847 Lemma D.5 states that the total amount of noise $\|\varepsilon\|_{[n]}^2$ concentrates around its mean $n\sigma^2$. More
 848 precisely, we will use that for any $c_\varepsilon \in (0, 1)$ the event (over the draw of D)

$$E_5 := \left\{ \frac{1}{n} \|\mathbf{f}^* - \mathbf{y}\|_{[n]}^2 \geq c_\varepsilon \cdot \sigma^2 \right\},$$

849 fulfills $P(E_5) \geq 1 - \exp\left(-n \cdot \left(\frac{1-c_\varepsilon}{2}\right)^2\right)$.

850 **Step 2c: Lower bounding $\|\varepsilon\|_{\tilde{\mathcal{P}}^c}^2$.**

851 Given some function $f \in \mathcal{H}$, assume in steps 2c and 2d that event E holds and that $\tilde{\mathcal{P}} \subseteq [n]$ denotes
 852 a subset of the training set that fulfills $|\tilde{\mathcal{P}}| \geq \lceil (1 - \kappa'')n \rceil$ and $|f^*(\mathbf{x}_i) - f(\mathbf{x}_i)| < c'' \quad \forall i \in \tilde{\mathcal{P}}$.

853 In step 2c, assume we choose $\tilde{c}_{\text{fit}} > 0$ such that $\tilde{c}_{\text{fit}} \|\mathbf{f}^* - \mathbf{y}\|_{\tilde{\mathcal{P}}}^2 \leq \|\mathbf{f} - \mathbf{y}\|_{\tilde{\mathcal{P}}}^2$. Then by the overfitting
 854 Assumption (O) it holds that

$$\frac{1}{n} (\tilde{c}_{\text{fit}} \|\mathbf{f}^* - \mathbf{y}\|_{\tilde{\mathcal{P}}}^2 + \|\mathbf{f} - \mathbf{y}\|_{\tilde{\mathcal{P}}^c}^2) \leq \frac{1}{n} (\|\mathbf{f} - \mathbf{y}\|_{\tilde{\mathcal{P}}}^2 + \|\mathbf{f} - \mathbf{y}\|_{\tilde{\mathcal{P}}^c}^2) \leq (1 - c_{\text{fit}}) \sigma^2. \quad (\text{D.4})$$

If we restrict ourselves to event E_5 , dropping the term $\|\mathbf{f} - \mathbf{y}\|_{\tilde{\mathcal{P}}^c}^2$ in (D.4), then dividing by \tilde{c}_{fit} and
 subtracting the result from the inequality in the definition of event E_5 yields

$$\frac{1}{n} \|\varepsilon\|_{\tilde{\mathcal{P}}^c}^2 = \frac{1}{n} \|\mathbf{f}^* - \mathbf{y}\|_{\tilde{\mathcal{P}}^c}^2 \geq c_\varepsilon \sigma^2 - \frac{1 - c_{\text{fit}}}{\tilde{c}_{\text{fit}}} \sigma^2. \quad (\text{D.5})$$

855 **Step 2d: Choosing the constants.**

If we choose $c_\varepsilon := 1 - \frac{c_{\text{fit}}}{4}$ and $\tilde{c}_{\text{fit}} := \frac{2-2c_{\text{fit}}}{2-c_{\text{fit}}} \in (0, 1)$, then (D.5) becomes

$$\frac{1}{n} \|\varepsilon\|_{\tilde{\mathcal{P}}^c}^2 \geq \frac{c_{\text{fit}}}{4} \sigma^2.$$

856 Now it is left to show that the condition $\tilde{c}_{\text{fit}} \|\mathbf{f}^* - \mathbf{y}\|_{\tilde{\mathcal{P}}}^2 \leq \|\mathbf{f} - \mathbf{y}\|_{\tilde{\mathcal{P}}}^2$, that is required for Step 2c,
 857 holds with high probability with our choice of \tilde{c}_{fit} .

With some arbitrary but fixed $\varepsilon_{\text{lower}} \in (0, \sqrt{c_{\text{lower}}})$, choose $c'' := (1 - \sqrt{\tilde{c}_{\text{fit}}}) \left(\sqrt{\frac{c_{\text{lower}}}{1-\kappa''}} - \frac{\varepsilon_{\text{lower}}}{\sqrt{1-\kappa''}} \right) \sigma$.
 Then on event E_{6ii} , for n large enough, it holds that

$$(1 - \sqrt{\tilde{c}_{\text{fit}}}) \frac{1}{\sqrt{n}} \|\mathbf{f}^* - \mathbf{y}\|_{\tilde{\mathcal{P}}} \geq (1 - \sqrt{\tilde{c}_{\text{fit}}}) \sqrt{c_{\text{lower}}} \sigma \geq \sqrt{1 - \kappa''} \cdot c'' + \frac{c''}{\sqrt{n}}. \quad (\text{D.6})$$

858 By definition of $\tilde{\mathcal{P}}$, it holds that

$$\|\mathbf{f} - \mathbf{f}^*\|_{\tilde{\mathcal{P}}}^2 = \sum_{i \in \tilde{\mathcal{P}}} (f(\mathbf{x}_i) - f^*(\mathbf{x}_i))^2 < \lceil (1 - \kappa'')n \rceil (c'')^2,$$

so that

$$\frac{1}{\sqrt{n}} \|\mathbf{f} - \mathbf{f}^*\|_{\tilde{\mathcal{P}}} < \sqrt{1 - \kappa''} \cdot c'' + \frac{c''}{\sqrt{n}}. \quad (\text{D.7})$$

Now, using the triangle inequality, (D.7) and (D.6) yields the condition required for Step 2c,

$$\begin{aligned} & \frac{1}{\sqrt{n}} \|\mathbf{f} - \mathbf{y}\|_{\tilde{\mathcal{P}}} \\ & \geq \frac{1}{\sqrt{n}} \|\mathbf{f}^* - \mathbf{y}\|_{\tilde{\mathcal{P}}} - \frac{1}{\sqrt{n}} \|\mathbf{f} - \mathbf{f}^*\|_{\tilde{\mathcal{P}}} \\ & \geq \frac{1}{\sqrt{n}} \|\mathbf{f}^* - \mathbf{y}\|_{\tilde{\mathcal{P}}} - \sqrt{1 - \kappa''} \cdot c'' - \frac{c''}{\sqrt{n}} \\ & \geq \sqrt{\tilde{c}_{\text{fit}}} \frac{1}{\sqrt{n}} \|\mathbf{f}^* - \mathbf{y}\|_{\tilde{\mathcal{P}}}. \end{aligned}$$

859 **Step 2e: Upper bounding the probability of E .**

To conclude, we have seen in steps 2c and 2d that on $E \cap E_{6ii} \cap E_5$, it holds that

$$\frac{1}{n} \|\boldsymbol{\varepsilon}\|_{\tilde{\mathcal{P}}^c}^2 \geq \frac{c_{\text{fit}}}{4} \sigma^2.$$

860 On E_{6i} , it holds that

$$\frac{1}{n} \|\boldsymbol{\varepsilon}\|_{\tilde{\mathcal{P}}^c}^2 < \frac{c_{\text{fit}}}{4} \sigma^2.$$

Hence $E_{6i} \cap E \cap E_{6ii} \cap E_5 = \emptyset$. This implies $E \subseteq (E_5 \cap E_{6i} \cap E_{6ii})^c$, where the right hand side is independent of $f \in \mathcal{H}$ and just depends on the training data D . Since $P(E_{6i}) \geq 1 - n^{-1}$ and $P(E_{6ii} \cap E_5) \geq 1 - n^{-1} - \exp\left(-n \cdot \left(\frac{1-c_\varepsilon}{2}\right)^2\right)$, it must hold that, for n large enough,

$$P(E) \leq P((E_5 \cap E_{6i} \cap E_{6ii})^c) \leq 2n^{-1} + \exp\left(-n \cdot \left(\frac{1-c_\varepsilon}{2}\right)^2\right) \leq 3n^{-1}. \quad \square$$

861 The following proposition generalizes the inconsistency result for small bandwidths, Proposition 5
862 in Buchholz (2022), beyond interpolating estimators to estimators whose RKHS norm is at most a
863 constant factor larger than the RKHS norm of the minimum-norm interpolant. The intuition is that
864 if the bandwidth is too small, then the minimum-norm interpolant $g_{D,\gamma}$ returns to 0 between the
865 training points. Then $\|g_{D,\gamma}\|_{L_2(\rho)}$ is smaller and bounded away from $\|f^*\|_{L_2(\rho)}$. We can replace
866 $g_{D,\gamma}$ by any other function $f \in \mathcal{H}$ that fulfills Assumption (N).

Proposition D.4 (Inconsistency for small bandwidths). *Under the assumptions of Proposition D.2, there exist constants $B, c > 0$ such that, with probability $1 - O(n^{-1})$ over the draw of D : For any function $f \in \mathcal{H}$ that fulfills Assumption (N) but not necessarily Assumption (O), the excess risk satisfies*

$$\mathbb{E}_{\mathbf{x}}(f(\mathbf{x}) - f^*(\mathbf{x}))^2 \geq c > 0,$$

867 where c depends neither on n nor on $\gamma < Bn^{-1/d}$.

Proof. Denote the upper bound on the Lebesgue density of P_X by C_u . The triangle inequality implies

$$\begin{aligned} \|f^* - f\|_{L_2(P_X)} & \geq \|f^*\|_{L_2(P_X)} - \|f\|_{L_2(P_X)} \geq \|f^*\|_{L_2(P_X)} - \sqrt{C_u} \|f\|_2 \\ & \geq \|f^*\|_{L_2(P_X)} - \sqrt{C_u} \|f\|_{\mathcal{H}} \geq \|f^*\|_{L_2(P_X)} - C_{\mathcal{H}} \sqrt{C_u} \|g_{D,\gamma}\|_{\mathcal{H}}, \end{aligned}$$

868 where $\|f\|_2 \leq \|f\|_{\mathcal{H}}$ follows from the fact that the Fourier transform \hat{k} of the kernel satisfies $\hat{k}(\xi) \leq 1$.
869 Now in the proof of Lemma 17 in Buchholz (2022) $a > 0$ can be chosen smaller to generalize the
870 statement to

$$\|g_{D,\gamma}\|_{\mathcal{H}}^2 \leq \frac{1}{6C_{\mathcal{H}}^2 C_u} \|f^*\|_{L_2(P_X)}^2 + c_9(\gamma^2 n^{2/d} + \gamma^{2s} n^{2s/d}),$$

where c_9 depends on c_u, f^*, d, s and C_{norm} . Finally we can choose B small enough such that Eq. (32) in Buchholz (2022) can be replaced by $C_{\mathcal{H}} \sqrt{C_u} \|g_{D,\gamma}\|_{\mathcal{H}} \leq \frac{2}{3} \|f^*\|_{L_2(P_X)}$ so that we get

$$\|f^* - f\|_{L_2(P_X)} \geq \frac{1}{3} \|f^*\|_{L_2(P_X)} > 0. \quad \square$$

871 **D.1 Auxiliary results for the proof of Theorem 1**

872 **Lemma D.5** (Concentration of χ_n^2 variables). *Let U be a chi-squared distributed random variable*
 873 *with n degrees of freedom. Then, for any $c \in (0, 1)$ it holds that*

$$P\left(\frac{U}{n} \leq c\right) \leq \exp\left(-n \cdot \left(\frac{1-c}{2}\right)^2\right).$$

874 *Proof.* Lemma 1 in Laurent and Massart (2000) implies for any $x > 0$,

$$P\left(\frac{U}{n} \leq 1 - 2\sqrt{\frac{x}{n}}\right) \leq \exp(-x).$$

875 Solving $c = 1 - 2\sqrt{\frac{x}{n}}$ for x yields $x = n \cdot \left(\frac{1-c}{2}\right)^2$. □

876 **Lemma D.6.** *Let $\varepsilon_1, \dots, \varepsilon_n$ be i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables, $\sigma^2 > 0$. Let $(\varepsilon^2)^{(i)}$ denote the i -th*
 877 *largest of $\varepsilon_1^2, \dots, \varepsilon_n^2$.*

878 (i) **A constant fraction of noise cannot concentrate on less than $\Theta(n)$ points:** *For all constants*
 879 *$\alpha, c > 0$ there exists a constant $C \in (0, 1)$ such that with probability at least $1 - n^{-\alpha}$, for*
 880 *n large enough,*

$$\frac{1}{n} \sum_{i=1}^{\lfloor Cn \rfloor} (\varepsilon^2)^{(i)} < c\sigma^2.$$

881 (ii) **$\Theta(n)$ points amount to a constant fraction of noise:** *For all constants $\alpha > 0$ and $\kappa \in (0, 1)$*
 882 *there exists a constant $c > 0$ such that with probability at least $1 - n^{-\alpha}$, for n large enough,*

$$\frac{1}{n} \sum_{i=1}^{\lfloor (1-\kappa)n \rfloor} (\varepsilon^2)^{(n-i+1)} \geq c\sigma^2.$$

883 *Proof.* Without loss of generality, we can assume $\sigma^2 = 1$.

884 (i) For a constant $C \in (0, 1)$ yet to be chosen, consider the sum

$$S_{C,n} := \frac{1}{n} \sum_{i=1}^{\lfloor Cn \rfloor} (\varepsilon^2)^{(i)}.$$

885 For $T > 0$ yet to be chosen, we consider the random set $\mathcal{I}_T := \{i \in [n] \mid \varepsilon_i^2 > T\}$
 886 and denote its size by $K := |\mathcal{I}_T|$. To bound K , we note that $K = \xi_1 + \dots + \xi_n$, where
 887 $\xi_i = \mathbb{1}_{\varepsilon_i^2 > T}$. We first want to bound $p_T := \mathbb{E}\xi_i = P(\varepsilon_i^2 > T)$.

The random variables ε_i^2 follow a χ_1^2 -distribution, whose CDF we denote by $F(t)$ and whose PDF is

$$f(t) = \mathbb{1}_{(0,\infty)}(t) C_1 t^{-1/2} \exp(-t/2) \tag{D.8}$$

888 for some absolute constant C_1 . Moreover, we use Φ and ϕ to denote the CDF and PDF of
 889 $\mathcal{N}(0, 1)$, respectively.

Step 1: Tail bounds. Following Duembgen (2010), we have for $x > 0$:

$$1 - \Phi(x) > \frac{2\phi(x)}{\sqrt{4+x^2}+x} \geq \frac{2\phi(x)}{2+x+x} = \frac{\phi(x)}{1+x}$$

$$1 - \Phi(x) < \frac{2\phi(x)}{\sqrt{2+x^2}+x} \leq \frac{2\phi(x)}{1+x}.$$

Hence, for $t > 0$, we have

$$1 - F(t) = 2(1 - \Phi(\sqrt{t})) > \frac{2\phi(\sqrt{t})}{1 + \sqrt{t}} = \sqrt{\frac{2}{\pi}} \frac{\exp(-t/2)}{1 + \sqrt{t}}$$

$$1 - F(t) = 2(1 - \Phi(\sqrt{t})) < \frac{4\phi(\sqrt{t})}{1 + \sqrt{t}} = \sqrt{\frac{8}{\pi}} \frac{\exp(-t/2)}{1 + \sqrt{t}}.$$

By choosing $T := -2 \log(C\sqrt{\pi/32}) > 0$, we obtain

$$p_T = 1 - F(T) < \sqrt{\frac{8}{\pi}} \exp(-T/2) = C/2.$$

Step 2: Bounding K . The random variables ξ_i from above satisfy $\xi_i \in [0, 1]$. By Hoeffding's inequality (Steinwart and Christmann, 2008, Theorem 6.10), we have for $\tau > 0$

$$P\left(\frac{1}{n} \sum_{i=1}^n (\xi_i - \mathbb{E}\xi_i) \geq (1-0)\sqrt{\frac{\tau}{2n}}\right) \leq \exp(-\tau).$$

We choose $\tau := C^2 n/2$, such that with probability $\geq 1 - \exp(-C^2 n/2)$, we have

$$K/n - p_T = \frac{1}{n} \sum_{i=1}^n (\xi_i - \mathbb{E}\xi_i) \leq \sqrt{\frac{C^2 n/2}{2n}} = C/2.$$

Suppose that this holds. Then, $K \leq np_T + Cn/2 < Cn$ and, since K is an integer, $K \leq \lfloor Cn \rfloor$. This implies

$$S_{C,n} \leq \frac{1}{n} \left(\sum_{i=1}^K (\varepsilon^2)^{(i)} + (\lfloor Cn \rfloor - K)T \right) \leq CT + \frac{1}{n} \sum_{i=1}^K (\varepsilon^2)^{(i)}. \quad (\text{D.9})$$

We now want to bound $\sum_{i=1}^K (\varepsilon^2)^{(i)}$. To this end, we note that conditioned on $K = k$ for some $k \in [n]$, the k random variables $(\varepsilon_i)_{i \in \mathcal{I}_T}$ are i.i.d. drawn from the distribution of ε^2 given $\varepsilon^2 > T$, for $\varepsilon \sim \mathcal{N}(0, 1)$. By X, X_1, X_2, \dots , we denote i.i.d. random variables drawn from the distribution of $\varepsilon^2 - T \mid \varepsilon^2 > T$. This means that conditioned on $K = k$,

$$\sum_{i=1}^k (\varepsilon^2)^{(i)} = \sum_{i \in \mathcal{I}_T} \varepsilon_i^2 \text{ is distributed as } kT + \sum_{i=1}^k X_i. \quad (\text{D.10})$$

Step 3: Conditional expectation. The density of X is given by

$$\begin{aligned} p_X(t) &= \mathbb{1}_{t>0} \frac{f(T+t)}{1-F(T)} \stackrel{(\text{D.8})}{\leq} \mathbb{1}_{t>0} \frac{C_1(T+t)^{-1/2} \exp(-(t+T)/2)}{\sqrt{2/\pi} \exp(-T/2)/(1+\sqrt{T})} \\ &\leq \mathbb{1}_{t>0} C_2 \exp(-t/2), \end{aligned}$$

890

where we have used that for $t > 0$,

$$\frac{1+\sqrt{T}}{\sqrt{T+t}} \leq \frac{1+\sqrt{T}}{\sqrt{T}} = 1 + \frac{1}{\sqrt{T}} \leq 2$$

since $T = -2 \log(C\sqrt{\pi/32}) \geq -2 \log(\sqrt{\pi/32}) \approx 1.008$. We can now bound

$$\begin{aligned} \mathbb{E}[X] &= \int_0^\infty t p_X(t) dt \\ &\leq \int_0^\infty C_2 t \exp(-t/2) dt = 4C_2. \end{aligned} \quad (\text{D.11})$$

Step 4: Conditional subgaussian norm. For $t \geq 0$,

$$\begin{aligned} P(|X| > t) &= P(X > t) = \frac{1-F(T+t)}{1-F(T)} \leq 2 \frac{1+\sqrt{T}}{1+\sqrt{T+t}} \frac{\exp(-(T+t)/2)}{\exp(-T/2)} \\ &\leq 2 \exp(-t/2). \end{aligned}$$

891

Since the denominator $2 \exp(-t/2)$ is constant, by Proposition 2.7.1 and Definition 2.7.5 in Vershynin (2018), the subexponential norm $\|X\|_{\psi_1}$ is therefore bounded by an absolute constant C_3 . Moreover, by Exercise 2.7.10 in Vershynin (2018), we have $\|X - \mathbb{E}X\|_{\psi_1} \leq C_4 \|X\|_{\psi_1} \leq C_5$ for absolute constants C_4, C_5 .

892

893

894

Step 5: Conditional concentration. Now, Bernstein's inequality for subexponential random variables (Vershynin, 2018, Corollary 2.8.1) yields for $t \geq 0$ and some absolute constant $C_6 > 0$:

$$P\left(\left|\sum_{i=1}^k X_i - \mathbb{E}X_i\right| \geq t\right) \leq 2 \exp\left(-C_6 \min\left(\frac{t^2}{kC_5^2}, \frac{t}{C_5}\right)\right). \quad (\text{D.12})$$

We choose $t = C_5 C n$ and obtain for $k \leq C n$

$$\begin{aligned}
& P \left(\sum_{i=1}^k (\varepsilon^2)^{(i)} \geq kT + 4C_2 k + C_5 C n \mid K = k \right) \\
& \stackrel{\text{(D.10)}}{=} P \left(\sum_{i=1}^k X_i \geq 4C_2 k + C_5 C n \mid K = k \right) \\
& \stackrel{\text{(D.11)}}{\leq} P \left(\left| \sum_{i=1}^k X_i - \mathbb{E} X_i \right| \geq t \right) \\
& \stackrel{\text{(D.12)}}{\leq} 2 \exp(-C_6 C n) .
\end{aligned}$$

Step 6: Final bound. From Step 2, we know that $K \leq \lfloor C n \rfloor$ with probability $\geq 1 - \exp(-C^2 n/2)$. Moreover, in this case, Step 5 yields

$$\sum_{i=1}^K (\varepsilon^2)^{(i)} < K T + 4C_2 K + C_5 C n \leq C n (T + 4C_2 + C_5)$$

with probability $\geq 1 - \exp(-C_6 C n)$. By Eq. (D.9), we therefore have

$$S_{C,n} < C T + C (T + 4C_2 + C_5) = -4C \log(C \sqrt{\pi/32}) + C_7 C .$$

895 Since $\lim_{C \searrow 0} -C \log(C) = 0$, we can choose $C \in (0, 1)$ such that $-4C \log(C \sqrt{\pi/32}) +$
896 $C_7 C < c$ for the given constant $c > 0$ from the theorem statement, and obtain the desired
897 bound with high probability in n .

(ii) Since the ε_i^2 are non-negative and their distribution has a density, there must exist $T > 0$ with $P(\varepsilon_i^2 < T) \leq (1 - \kappa)/4$. Similar to the proof of (i), we then want to bound $K := |\{i \in [n] \mid \varepsilon_i^2 < T\}| = \xi_1 + \dots + \xi_n$ with $\xi_i = \mathbb{1}_{\varepsilon_i^2 < T}$. The $\xi_i \in [0, 1]$ are independent with $\mathbb{E} \xi_i = P(\varepsilon_i^2 < T) \leq (1 - \kappa)/4$. As in Step 2 of (i), Hoeffding's inequality then yields for $\tau > 0$:

$$P \left(\frac{1}{n} \sum_{i=1}^n (\xi_i - \mathbb{E} \xi_i) \geq (1 - 0) \sqrt{\frac{\tau}{2n}} \right) \leq \exp(-\tau) .$$

We set $\tau := (1 - \kappa)^2 n/2$, such that with probability $\geq 1 - \exp(-((1 - \kappa)^2 n/2))$, we have

$$\begin{aligned}
K/n - (1 - \kappa)/4 & \leq K/n - P(\varepsilon_i^2 < T) = \frac{1}{n} \sum_{i=1}^n (\xi_i - \mathbb{E} \xi_i) < \sqrt{\frac{(1 - \kappa)^2 n/2}{2n}} \\
& = \frac{1 - \kappa}{2} .
\end{aligned}$$

In this case, we have

$$\begin{aligned}
\frac{1}{n} \sum_{i=1}^{\lfloor (1-\kappa)n \rfloor} (\varepsilon^2)^{(n-i+1)} & \geq \frac{1}{n} (\lfloor (1 - \kappa)n \rfloor - K) T \geq \frac{1}{n} (((1 - \kappa)n - 1) - K) T \\
& \geq \left(\frac{1 - \kappa}{4} - \frac{1}{n} \right) T ,
\end{aligned}$$

898 where the right-hand side is lower bounded by $c := (1 - \kappa)T/8$ for n large enough. \square

899 The next lemma is a generalization of Lemma 9 in Buchholz (2022) to arbitrary fractions κ of the
900 training points. Therefore, for any $\kappa \in (0, 1)$ define

$$\delta_{\min}(\kappa) = n^{-1/d} \left(\frac{\kappa}{C_\rho \omega_d} \right)^{1/d} ,$$

901 **Lemma D.7** (Generalization of Lemma 9 in Buchholz (2022)). *Let $\kappa, \nu \in (0, 1)$, and let $c_\Omega > 0$*
902 *be a constant that satisfies $P_X(\text{dist}(\mathbf{x}, \partial\Omega) < c_\Omega) \leq \kappa$. Let $\mathcal{P} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ be i.i.d. points*
903 *distributed according to the measure P_X , which has lower and upper bounded density on its entire*

904 bounded open Lipschitz domain $\Omega \subseteq \mathbb{R}^d$, $C_l \leq p_X(\mathbf{x}) \leq C_u$. Then there exists a constant $\Theta > 0$
905 depending on d, C_u, ν such that with probability at least $1 - \exp\left(-\frac{3\kappa n}{7}\right)$ there exists a good subset
906 $\mathcal{P}' \subseteq \mathcal{P}$, $|\mathcal{P}'| \geq (1 - 7\kappa)n$, with the following properties: For $\mathbf{x} \in \mathcal{P}'$ we have $\text{dist}(\mathbf{x}, \partial\Omega) \geq c_\Omega$,
907 $|\mathbf{x} - \mathbf{y}| > \delta_{\min}(\kappa)$ for $\mathbf{x} \neq \mathbf{y} \in \mathcal{P}'$, and for all $\mathbf{x} \in \mathcal{P}'$ we have

$$\sum_{\mathbf{y} \in \mathcal{P}' \setminus \{\mathbf{x}\}} |\mathbf{x} - \mathbf{y}|^{-d-2\nu} \leq \frac{2\Theta\delta_{\min}(\kappa)^{-2\nu}n}{\kappa^2}.$$

Proof. First by the definition of δ_{\min} , it holds that

$$P\left(\mathbf{x}_j \in \bigcup_{i < j} B(\mathbf{x}_i, \delta_{\min})\right) \leq C_u \omega_d \delta_{\min}^d n \leq \kappa$$

Also for all $\mathbf{y} \in \Omega$

$$\begin{aligned} \mathbb{E}_{\mathbf{x}}\left((\mathbf{x} - \mathbf{y})^{-d-2s} \mathbf{1}(|\mathbf{x} - \mathbf{y}| \geq \delta_{\min})\right) &= \int_{B(\mathbf{y}, \delta_{\min})^c} |\mathbf{x} - \mathbf{y}|^{-d-2\nu} p_X(\mathbf{x}) d\mathbf{x} \\ &\leq C_u \int_{B(\mathbf{x}, \delta_{\min})^c} |\mathbf{x} - \mathbf{y}|^{-d-2\nu} d\mathbf{y} \leq \Theta \delta_{\min}^{-2\nu} \end{aligned}$$

for some $\Theta > 0$ depending only on C_u, d and ν . We conclude that for each j

$$P\left(\sum_{i < j} |\mathbf{x}_i - \mathbf{x}_j|^{-d-2\nu} \mathbf{1}(|\mathbf{x}_i - \mathbf{x}_j| > \delta_{\min}) > \frac{\Theta \delta_{\min}^{-2\nu} n}{\kappa}\right) \leq \kappa.$$

Also $P(\text{dist}(\mathbf{x}_j, \partial\Omega) < c_\Omega) < \kappa$. The union bound implies that

$$\begin{aligned} &P\left(\mathbf{x}_j \notin \bigcup_{i < j} B(\mathbf{x}_i, \delta_{\min}), \sum_{i < j} |\mathbf{x}_i - \mathbf{x}_j|^{-d-2\nu} \mathbf{1}_{|\mathbf{x}_i - \mathbf{x}_j| > \delta_{\min}} < \frac{\Theta \delta_{\min}^{-2\nu} n}{\kappa}, \text{dist}(\mathbf{x}_j, \partial\Omega) > c_\Omega\right) \\ &= P\left(\mathbf{x}_j \notin \bigcup_{i < j} B(\mathbf{x}_i, \delta_{\min}), \sum_{i < j} |\mathbf{x}_i - \mathbf{x}_j|^{-d-2\nu} < \frac{\Theta \delta_{\min}^{-2\nu} n}{\kappa}, \text{dist}(\mathbf{x}_j, \partial\Omega) > c_\Omega\right) \geq 1 - 3\kappa. \end{aligned}$$

908 We use a martingale construction similar to the one in Lemma 7 of Buchholz (2022) by defining

$$E_j := \left\{ \mathbf{x}_j \in \bigcup_{i < j} B(\mathbf{x}_i, \delta_{\min}), \text{ or } \sum_{i < j} |\mathbf{x}_i - \mathbf{x}_j|^{-d-2\nu} \geq \frac{\Theta \delta_{\min}^{-2\nu} n}{\kappa}, \text{ or } \text{dist}(\mathbf{x}_j, \partial\Omega) \leq c_\Omega \right\}.$$

909 Now define $S_n := \sum_{i=1}^n \mathbf{1}_{E_i}$. Using the filtration $\mathcal{F}_i = \sigma(\mathbf{x}_1, \dots, \mathbf{x}_i)$, S_n can be decomposed into
910 $S_n = M_n + A_n$, where M_n is a martingale and A_n is predictable with respect to \mathcal{F}_n . We then get
911 $A_n \leq \sum_{i=1}^n P(E_i | \mathcal{F}_{i-1}) \leq 3\kappa n$ as well as $\text{Var}(M_i | \mathcal{F}_{i-1}) \leq 3\kappa$. Hence Freedman's inequality
912 Theorem D.8 yields

$$P(S_n \geq 6\kappa n) \leq P(A_n \geq 3\kappa n) + P(M_n \geq 3\kappa n) \leq \exp\left(-\frac{3\kappa n}{7}\right).$$

This implies that with probability at least $1 - \exp\left(-\frac{3\kappa n}{7}\right)$ we can find a subset $\mathcal{P}_s = \{\mathbf{z}_1, \dots, \mathbf{z}_m\}$
with $|\mathcal{P}_s| \geq (1 - 6\kappa)n$ on which it holds that $\min_{i \neq j} |\mathbf{z}_i - \mathbf{z}_j| \geq \delta_{\min}$, $\text{dist}(\mathbf{z}_j, \partial\Omega) \geq c_\Omega$ and

$$\sum_{i \neq j} |\mathbf{z}_i - \mathbf{z}_j|^{-d-2\nu} \leq \frac{2\Theta \delta_{\min}^{-2\nu} n^2}{\kappa}.$$

Using Markov's inequality we see that there are at most κn points in \mathcal{P}_s such that

$$\sum_{\mathbf{z}' \in \mathcal{P}_s, \mathbf{z}' \neq \mathbf{z}} |\mathbf{z} - \mathbf{z}'|^{-d-2\nu} \geq \frac{2\Theta \delta_{\min}^{-2\nu} n}{\kappa^2}.$$

913 Removing those points we find a subset $\mathcal{P}' \subset \mathcal{P}_s$ such that $|\mathcal{P}'| \geq (1 - 7\kappa)n$ and for each $z \in \mathcal{P}'$

$$\sum_{z' \in \mathcal{P}_s, z \neq z'} |z - z'|^{-d-2\nu} \leq \frac{2\Theta \delta_{\min}^{-2\nu} n}{\kappa^2}. \quad \square$$

Theorem D.8 (Freedman’s inequality, Theorem 6.1 in Chung and Lu (2006)). *Let M_i be a discrete martingale adapted to the filtration \mathcal{F}_i with $M_0 = 0$ that satisfies for all $i \geq 0$*

$$\begin{aligned} |M_{i+1} - M_i| &\leq K \\ \text{Var}(M_i | \mathcal{F}_{i-1}) &\leq \sigma_i^2. \end{aligned}$$

Then

$$P(M_n - \mathbb{E}(M_n) \geq \lambda) \leq e^{-\frac{\lambda^2}{2 \sum_{i=1}^n \sigma_i^2 + K\lambda/3}}.$$

914 E Translating between \mathbb{R}^d and \mathbb{S}^d

915 Since the RKHS of the ReLU NTK and NNGP kernels mentioned in Theorem 3 are equivalent to
 916 the Sobolev spaces $H^{(d+1)/2}(\mathbb{S}^d)$ and $H^{(d+3)/2}(\mathbb{S}^d)$, respectively (Chen and Xu, 2021, Bietti and
 917 Bach, 2021) (detailed summary in Appendix B.4). Inconsistency of functions in these RKHS that
 918 fulfill Assumptions (O) and (N), as in Theorem 1, follows immediately by adapting Theorem 1 via
 919 Lemma E.1. In particular, inconsistency holds for the gradient flow and gradient descent estimators
 920 $f_{t,\rho}$ and $f_{t,\rho}^{\text{GD}}$ as soon as they overfit with lower bounded probability.

921 For arbitrary open sphere caps $T := \{\mathbf{x} \in \mathbb{S}^d \mid x_{d+1} < v\}$, $v \in (-1, 1)$, and the open unit ball
 922 $B_1(0) := \{\mathbf{y} \in \mathbb{R}^d \mid \|\mathbf{y}\|_2 < 1\}$, define the scaled stereographic projection $\phi : T \rightarrow B_1(0) \subseteq \mathbb{R}^d$ as

$$\phi(x_1, \dots, x_{d+1}) = \left(\frac{c_v x_1}{1 - x_{d+1}}, \dots, \frac{c_v x_d}{1 - x_{d+1}} \right),$$

923 where the normalization constant $c_v = \sqrt{\frac{1-v}{1+v}}$ ensures surjectivity.

924 Straightforward calculations show that ϕ defines a diffeomorphism. Its inverse $\phi^{-1} : B_1(0) \rightarrow T$ is
 925 given by

$$\phi^{-1}(y_1, \dots, y_d) = \left(\frac{2c_v^{-1}y_1}{c_v^{-2}\|\mathbf{y}\|_2^2 + 1}, \dots, \frac{2c_v^{-1}y_d}{c_v^{-2}\|\mathbf{y}\|_2^2 + 1}, \frac{c_v^{-2}\|\mathbf{y}\|_2^2 - 1}{c_v^{-2}\|\mathbf{y}\|_2^2 + 1} \right).$$

926 We can translate kernel learning with the kernel k on \mathbb{S}^d and the probability distribution P , where
 927 P_X is supported on T , to kernel learning with a transformed kernel \tilde{k} and \tilde{P} using a sufficiently
 928 smooth diffeomorphism like $\phi : T \rightarrow B_1(0) \subseteq \mathbb{R}^d$. If the RKHS of k is equivalent to $H^s(\mathbb{S}^d)$ then
 929 the RKHS of \tilde{k} is equivalent to $H^s(B_1(0))$. We formalize this argument in the following lemma. As
 930 a consequence it suffices to prove all inconsistency results for Sobolev kernels on $B_1(0)$.

931 **Lemma E.1 (Transfer to sphere caps).** *Let k be a kernel on \mathbb{S}^d whose RKHS is equivalent to a*
 932 *Sobolev space $H^s(\mathbb{S}^d)$. For fixed $v \in (-1, 1)$, consider an “open sphere cap” $T := \{\mathbf{x} \in \mathbb{S}^d \mid$
 933 $x_{d+1} < v\}$. Furthermore, consider a distribution P such that P_X is supported on T and has lower
 934 and upper bounded density p_X on T , i.e. $0 < C_l \leq p_X(\mathbf{x}) \leq C_u < \infty$ for all $\mathbf{x} \in T$. Then*

- 935 • $\tilde{k}(\mathbf{x}, \mathbf{x}') := k(\phi^{-1}(\mathbf{x}), \phi^{-1}(\mathbf{x}'))$ defines a positive definite kernel on $B_1(0) \subseteq \mathbb{R}^d$ whose
- 936 RKHS is equivalent to the Sobolev space $H^s(B_1(0))$,
- 937 • $\tilde{P} := P \circ \psi^{-1}$ with $\psi(\mathbf{x}, y) := (\phi(\mathbf{x}), y)$ defines a probability distribution such that $\tilde{P}_{\tilde{X}}$
- 938 has lower and upper bounded density on $B_1(0) \subseteq \mathbb{R}^d$,

939 and kernel learning with (k, P) or with (\tilde{k}, \tilde{P}) is equivalent in the following sense:

940 For every function $f \in \mathcal{H}(k|_T)$ the transformed function $\tilde{f} = f \circ \phi^{-1} \in \mathcal{H}(\tilde{k})$ has the same RKHS
 941 norm, i.e. $\|f\|_{\mathcal{H}(k|_T)} = \|\tilde{f}\|_{\mathcal{H}(\tilde{k})}$. Furthermore, the excess risks of f over P and \tilde{f} over \tilde{P} coincide,
 942 i.e.

$$\mathbb{E}_{\mathbf{x} \sim P_X} (f(\mathbf{x}) - f_P^*(\mathbf{x}))^2 = \mathbb{E}_{\tilde{\mathbf{x}} \sim \tilde{P}_X} (\tilde{f}(\tilde{\mathbf{x}}) - \tilde{f}_P^*(\tilde{\mathbf{x}}))^2,$$

943 where $\tilde{f}_P^*(\tilde{\mathbf{x}}) = \mathbb{E}_{(\tilde{X}, \tilde{Y}) \sim \tilde{P}} (\tilde{Y} | \tilde{X} = \tilde{\mathbf{x}})$ denotes the Bayes optimal predictor under \tilde{P} .

944 **Remark E.2.** Many kernel regression estimators can be explicitly written as $f_D^k(\mathbf{x}) =$
 945 $\hat{f}_n(k(\mathbf{x}, \mathbf{X}), k(\mathbf{X}, \mathbf{X}), \mathbf{y})$ where $\hat{f}_n : \mathbb{R}^n \times \mathbb{R}^{n \times n} \times \mathbb{R}^n \rightarrow \mathbb{R}$ denotes a measurable function
 946 for all $n \in \mathbb{N}$. Then the explicit form is preserved under the transformation, i.e. $f \circ \phi^{-1} = f_{\tilde{D}}^{\tilde{k}}$ with
 947 the transformed data set $\tilde{D} = \{(\phi(\mathbf{x}_i), y_i)\}_{i \in [n]}$. ◀

948 *Proof of Lemma E.1. Step 1: Bounded density.* For $i \in [d], j \in [d+1]$, the partial derivatives of ϕ
 949 are given by

$$\partial_{x_j} \phi_i(\mathbf{x}) = \begin{cases} \frac{c_v}{1-x_{d+1}}, & \text{for } i = j, \\ \frac{c_v x_i}{(1-x_{d+1})^2}, & \text{for } i \in [d], j = d+1, \\ 0, & \text{otherwise.} \end{cases}$$

950 Given an arbitrary multi-index α , the partial derivatives $\partial_\alpha \phi_i \in L^2(T)$, $\partial_\alpha \phi_j^{-1} \in L^2(B_1(0))$ are
 951 bounded for all $i \in [d], j \in [d+1]$, using $x_{d+1} \leq v < 1$ and the inverse function theorem.

952 Now define $\tilde{k}(\mathbf{x}, \mathbf{x}') := k(\phi^{-1}(\mathbf{x}), \phi^{-1}(\mathbf{x}'))$, $\psi(\mathbf{x}, y) := (\phi(\mathbf{x}), y)$ and $\tilde{P} := P \circ \psi^{-1}$. Then using
 953 integration by substitution (Stroock et al., 2011, Theorem 5.2.16), the Lebesgue density of \tilde{P}_X is
 954 given by

$$p_{\tilde{X}}(\tilde{\mathbf{x}}) = p_X(\phi^{-1}(\tilde{\mathbf{x}})) J\phi^{-1}(\tilde{\mathbf{x}}),$$

955 where

$$J\phi^{-1}(\tilde{\mathbf{x}}) := \left[\det \left(\left(\langle \partial_i \phi^{-1}(\tilde{\mathbf{x}}), \partial_j \phi^{-1}(\tilde{\mathbf{x}}) \rangle_{\mathbb{R}^{d+1}} \right)_{i,j \in \{1, \dots, d\}} \right) \right]^{1/2}.$$

956 $J\phi$ and $J\phi^{-1}$ can be continuously extended to \bar{T} and $\bar{B}_1(0)$, respectively. Then, since $J\phi^{-1}$ is
 957 continuous on a compact set and because ϕ with the extended domain remains a diffeomorphism so
 958 that $J\phi^{-1}$ cannot attain the value 0, there exists a constant $C_\phi > 0$ such that $\frac{1}{C_\phi} \leq J\phi^{-1}(\tilde{\mathbf{x}}) \leq C_\phi$
 959 for all $\tilde{\mathbf{x}} \in B_1(0)$. Hence, $p_{\tilde{X}}$ is lower and upper bounded.

960 **Step 2: Excess risks coincide.** If $(\tilde{X}, \tilde{Y}) \sim \tilde{P}$, the Bayes predictor of \tilde{Y} given \tilde{X} is given by
 961 $\tilde{f}^*(\tilde{\mathbf{x}}) = \mathbb{E}(\tilde{Y} | \tilde{X} = \tilde{\mathbf{x}}) = f^*(\phi^{-1}(\tilde{\mathbf{x}}))$.

Let $\pi_1(\mathbf{x}, y) = \mathbf{x}$ be the projection onto the first component. Then, $\phi(\pi_1(\mathbf{x}, y)) = \phi(\mathbf{x}) =$
 $\pi_1(\phi(\mathbf{x}), y) = \pi_1(\psi(\mathbf{x}, y))$ and hence

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim P_X} (f(\mathbf{x}) - f^*(\mathbf{x}))^2 &= \mathbb{E}_{(\mathbf{x}, y) \sim P} (f(\pi_1(\mathbf{x}, y)) - f^*(\pi_1(\mathbf{x}, y)))^2 \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim P} (f(\phi^{-1}(\phi(\pi_1(\mathbf{x}, y)))) - f^*(\phi^{-1}(\phi(\pi_1(\mathbf{x}, y))))^2 \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim P} (\tilde{f}(\pi_1(\psi(\mathbf{x}, y))) - \tilde{f}^*(\pi_1(\psi(\mathbf{x}, y))))^2 \\ &= \mathbb{E}_{(\mathbf{x}, y) \sim \tilde{P}} (\tilde{f}(\pi_1(\mathbf{x}, y)) - \tilde{f}^*(\pi_1(\mathbf{x}, y)))^2 \\ &= \mathbb{E}_{\tilde{\mathbf{x}} \sim \tilde{P}_{\tilde{X}}} (\tilde{f}(\tilde{\mathbf{x}}) - \tilde{f}^*(\tilde{\mathbf{x}}))^2. \end{aligned}$$

962 **Step 3: Transformed RKHS.** We want to show that $\mathcal{H}(k|_T) \rightarrow \mathcal{H}(\tilde{k}), f \mapsto f \circ \phi^{-1}$ defines
 963 an isometric isomorphism, which especially shows the statement $\|f\|_{\mathcal{H}(k|_T)} = \|\tilde{f}\|_{\mathcal{H}(\tilde{k})}$ from the
 964 proposition. For this, we use the following theorem characterizing RKHSs:

Theorem E.3 (Theorem 4.21 in Steinwart and Christmann (2008)). *Let $k : X \times X \rightarrow \mathbb{R}$ be a positive
 definite kernel function with feature space H_0 and feature map $\Phi_0 : X \rightarrow H_0$. Then*

$$H = \{f : X \rightarrow \mathbb{R} \mid \exists w \in H_0 : f = \langle w, \Phi_0(\cdot) \rangle_{H_0}\} \text{ with} \\ \|f\|_H := \inf\{\|w\|_{H_0} : f = \langle w, \Phi_0(\cdot) \rangle_{H_0}\},$$

965 *is the only RKHS for which k is a reproducing kernel.*

A feature map for $k|_T$ is given by $\Phi : T \rightarrow \mathcal{H}(k|_T)$, $\Phi(\mathbf{x}) = k(\mathbf{x}, \cdot)$. Hence a feature map for \tilde{k} is
 given by $\Phi \circ \phi^{-1} : B_1(0) \rightarrow \mathcal{H}(k|_T)$. Theorem E.3 states that

$$\mathcal{H}(k|_T) = \{f : T \rightarrow \mathbb{R} \mid \exists w \in \mathcal{H}(k|_T) : f = \langle w, \Phi(\cdot) \rangle_{\mathcal{H}(k|_T)}\} \text{ with} \quad (\text{E.1}) \\ \|f\|_{\mathcal{H}(k|_T)} := \inf\{\|w\|_{\mathcal{H}(k|_T)} : f = \langle w, \Phi(\cdot) \rangle_{\mathcal{H}(k|_T)}\},$$

as well as

$$\begin{aligned} \mathcal{H}(\tilde{k}) &= \left\{ \tilde{f} : B_1(0) \rightarrow \mathbb{R} \mid \exists w \in \mathcal{H}(k|_T) : \tilde{f} = \langle w, \Phi \circ \phi^{-1}(\cdot) \rangle_{\mathcal{H}(k|_T)} \right\} \text{ with} \quad (\text{E.2}) \\ \|\tilde{f}\|_{\mathcal{H}(\tilde{k})} &:= \inf \{ \|w\|_{\mathcal{H}(k|_T)} : \tilde{f} = \langle w, \Phi \circ \phi^{-1}(\cdot) \rangle_{\mathcal{H}(k|_T)} \}. \end{aligned}$$

966 As ϕ^{-1} is bijective, this characterization induces an isometric isomorphism between $\mathcal{H}(k|_T)$ and
 967 $\mathcal{H}(\tilde{k})$ by mapping $f = \langle w, \Phi(\cdot) \rangle_{\mathcal{H}(k|_T)} \in \mathcal{H}(k|_T)$ to $\tilde{f} = f \circ \phi^{-1} = \langle w, \Phi \circ \phi^{-1}(\cdot) \rangle_{\mathcal{H}(k|_T)} \in \mathcal{H}(\tilde{k})$.
 968 This shows $\|f\|_{\mathcal{H}(k|_T)} = \|\tilde{f}\|_{\mathcal{H}(\tilde{k})}$.

Step 4: RKHS of \tilde{k} . We now show that the RKHS of \tilde{k} , denoted as $\mathcal{H}(\tilde{k})$, is equivalent to $H^s(B_1(0))$. To this end, denoting $\mathcal{A} \circ \phi := \{f \circ \phi \mid f \in \mathcal{A}\}$ and $\mathcal{A}|_T := \{f|_T \mid f \in \mathcal{A}\}$, we show the following equality of sets (ignoring the norms):

$$\mathcal{H}(\tilde{k}) \circ \phi \stackrel{\text{(I)}}{=} \mathcal{H}(k|_T) \stackrel{\text{(II)}}{=} \mathcal{H}(k)|_T \stackrel{\text{(III)}}{=} H^s(\mathbb{S}^d)|_T \stackrel{\text{(IV)}}{=} H^s(B_1(0)) \circ \phi.$$

969 Since ϕ is bijective, this implies $\mathcal{H}(\tilde{k}) = H^s(B_1(0))$ as sets, and the norm equivalence then follows
 970 from Lemma F.9.

971 Equality (I) follows from Step 3. Equality (II) follows from Theorem E.3 by observing that if Φ is a
 972 feature map for k , then $\Phi|_T$ is a feature map for $k|_T$. Equality (III) holds by assumption. To show
 973 (IV), we need a characterization of $H^s(\mathbb{S}^d)$ that allows to work with charts like ϕ .

974 **Step 4.1: Chart-based characterization of $H^s(\mathbb{S}^d)$.** A trivialization of a Riemannian manifold
 975 (M, g) with bounded geometry of dimension d consists of a locally finite open covering $\{U_\alpha\}_{\alpha \in I}$
 976 of M , smooth diffeomorphisms $\kappa_\alpha : V_\alpha \subset \mathbb{R}^d \rightarrow U_\alpha$, also called charts, and a partition of unity
 977 $\{h_\alpha\}_{\alpha \in I}$ of M that fulfills $\text{supp}(h_\alpha) \subseteq U_\alpha$, $0 \leq h_\alpha \leq 1$ and $\sum_{\alpha \in I} h_\alpha = 1$. An admissible
 978 trivialization of (M, g) is a uniformly locally finite trivialization of M that is compatible with
 979 geodesic coordinates, for details see (Schneider and Große, 2013, Definition 12).

In our case, define an open neighborhood of T by $U_1 := \{\mathbf{x} \in \mathbb{S}^d \mid x_{d+1} < v + \varepsilon\}$ with some
 $\varepsilon \in (0, 1 - v)$ arbitrary but fixed, and $U_2 := \{\mathbf{x} \in \mathbb{S}^d \mid x_{d+1} > v + \varepsilon/2\}$. It holds that $U_1 \cup U_2 = \mathbb{S}^d$.
 Moreover, there exists an appropriate partition of unity consisting of C^∞ functions $h_1, h_2 : \mathbb{S}^d \rightarrow$
 $[0, 1]$. Especially, we have $h_1(T) \subseteq h_1(U_2^c) = \{1\}$. Let $\phi_1 : U_1 \rightarrow B_{r_1}(0)$ denote the stereographic
 projection with respect to $\mathbf{x}_0 = (0, \dots, 0, 1)$ as above, scaled such that $\phi_1|_T = \phi$ and hence
 $\phi_1(T) = B_1(0)$. Similarly, let $\phi_2 : U_2 \rightarrow B_{r_2}(0)$ denote an arbitrarily scaled stereographic
 projection with respect to $\mathbf{x}_0 = (0, \dots, 0, -1)$. Then $(\{U_1, U_2\}, \{\phi_1^{-1}, \phi_2^{-1}\}, \{h_1, h_2\})$ yields an
 admissible trivialization of \mathbb{S}^d consisting of only two charts. A detailed derivation can be found in
 (Hubbert et al., 2015, Section 1.7). Therefore (Schneider and Große, 2013, Theorem 14) lets us define
 the Sobolev norm on \mathbb{S}^d (up to equivalence) as¹

$$\begin{aligned} \|g\|_{H^s(\mathbb{S}^d)} &:= \left(\sum_{\alpha \in I} \|(h_\alpha g) \circ \kappa_\alpha\|_{H^s(\mathbb{R}^d)}^2 \right)^{1/2} \\ &= \left(\|(h_1 g) \circ \phi_1^{-1}\|_{H^s(\mathbb{R}^d)}^2 + \|(h_2 g) \circ \phi_2^{-1}\|_{H^s(\mathbb{R}^d)}^2 \right)^{1/2}, \end{aligned}$$

980 for any distribution $g \in \mathcal{D}'(\mathbb{S}^d)$ (i.e. any continuous linear functional on $C_c^\infty(\mathbb{S}^d)$). Then $g \in H^s(\mathbb{S}^d)$
 981 if and only if $\|g\|_{H^s(\mathbb{S}^d)} < \infty$.

Step 4.2: Showing (IV). First, let $g \in H^s(\mathbb{S}^d)$. Then, as we saw in Step 4.1, we must have
 $\|(h_1 g) \circ \phi_1^{-1}\|_{H^s(\mathbb{R}^d)} < \infty$ and thus $(h_1 g) \circ \phi_1^{-1} \in H^s(\mathbb{R}^d)$. By our discussion in Appendix B.1,
 we then have

$$(g|_T) \circ \phi^{-1} = ((h_1 g) \circ \phi_1^{-1})|_{B_1(0)} \in H^s(B_1(0)),$$

982 which shows $g|_T \in H^s(B_1(0)) \circ \phi$.

Now, let $f \in H^s(B_1(0))$. Then, again following our discussion in Appendix B.1, there exists an
 extension $\tilde{f} \in H^s(\mathbb{R}^d)$ with $\tilde{f}|_{B_1(0)} = f$. The set $\mathcal{B} := \phi_1(U_1 \setminus U_2)$ is a closed ball $\overline{B_r(0)}$ of radius

¹Here, the norms are taken on $H^s(\mathbb{R}^d)$ since the respective functions can be extended to \mathbb{R}^d by zero outside
 of their domain of definition, thanks to the properties of the partition of unity.

$1 < r < r_1$. Hence, we can find $\varphi \in C^\infty(\mathbb{R}^d)$ with $\varphi(B_1(0)) = \{1\}$ and $\varphi(\overline{(B_r(0))^c}) = \{0\}$. Since φ is smooth with compact support, we have $\varphi \cdot \bar{f} \in H^s(\mathbb{R}^d)$. Define

$$f_{\mathbb{S}^d} : \mathbb{S}^d \rightarrow \mathbb{R}, \mathbf{x} \mapsto \begin{cases} (\varphi \cdot \bar{f})(\phi(\mathbf{x})) & , \mathbf{x} \in U_1 \\ 0 & , \mathbf{x} \notin U_1 . \end{cases}$$

By construction, we have $f_{\mathbb{S}^d}(\mathbf{x}) = 0$ for all $\mathbf{x} \in U_2$. Hence, the equivalent Sobolev norm from Step 4.1 is

$$\begin{aligned} \|f_{\mathbb{S}^d}\|_{H^s(\mathbb{S}^d)} &= \left(\|(h_1 f_{\mathbb{S}^d}) \circ \phi_1^{-1}\|_{H^s(\mathbb{R}^d)}^2 + \|(h_2 f_{\mathbb{S}^d}) \circ \phi_2^{-1}\|_{H^s(\mathbb{R}^d)}^2 \right)^{1/2} \\ &= \|(h_1 \circ \phi_1^{-1}) \cdot \varphi \cdot \bar{f}\|_{H^s(\mathbb{R}^d)} < \infty , \end{aligned}$$

983 which shows $f_{\mathbb{S}^d} \in H^s(\mathbb{S}^d)$. But then, $f \circ \phi = f_{\mathbb{S}^d}|_T \in H^s(\mathbb{S}^d)|_T$.

984 In total, we obtain $H^s(\mathbb{S}^d)|_T = H^s(B_1(0)) \circ \phi$, which shows (IV). \square

985 F Spectral lower bound

986 F.1 General lower bounds

987 A common first step to analyze the expected excess risk caused by label noise is to perform a bias-
988 variance decomposition and integrate over \mathbf{y} first (see e.g. Liang and Rakhlin, 2020, Hastie et al.,
989 2022, Holzmüller, 2021), which is also used in the following lemma.

Lemma F.1. *Consider an estimator of the form $f_{\mathbf{X},\mathbf{y}}(\mathbf{x}) = (\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top \mathbf{y}$. If $\text{Var}_P(y|\mathbf{x}) \geq \sigma^2$ for P_X -almost all \mathbf{x} , then the expected excess risk satisfies*

$$\mathbb{E}_D R_P(f_{\mathbf{X},\mathbf{y}}) - R_P^* \geq \sigma^2 \mathbb{E}_{\mathbf{X},\mathbf{x}} \text{tr}(\mathbf{v}_{\mathbf{X},\mathbf{x}}(\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top) .$$

Proof. A standard bias-variance decomposition lets us lower-bound the expected excess risk by the estimator variance due to the label noise, which can then be further simplified:

$$\begin{aligned} \mathbb{E}_D R_P(f_{\mathbf{X},\mathbf{y}}) - R_P^* &\geq \mathbb{E}_{\mathbf{X},\mathbf{x}} \left(\mathbb{E}_{\mathbf{y}|\mathbf{X}} [f_{\mathbf{X},\mathbf{y}}(\mathbf{x})^2] - (\mathbb{E}_{\mathbf{y}|\mathbf{X}} [f_{\mathbf{X},\mathbf{y}}(\mathbf{x})])^2 \right) . \\ &= \mathbb{E}_{\mathbf{X},\mathbf{x}} \mathbb{E}_{\mathbf{y}|\mathbf{X}} (f_{\mathbf{X},\mathbf{y}}(\mathbf{x}) - \mathbb{E}_{\mathbf{y}|\mathbf{X}} f_{\mathbf{X},\mathbf{y}}(\mathbf{x}))^2 \\ &= \mathbb{E}_{\mathbf{X},\mathbf{x}} \mathbb{E}_{\mathbf{y}|\mathbf{X}} (\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top (\mathbf{y} - \mathbb{E}_{\mathbf{y}|\mathbf{X}} \mathbf{y}) (\mathbf{y} - \mathbb{E}_{\mathbf{y}|\mathbf{X}} \mathbf{y})^\top \mathbf{v}_{\mathbf{X},\mathbf{x}} \\ &= \mathbb{E}_{\mathbf{X},\mathbf{x}} (\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top \left[\mathbb{E}_{\mathbf{y}|\mathbf{X}} (\mathbf{y} - \mathbb{E}_{\mathbf{y}|\mathbf{X}} \mathbf{y}) (\mathbf{y} - \mathbb{E}_{\mathbf{y}|\mathbf{X}} \mathbf{y})^\top \right] \mathbf{v}_{\mathbf{X},\mathbf{x}} \\ &= \mathbb{E}_{\mathbf{X},\mathbf{x}} (\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top \text{Cov}(\mathbf{y}|\mathbf{X}) \mathbf{v}_{\mathbf{X},\mathbf{x}} . \end{aligned}$$

Here, the conditional covariance matrix can be lower bounded in terms of the Loewner order (which is defined as $A \succeq B \Leftrightarrow A - B$ positive semi-definite):

$$\text{Cov}(\mathbf{y}|\mathbf{X}) = \begin{pmatrix} \text{Var}(y_1|\mathbf{x}_1) & & \\ & \ddots & \\ & & \text{Var}(y_n|\mathbf{x}_n) \end{pmatrix} \succeq \sigma^2 \mathbf{I}_n$$

since the labels y_i are conditionally independent given \mathbf{X} . We therefore obtain

$$\begin{aligned} \mathbb{E}_D R_P(f_{\mathbf{X},\mathbf{y}}) - R_P^* &\geq \mathbb{E}_{\mathbf{X},\mathbf{x}} (\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top \text{Cov}(\mathbf{y}|\mathbf{X}) \mathbf{v}_{\mathbf{X},\mathbf{x}} \\ &\geq \sigma^2 \mathbb{E}_{\mathbf{X},\mathbf{x}} \text{tr}((\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top \mathbf{v}_{\mathbf{X},\mathbf{x}}) \\ &= \sigma^2 \mathbb{E}_{\mathbf{X},\mathbf{x}} \text{tr}(\mathbf{v}_{\mathbf{X},\mathbf{x}}(\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top) . \end{aligned} \quad \square$$

Proposition 4 (Spectral lower bound). *Assume that the kernel matrix $k(\mathbf{X}, \mathbf{X})$ is almost surely positive definite, and that $\text{Var}(y|\mathbf{x}) \geq \sigma^2$ for P_X -almost all \mathbf{x} . Then, the expected excess risk satisfies*

$$\mathbb{E}_D R_P(f_{t,\rho}) - R_P^* \geq \frac{\sigma^2}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{X}} \frac{\lambda_i(k_*(\mathbf{X}, \mathbf{X})/n) (1 - e^{-2t(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + \rho)})^2}{(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + \rho)^2} . \quad (3)$$

Proof. Recall from Eq. (1) that

$$\begin{aligned} f_{t,\rho}(\mathbf{x}) &= k(\mathbf{x}, \mathbf{X}) \mathbf{A}_{t,\rho}(\mathbf{X}) \mathbf{y}, \\ \mathbf{A}_{t,\rho}(\mathbf{X}) &:= \left(\mathbf{I}_n - e^{-\frac{2}{n}t(k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n)} \right) (k(\mathbf{X}, \mathbf{X}) + \rho n \mathbf{I}_n)^{-1}. \end{aligned}$$

By setting $(\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top := k(\mathbf{x}, \mathbf{X}) \mathbf{A}_{t,\rho}(\mathbf{X})$, we can write $f_{\mathbf{X},\mathbf{y},t,\rho}(\mathbf{x}) := f_{t,\rho}(\mathbf{x}) = (\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top \mathbf{y}$. Using Lemma F.1, we then obtain

$$\begin{aligned} \mathbb{E}_D R_P(f_{\mathbf{X},\mathbf{y},t,\rho}) - R_P^* &\geq \sigma^2 \mathbb{E}_{\mathbf{X},\mathbf{x}} \text{tr}(\mathbf{v}_{\mathbf{X},\mathbf{x}}(\mathbf{v}_{\mathbf{X},\mathbf{x}})^\top) \\ &= \sigma^2 \mathbb{E}_{\mathbf{X},\mathbf{x}} \text{tr}(\mathbf{A}_{t,\rho}(\mathbf{X})^\top k(\mathbf{X}, \mathbf{x}) k(\mathbf{x}, \mathbf{X}) \mathbf{A}_{t,\rho}(\mathbf{X})). \end{aligned}$$

Since

$$(\mathbb{E}_{\mathbf{x}} k(\mathbf{X}, \mathbf{x}) k(\mathbf{x}, \mathbf{X}))_{ij} = \mathbb{E}_{\mathbf{x}} k(\mathbf{x}_i, \mathbf{x}) k(\mathbf{x}, \mathbf{x}_j) = k_*(\mathbf{x}_i, \mathbf{x}_j) = k_*(\mathbf{X}, \mathbf{X})_{ij},$$

we conclude

$$\begin{aligned} \mathbb{E}_D R_P(f_{\mathbf{X},\mathbf{y},t,\rho}) - R_P^* &\geq \sigma^2 \mathbb{E}_{\mathbf{X}} \text{tr}(\mathbf{A}_{t,\rho}^\top k_*(\mathbf{X}, \mathbf{X}) \mathbf{A}_{t,\rho}) \\ &= \sigma^2 \mathbb{E}_{\mathbf{X}} \text{tr}(k_*(\mathbf{X}, \mathbf{X}) \mathbf{A}_{t,\rho}(\mathbf{X}) \mathbf{A}_{t,\rho}(\mathbf{X})^\top). \end{aligned}$$

Richter (1958) showed (see also Mirsky, 1959) that for two symmetric matrices \mathbf{B}, \mathbf{C} , we have $\text{tr}(\mathbf{B}\mathbf{C}) \geq \sum_{i=1}^n \lambda_i(\mathbf{B}) \lambda_{n+1-i}(\mathbf{C})$. We can therefore conclude

$$\mathbb{E}_D R_P(f_{\mathbf{X},\mathbf{y},t,\rho}) - R_P^* \geq \sigma^2 \mathbb{E}_{\mathbf{X}} \sum_{i=1}^n \lambda_i(k_*(\mathbf{X}, \mathbf{X})) \lambda_{n+1-i}(\mathbf{A}_{t,\rho}(\mathbf{X}) \mathbf{A}_{t,\rho}(\mathbf{X})^\top).$$

As $\mathbf{A}_{t,\rho}(\mathbf{X}) \mathbf{A}_{t,\rho}(\mathbf{X})^\top$ is built only out of the matrices $k(\mathbf{X}, \mathbf{X})$ and \mathbf{I}_n , it is not hard to see that $\mathbf{A}_{t,\rho}(\mathbf{X}) \mathbf{A}_{t,\rho}(\mathbf{X})^\top$ has the same eigenbasis as $k(\mathbf{X}, \mathbf{X})$ with eigenvalues

$$\tilde{\lambda}_i := \left(\frac{1 - e^{-\frac{2}{n}t(\lambda_i(k(\mathbf{X}, \mathbf{X}) + \rho n)}}{\lambda_i(k(\mathbf{X}, \mathbf{X}) + \rho n)} \right)^2 = \frac{1}{n^2} \left(\frac{1 - e^{-2t(\lambda_i(k(\mathbf{X}, \mathbf{X})/n + \rho)}}{\lambda_i(k(\mathbf{X}, \mathbf{X})/n + \rho)} \right)^2.$$

It remains to order these eigenvalues correctly. To this end, we observe that for $\lambda > 0$, the function $g(\lambda) := \frac{1 - e^{-2t\lambda}}{\lambda}$ satisfies

$$g'(\lambda) = \frac{2t\lambda e^{-2t\lambda} - (1 - e^{-2t\lambda})}{\lambda^2} = \frac{(2t\lambda + 1)e^{-2t\lambda} - 1}{\lambda^2} \leq \frac{e^{2t\lambda} e^{-2t\lambda} - 1}{\lambda^2} = 0.$$

Therefore, g is nonincreasing, hence the sequence $(\tilde{\lambda}_i)$ is nondecreasing and thus

$$\lambda_{n+1-i}(\mathbf{A}_{t,\rho} \mathbf{A}_{t,\rho}^\top) = \tilde{\lambda}_i,$$

990 from which the claim follows. \square

Theorem F.2. *Let k be a kernel on a compact set Ω and let P_X be supported on Ω . Suppose that $k(\mathbf{X}, \mathbf{X})$ is almost surely positive definite and that $\text{Var}(y|\mathbf{x}) \geq \sigma^2$ for P_X -almost all \mathbf{x} . Fix constants $c > 0$ and $q, C \geq 1$. Suppose that $\lambda_i := \lambda_i(T_{k, P_X}) \geq ci^{-q}$. Let $\mathcal{I}(n)$ be the set of all $i \in [n]$ for which*

$$\begin{aligned} \lambda_i/C &\leq \lambda_i(k(\mathbf{X}, \mathbf{X})/n) \leq C\lambda_i \\ \lambda_i^2/C &\leq \lambda_i(k_*(\mathbf{X}, \mathbf{X})/n) \end{aligned} \tag{F.1}$$

both hold at the same time with probability $\geq 1/2$. Moreover, let $I(n) := \max\{m \in [n] \mid [m] \subseteq \mathcal{I}(n)\}$. Then, there exists a constant $c' > 0$ depending only on c, C such that for all $\rho \in [0, \infty)$ and $t \in (0, \infty]$, the following two bounds hold:

$$\begin{aligned} \mathbb{E}_D R_P(f_{\mathbf{X},\mathbf{y},t,\rho}) - R_P^* &\geq c' \sigma^2 \frac{1}{1 + (\rho + t^{-1})n^q} \cdot \frac{|\mathcal{I}(n)|}{n}, \\ \mathbb{E}_D R_P(f_{\mathbf{X},\mathbf{y},t,\rho}) - R_P^* &\geq c' \sigma^2 \min \left\{ \frac{(\rho + t^{-1})^{-2}}{n}, \frac{(\rho + t^{-1})^{-1/q}}{n}, \frac{I(n)}{n} \right\}. \end{aligned}$$

991 **Remark F.3.** Theorem F.2 provides two lower bounds, one for general “concentration sets” $\mathcal{I}(n)$ and
 992 one that applies if concentration holds for a sequence of “head eigenvalues” $\{1, \dots, I(n)\} \subseteq \mathcal{I}(n)$.
 993 If $I(n) \approx |\mathcal{I}(n)|$, the latter bound is stronger for larger regularization levels, and this bound would
 994 be particularly suitable for typical forms of relative concentration inequalities for kernel matrices.
 995 However, in this paper, we obtain concentration only for “middle eigenvalues” $\mathcal{I}(n) = \{i \mid \varepsilon n \leq i \leq$
 996 $(1 - \varepsilon)n\}$, and therefore we only use the first bound in the proof of Theorem 5. \blacktriangleleft

Proof of Theorem F.2. Step 1: Miscellaneous inequalities. For $x > 0$,

$$1 - e^{-x} = 1 - \frac{1}{e^x} \geq 1 - \frac{1}{x+1} = \frac{x}{x+1} = \frac{1}{1+x^{-1}}. \quad (\text{F.2})$$

Moreover, since $(1+a)^2 \leq (1+a)^2 + (1-a)^2 = 2 + 2a^2$, we have for $a \neq -1$:

$$\left(\frac{1}{1+a}\right)^2 \geq \frac{1}{2(1+a^2)}. \quad (\text{F.3})$$

Step 2: Applying the eigenvalue bound. Define

$$S_i(\mathbf{X}) := \frac{\lambda_i(k_*(\mathbf{X}, \mathbf{X})/n) (1 - e^{-2t(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + \rho)})^2}{(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + \rho)^2}.$$

By Proposition 4, we have

$$\mathbb{E}_D R_P(f_{\mathbf{X}, \mathbf{y}, t, \rho}) - R_P^* \geq \frac{\sigma^2}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{X}} S_i(\mathbf{X}) \geq \frac{\sigma^2}{n} \sum_{i \in \mathcal{I}(n)} \mathbb{E}_{\mathbf{X}} S_i(\mathbf{X}). \quad (\text{F.4})$$

Since $S_i(\mathbf{X})$ is almost surely positive, we can focus on the case where (F.1) and (F.2) hold, which is true with probability $\geq 1/2$ by assumption for $i \in \mathcal{I}(n)$. Hence,

$$\mathbb{E}_{\mathbf{X}} S_i(\mathbf{X}) \geq \frac{1}{2} \frac{\lambda_i^2/C \cdot (1 - e^{-2t(\lambda_i/C + \rho)})^2}{(C\lambda_i + \rho)^2} \stackrel{(\text{F.2})}{\geq} \frac{1}{2} \frac{\lambda_i^2/C}{(C\lambda_i + \rho)^2 (1 + (2t(\lambda_i/C + \rho))^{-1})^2}.$$

We can upper-bound the denominator, using $C \geq 1$, as

$$\begin{aligned} (C\lambda_i + \rho) \left(1 + \frac{1}{2t(\lambda_i/C + \rho)}\right) &\leq C\lambda_i + \rho + \frac{C^2\lambda_i + C\rho}{(\lambda_i + C\rho)t} \leq C\lambda_i + \rho + \frac{C^2\lambda_i + C^3\rho}{(\lambda_i + C\rho)t} \\ &\leq C^2(\lambda_i + \rho + t^{-1}), \end{aligned}$$

which yields

$$\begin{aligned} \mathbb{E}_{\mathbf{X}} S_i(\mathbf{X}) &\geq \frac{1}{2C^5} \frac{\lambda_i^2}{(\lambda_i + \rho + t^{-1})^2} = \frac{1}{2C^5} \frac{1}{\left(1 + \frac{\rho + t^{-1}}{\lambda_i}\right)^2} \stackrel{(\text{F.3})}{\geq} \frac{1}{4C^5} \frac{1}{1 + \left(\frac{\rho + t^{-1}}{\lambda_i}\right)^2} \\ &\geq \frac{1}{4C^5} \frac{1}{1 + \left(\frac{\rho + t^{-1}}{c} i^q\right)^2}. \end{aligned} \quad (\text{F.5})$$

Step 3: Analyzing the sum. We want to analyze the behavior of the sum

$$S(\beta) := \sum_{i \in \mathcal{I}(n)} \frac{1}{1 + (\beta i^q)^2}$$

997 for $\beta := \frac{\rho + t^{-1}}{c} > 0$. We first obtain the trivial bound

$$S(\beta) \geq |\mathcal{I}(n)| \frac{1}{1 + (\beta n^q)^2}.$$

Moreover, we can bound

$$S(\beta) \geq \sum_{i=1}^{I(n)} \frac{1}{1 + (\beta i^q)^2}$$

998 and distinguish three cases:

(a) If $\beta \geq 1$, we bound

$$S(\beta) \geq \sum_{i=1}^{I(n)} \frac{1}{2(\beta i^q)^2} \geq \frac{1}{2\beta^2}.$$

(b) If $\beta \in (I(n)^{-q}, 1)$, we observe that

$$J(\beta) := \lfloor \beta^{-1/q} \rfloor \geq \lceil \beta^{-1/q} \rceil - 1 \geq \frac{1}{2} \lceil \beta^{-1/q} \rceil \geq \frac{\beta^{-1/q}}{2}$$

and therefore

$$S(\beta) \geq \sum_{i=1}^{J(\beta)} \frac{1}{1 + (\beta i^q)^2} \geq \sum_{i=1}^{J(\beta)} \frac{1}{1 + 1} = \frac{J(\beta)}{2} \geq \frac{\beta^{-1/q}}{4}.$$

(c) If $\beta \in (0, I(n)^{-q}]$, we similarly find that

$$S(\beta) \geq \sum_{i=1}^{I(n)} \frac{1}{1 + 1} = \frac{I(n)}{2}.$$

999 Moreover, there is an absolute constant $c_1 > 0$ such that for any $\beta > 0$,

$$S(\beta) \geq c_1 \min\{\beta^{-2}, \beta^{-1/q}, I(n)\}, \quad (\text{F.6})$$

1000 because

- 1001 (a) $\beta^{-2} = \min\{\beta^{-2}, \beta^{-1/q}, I(n)\}$ for $\beta \geq 1$,
- 1002 (b) $\beta^{-1/q} = \min\{\beta^{-2}, \beta^{-1/q}, I(n)\}$ for $\beta \in (I(n)^{-q}, 1)$, and
- 1003 (c) $I(n) = \min\{\beta^{-2}, \beta^{-1/q}, I(n)\}$ for $\beta \in (0, I(n)^{-q}]$.

Step 4: Putting it together. Combining the trivial bound in Step 3 with Eq. (F.4) and Eq. (F.5), we obtain

$$\begin{aligned} \mathbb{E}_D R_P(f_{\mathbf{X}, \mathbf{y}, t, \rho}) - R_P^* &\geq \frac{\sigma^2}{n} \sum_{i \in \mathcal{I}(n)} \mathbb{E}_{\mathbf{X}} S_i(\mathbf{X}) \geq \frac{\sigma^2}{n} \cdot \frac{1}{4C^5} S(\beta) \\ &\geq c' \sigma^2 \frac{1}{1 + (\rho + t^{-1})n^q} \cdot \frac{|\mathcal{I}(n)|}{n} \end{aligned} \quad (\text{F.7})$$

1004 for a suitable constant $c' > 0$ depending only on c and C .

Moreover, from Eq. (F.6), we obtain

$$S(\beta) \geq \tilde{c}_1 \min\{\beta^{-2}, \beta^{-1/q}, I(n)\} \geq \tilde{c}'' \min\{(\rho + t^{-1})^{-2}, (\rho + t^{-1})^{-1/q}, I(n)\}$$

for a suitable constant $\tilde{c}'' > 0$ depending only on c . Again, (F.4) and (F.5) yield

$$\begin{aligned} \mathbb{E}_D R_P(f_{\mathbf{X}, \mathbf{y}, t, \rho}) - R_P^* &\geq \frac{\sigma^2}{n} \cdot \frac{1}{4C^5} S(\beta) \\ &\geq \frac{\tilde{c}''}{4C^5} \sigma^2 \min \left\{ \frac{(\rho + t^{-1})^{-2}}{n}, \frac{(\rho + t^{-1})^{-1/q}}{n}, \frac{I(n)}{n} \right\}. \quad \square \end{aligned}$$

1005 F.2 Equivalences of norms and eigenvalues

1006 Later, we will use concentration inequalities for kernel matrix eigenvalues proved for specific kernels,
1007 which we then want to transfer to other kernels with equivalent RKHSs. In this subsection, we show
1008 that this is possible.

Definition F.4 (C -equivalence of matrices and norms). Let $n \geq 1$ and let $\mathbf{K}, \tilde{\mathbf{K}} \in \mathbb{R}^{n \times n}$ be symmetric. For $C \geq 1$, we say that \mathbf{K} and $\tilde{\mathbf{K}}$ are C -equivalent if their ordered eigenvalues satisfy

$$C^{-1} \lambda_i(\mathbf{K}) \leq \lambda_i(\tilde{\mathbf{K}}) \leq C \lambda_i(\mathbf{K})$$

for all $i \in [n]$. Moreover, we say that two norms $\|\cdot\|_A, \|\cdot\|_B$ on a vector space V are C -equivalent if

$$C^{-1} \|\mathbf{v}\|_A \leq \|\mathbf{v}\|_B \leq C \|\mathbf{v}\|_A$$

1009 for all $\mathbf{v} \in V$. ◀

1010 **Lemma F.5.** Let $n \geq 1$ and let $\mathbf{K}, \tilde{\mathbf{K}} \in \mathbb{R}^{n \times n}$ be symmetric. Then, \mathbf{K} and $\tilde{\mathbf{K}}$ are C -equivalent iff
 1011 the Moore-Penrose pseudoinverses \mathbf{K}^+ and $\tilde{\mathbf{K}}^+$ are C -equivalent.

1012 *Proof.* This follows from the fact that if \mathbf{K} has eigenvalues $\lambda_1, \dots, \lambda_n$, then \mathbf{K}^+ has eigenvalues
 1013 $1/\lambda_1, \dots, 1/\lambda_n$, where we define $1/0 := 0$. (A detailed proof would be a bit technical due to the
 1014 sorting of eigenvalues.) \square

1015 **Lemma F.6.** Let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a kernel on a set \mathcal{X} . Then, for any $\mathbf{y} \in \mathbb{R}^n$,

$$\mathbf{y}^\top k(\mathbf{X}, \mathbf{X})^+ \mathbf{y} = \|f_{k, \mathbf{y}}^*\|_{\mathcal{H}_k}^2,$$

where \mathcal{H}_k is the RKHS associated with k and $f_{k, \mathbf{y}}^*$ is the minimum-norm regression solution

$$f_{k, \mathbf{y}}^* := \operatorname{argmin}_{f \in B} \|f\|_{\mathcal{H}_k}^2,$$

$$B := \left\{ f \in \mathcal{H}_k \mid \sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2 = \inf_{\tilde{f} \in \mathcal{H}_k} \sum_{i=1}^n (\tilde{f}(\mathbf{x}_i) - y_i)^2 \right\}.$$

Proof. It is well-known that $f_{k, \mathbf{y}}^*(\mathbf{x}) = \sum_{i=1}^n \alpha_i k(\mathbf{x}, \mathbf{x}_i)$, where $\boldsymbol{\alpha} := \mathbf{K}^+ \mathbf{y}$ (see e.g. Rangamani et al., 2023). We then have

$$\begin{aligned} \|f_{k, \mathbf{y}}^*\|_{\mathcal{H}_k}^2 &= \left\langle \sum_{i=1}^n \alpha_i k(\mathbf{x}_i, \cdot), \sum_{j=1}^n \alpha_j k(\mathbf{x}_j, \cdot) \right\rangle_{\mathcal{H}_k} = \sum_{i=1}^n \sum_{j=1}^n \alpha_i \mathbf{K}_{ij} \alpha_j \\ &= \mathbf{y}^\top \mathbf{K}^+ \mathbf{K} \mathbf{K}^+ \mathbf{y} = \mathbf{y}^\top \mathbf{K}^+ \mathbf{y}, \end{aligned}$$

1016 where the last step follows from a standard identity for the Moore-Penrose pseudoinverse (see e.g.
 1017 Section 1.1.1 in Wang et al., 2018). \square

1018 **Proposition F.7** (Equivalent kernels have equivalent kernel matrices). Let $k, \tilde{k} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be
 1019 kernels such that their RKHSs are equal as sets and the corresponding RKHS-norms are C -equivalent
 1020 as defined in Definition F.4. Then, for any $n \geq 1$ and any $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}$, the corresponding kernel
 1021 matrices $k(\mathbf{X}, \mathbf{X}), \tilde{k}(\mathbf{X}, \mathbf{X})$ are C^2 -equivalent.

Proof. Let $i \in [n]$. For $\mathbf{y} \in \mathbb{R}^n$ we have, using the notation of Lemma F.6:

$$\mathbf{y}^\top k(\mathbf{X}, \mathbf{X})^+ \mathbf{y} = \|f_{k, \mathbf{y}}^*\|_{\mathcal{H}_k}^2 \geq C^{-2} \|f_{k, \mathbf{y}}^*\|_{\mathcal{H}_{\tilde{k}}}^2 \geq C^{-2} \|f_{\tilde{k}, \mathbf{y}}^*\|_{\mathcal{H}_{\tilde{k}}}^2 = C^{-2} \mathbf{y}^\top \tilde{k}(\mathbf{X}, \mathbf{X})^+ \mathbf{y}.$$

Now, by the Courant-Fischer-Weyl theorem,

$$\begin{aligned} \lambda_i(k(\mathbf{X}, \mathbf{X})^+) &= \sup_{V: \dim V=i} \inf_{y \in V: \|y\|_2=1} \mathbf{y}^\top k(\mathbf{X}, \mathbf{X})^+ \mathbf{y} \\ &\geq C^{-2} \sup_{V: \dim V=i} \inf_{y \in V: \|y\|_2=1} \mathbf{y}^\top \tilde{k}(\mathbf{X}, \mathbf{X})^+ \mathbf{y} \\ &= C^{-2} \lambda_i(\tilde{k}(\mathbf{X}, \mathbf{X})^+). \end{aligned}$$

1022 By switching the roles of k and \tilde{k} , we obtain that $k(\mathbf{X}, \mathbf{X})^+$ and $\tilde{k}(\mathbf{X}, \mathbf{X})^+$ are C^2 -equivalent. By
 1023 Lemma F.5 $k(\mathbf{X}, \mathbf{X})$ and $\tilde{k}(\mathbf{X}, \mathbf{X})$ are then also C^2 -equivalent. \square

1024 To prove Theorem 5 for arbitrary input distributions $P_{\mathcal{X}}$ with lower and upper bounded densities, we
 1025 need the following theorem investigating the corresponding eigenvalues of the integral operator.

1026 **Lemma F.8** (Integral operators for equivalent densities have equivalent eigenvalues). Let $k : \mathcal{X} \times \mathcal{X} \rightarrow$
 1027 \mathbb{R} be a kernel and let μ, ν be finite measures on \mathcal{X} whose support is \mathcal{X} such that ν has an lower and
 1028 upper bounded density w.r.t. μ . Then, $\lambda_i(T_{k, \nu}) = \Theta(\lambda_i(T_{k, \mu}))$.

1029 *Proof.* Let p be such an upper bounded density, that is, $d\nu = p d\mu$ and there exist $c, C > 0$ such that
 1030 $c \leq p(\mathbf{x}) \leq C$ for all $\mathbf{x} \in \mathcal{X}$. For $f \in L_2(\nu)$, we have

$$\|p \cdot f\|_{L_2(\mu)}^2 = \int f^2 p^2 d\mu \leq C \int f^2 p d\mu = C \int f^2 d\nu = C \|f\|_{L_2(\nu)}^2.$$

1031 Hence, the linear operator

$$A : L_2(\nu) \rightarrow L_2(\mu), f \mapsto p \cdot f$$

1032 is well-defined and continuous. It is also easily verified that A is bijective. Moreover, we have

$$\langle Af, Af \rangle_{L_2(\mu)} = \int f^2 p^2 d\mu \geq c \int f^2 p d\mu = c \int f^2 d\nu = c \langle f, f \rangle_{L_2(\nu)} .$$

1033 and

$$\langle f, T_{k,\nu} f \rangle_{L_2(\nu)} = \int \int p(\mathbf{x}) f(\mathbf{x}) k(\mathbf{x}, \mathbf{x}') f(\mathbf{x}') p(\mathbf{x}') d\mu(\mathbf{x}) d\mu(\mathbf{x}') = \langle Af, T_{k,\mu} Af \rangle_{L_2(\mu)} .$$

Since $T_{k,\mu}$ and $T_{k,\nu}$ are compact, self-adjoint, and positive, we can use the Courant-Fischer minmax principle for operators (see e.g. Bell, 2014) to obtain

$$\begin{aligned} \lambda_i(T_{k,\nu}) &= \max_{\substack{V \subseteq L_2(\nu) \\ \dim V = i}} \min_{f \in V \setminus \{0\}} \frac{\langle f, T_{k,\nu} f \rangle_{L_2(\nu)}}{\langle f, f \rangle_{L_2(\nu)}} \\ &\geq c \max_{\substack{V \subseteq L_2(\nu) \\ \dim V = i}} \min_{f \in V \setminus \{0\}} \frac{\langle Af, T_{k,\mu} Af \rangle_{L_2(\mu)}}{\langle Af, Af \rangle_{L_2(\mu)}} \\ &= c \max_{\substack{\tilde{V} \subseteq L_2(\nu) \\ \dim \tilde{V} = i}} \min_{g \in \tilde{V} \setminus \{0\}} \frac{\langle g, T_{k,\mu} g \rangle_{L_2(\mu)}}{\langle g, g \rangle_{L_2(\mu)}} \\ &= c \lambda_i(T_{k,\mu}) . \end{aligned}$$

1034 Here, we have used that since A is bijective, the subspaces AV for $\dim(V) = i$ are exactly the
1035 i -dimensional subspaces of $L_2(\mu)$. Our calculation above shows that $\lambda_i(T_{k,\mu}) \leq O(\lambda_i(T_{k,\nu}))$. Since
1036 $d\mu = \frac{1}{p} d\nu$ with the lower and upper bounded density $1/p$, we can reverse the roles of ν and μ to
1037 also obtain $\lambda_i(T_{k,\nu}) \leq O(\lambda_i(T_{k,\mu}))$, which proves the claim. \square

1038 **Lemma F.9.** Let \mathcal{H}_1 and \mathcal{H}_2 be two RKHSs with $\mathcal{H}_1 \subset \mathcal{H}_2$. Then there exists a constant $C > 0$ such
1039 that $\|f\|_{\mathcal{H}_2} \leq C \|f\|_{\mathcal{H}_1}$.

1040 *Proof.* Let $I : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ be the inclusion map, i.e. $I_h := h$ for all $h \in \mathcal{H}_1$. Obviously, I is linear
1041 and we need to show that I is bounded. To this end, let $(h_n)_{n \geq 1} \subset \mathcal{H}_1$ be a sequence such that
1042 there exist $h \in \mathcal{H}_1$ and $g \in \mathcal{H}_2$ with $h_n \rightarrow h$ in \mathcal{H}_1 and $Ih_n \rightarrow g$ in \mathcal{H}_2 . This implies $h_n \rightarrow h$
1043 pointwise and $h_n = Ih_n \rightarrow g$ pointwise, which in turn gives $h = g$. The closed graph theorem, see
1044 e.g. (Megginson, 1998, Theorem 1.6.11), then shows that I is bounded. \square

1045 F.3 Kernel matrix eigenvalue bounds

1046 For upper bounds on the eigenvalues of kernel matrices, we use the following result:

Proposition F.10 (Kernel matrix eigenvalue upper bound in expectation). For $m \geq 1$, we have

$$\mathbb{E}_{\mathbf{X}} \sum_{i=m}^n \lambda_i(k(\mathbf{X}, \mathbf{X})/n) \leq \sum_{i=m}^{\infty} \lambda_i(T_k) . \quad (\text{F.8})$$

1047 *Proof.* Theorem 7.29 in Steinwart and Christmann (2008) shows that

$$\mathbb{E}_{D \sim \mu^n} \sum_{i=m}^{\infty} \lambda_i(T_{k,D}) \leq \sum_{i=m}^{\infty} \lambda_i(T_{k,\mu}) , \quad (\text{F.9})$$

where $T_{k,\mu} : L_2(\mu) \rightarrow L_2(\mu), f \mapsto \int k(x, \cdot) f(x) d\mu(x)$ is the integral operator corresponding to
the measure μ and $T_{k,D}$ is the corresponding discrete version thereof. We set $\mu := P_{\mathbf{X}}$ and need to
show that $k(\mathbf{X}, \mathbf{X})/n$ has the same eigenvalues as $T_{k,D}$ if D and \mathcal{X} contain the same data points
 $\mathbf{x}_1, \dots, \mathbf{x}_n$. Consider a fixed D . Then, we can write $T_{k,D}(f) = n^{-1} ABf$, where

$$A : \mathbb{R}^n \rightarrow L_2(D), \mathbf{v} \mapsto \sum_{i=1}^n v_i k(\mathbf{x}_i, \cdot)$$

$$B : L_2(D) \rightarrow \mathbb{R}^n, f \mapsto (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))^\top .$$

Then, $k(\mathbf{X}, \mathbf{X})/n$ is the matrix representation of $n^{-1}BA$ with respect to the standard basis of \mathbb{R}^n . But AB and BA have the same non-zero eigenvalues, which means that

$$\sum_{i=m}^n \lambda_i(k(\mathbf{X}, \mathbf{X})/n) = \sum_{i=m}^{\infty} \lambda_i(T_{k,D}) ,$$

1048 from which the claim follows. \square

1049 To obtain a lower bound, we want to leverage the lower bound by Buchholz (2022) for a certain radial
1050 basis function kernel with data generated from an open subset of \mathbb{R}^d . However, we want to consider
1051 different kernels and distributions on the whole sphere. The following theorem bridges the gap by
1052 going to subsets of the data on a sphere cap, projecting them to \mathbb{R}^d , and using the kernel equivalence
1053 results from Appendix F.2:

Theorem F.11 (Kernel matrix eigenvalue lower bound for Sobolev kernels on the sphere). *Let k be a kernel on \mathbb{S}^d such that its RKHS \mathcal{H}_k is equivalent to a Sobolev space $H^s(\mathbb{S}^d)$ with smoothness $s > d/2$. Moreover, let P_X be a probability distribution on \mathbb{S}^d with lower and upper bounded density. Let the rows of $\mathbf{X} \in \mathbb{R}^{n \times d}$ are drawn independently from P_X . Then, for $\varepsilon \in (0, 1/20)$, there exists a constant $c > 0$ and $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$,*

$$\lambda_m(k(\mathbf{X}, \mathbf{X})/n) \geq cn^{-2s/d}$$

1054 holds with probability $\geq 4/5$ for all $m \in \mathbb{N}$ with $1 \leq m \leq (1 - 11\varepsilon)n$.

1055 *Proof.* We can choose a suitably large sphere cap T such that $P_X(T) \geq 1 - \varepsilon$. Define the conditional
1056 distribution $P_T(\cdot) := P_X(\cdot|T)$. Out of the points $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, we can consider the submatrix
1057 $\mathbf{X}_T = (\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_N})^\top$ of the points lying in T . Conditioned on N , these points are i.i.d. samples
1058 from P_T . Moreover, by applying Markov's inequality to a Bernoulli distribution, we obtain $N \geq$
1059 $(1 - 10\varepsilon)n$ with probability $\geq 9/10$. We fix a value of $N \geq (1 - 10\varepsilon)n$ in the following and condition
1060 on it.

1061 We denote the centered unit ball in \mathbb{R}^d by $B_1(\mathbb{R}^d)$. Using a construction as in Lemma E.1, we can
1062 transport k and P_T from T to the unit ball $B_1(\mathbb{R}^d)$ using a rescaled stereographic projection feature
1063 map ϕ , such that we obtain a kernel k_ϕ and a distribution $P_\phi = (P_T)_\phi$ on $B_1(\mathbb{R}^d)$ that generate the
1064 same distribution of kernel matrices as k with P_T , and such that $\mathcal{H}_{k_\phi} \cong H^s(B_1(\mathbb{R}^d))$. The rows
1065 of $\mathbf{X}_\phi := \phi(\mathbf{X}_T)$ are i.i.d. samples from P_ϕ . Moreover, we know that P_ϕ has an lower and upper
1066 bounded density w.r.t. the Lebesgue measure on $B_1(\mathbb{R}^d)$.

1067 In order to apply the results from Buchholz (2022), we define a translation-invariant reference kernel
1068 on \mathbb{R}^d through the Fourier transform

$$\hat{k}_{\text{ref}}(\xi) = (1 + |\xi|^2)^{-2s} ,$$

1069 see Eq. (3) in Buchholz (2022). The RKHS of k_{ref} on \mathbb{R}^d is equivalent to the Sobolev space $H^s(\mathbb{R}^d)$.
1070 Therefore, the RKHS of $k_{\text{ref}}|_{B_1(\mathbb{R}^d), B_1(\mathbb{R}^d)}$ is $H^s(B_1(\mathbb{R}^d))$, cf. the remarks in Appendix B.1 and
1071 Lemma F.9.

1072 Now, let $1 \leq m \leq (1 - 11\varepsilon)n$, which implies

$$1 \leq m \leq (1 - 11\varepsilon)n \leq (1 - \varepsilon)(1 - 10\varepsilon)n \leq (1 - \varepsilon)N .$$

We apply Theorem 12 by Buchholz (2022) with bandwidth $\gamma = 1$ and $\alpha = 2s$ to λ_m and obtain with probability at least $1 - 2/N$:

$$\begin{aligned} \lambda_m(k_{\text{ref}}(\mathbf{X}_\phi, \mathbf{X}_\phi))^{-1} &\leq c_3 \left(\frac{N^{2(\alpha-d)/d}}{(N-m)^{(\alpha-d)/d}} + 1 \right) \leq c_3 \left(\frac{N^{2(\alpha-d)/d}}{(\varepsilon N)^{(\alpha-d)/d}} + 1 \right) \\ &\leq c_4 (n^{\alpha/d-1} + 1) \end{aligned}$$

as long as N is large enough such that $(1 - \varepsilon)N < N - 32 \ln(N)$, which is the case if n is large enough. Here, the constant c_3 from Buchholz (2022) does not depend on N or m , but only on α, d ,

and the upper and lower bounds on the density, which in our case depend on ε through the choice of T . Since $\alpha = 2s > d$, we have $n^{\alpha/d-1} > 1$ and therefore

$$\lambda_m(k_{\text{ref}}(\mathbf{X}_\phi, \mathbf{X}_\phi)/n) \geq c_5 n^{-\alpha/d} = c_5 n^{-2s/d}.$$

Now, we want to translate this to the kernel k . Since the RKHSs of k_ϕ and k_{ref} on $B_1(\mathbb{R}^d)$ are both equivalent to $H^s(B_1(\mathbb{R}^d))$, the kernels themselves are C -equivalent for some constant $C \geq 1$ as defined in Definition F.4. Therefore, Proposition F.7 shows that the corresponding kernel matrices are C^2 -equivalent, which implies

$$\lambda_m(k_{*,\phi}(\mathbf{X}_\phi, \mathbf{X}_\phi)/n) \geq c_5 C^{-2} n^{-2s/d}.$$

By Cauchy's interlacing theorem, we therefore have

$$\lambda_m(k_*(\mathbf{X}, \mathbf{X})/n) \geq \lambda_m(k_*(\mathbf{X}_T, \mathbf{X}_T)/n) = \lambda_m(k_{*,\phi}(\mathbf{X}_\phi, \mathbf{X}_\phi)/n) \geq c_5 C^{-2} n^{-2s/d}.$$

Denoting the event where $\lambda_m(k_*(\mathbf{X}, \mathbf{X})/n) \geq c_5 C^{-2} n^{-2s/d}$ by A , we thus have

$$\begin{aligned} P(A) &= P(A|N \geq (1-10\varepsilon)n)P(N \geq (1-10\varepsilon)n) \geq \frac{9}{10}P(A|N \geq (1-10\varepsilon)n) \\ &= \frac{9}{10} \sum_{\hat{N}=\lceil(1-10\varepsilon)n\rceil}^n P(N = \hat{N}|N \geq (1-10\varepsilon)n)P(A|N = \hat{N}) \\ &\geq \frac{9}{10} \sum_{\hat{N}=\lceil(1-10\varepsilon)n\rceil}^n P(N = \hat{N}|N \geq (1-10\varepsilon)n)(1-2/\hat{N}) \\ &\geq \frac{9}{10} \left(1 - \frac{2}{(1-10\varepsilon)n}\right) \sum_{\hat{N}=\lceil(1-10\varepsilon)n\rceil}^n P(N = \hat{N}|N \geq (1-10\varepsilon)n) \\ &= \frac{9}{10} \left(1 - \frac{2}{(1-10\varepsilon)n}\right) \geq \frac{4}{5}, \end{aligned}$$

1073 where the last step holds for sufficiently large n . □

1074 **F.4 Spectral lower bound for dot-product kernels on the sphere**

1075 An application of the spectral generalization bound in Proposition 4 requires a lower bound on
 1076 eigenvalues of the kernel matrix $k_*(\mathbf{X}, \mathbf{X})$. To achieve this, we need to understand the properties of
 1077 the convolution kernel k_* . Since the eigenvalues of T_{k_*, P_X} are the squared eigenvalues of T_{k, P_X} , one
 1078 might hope that if \mathcal{H}_k is equivalent to a Sobolev space H^s , then \mathcal{H}_{k_*} is equivalent to a Sobolev space
 1079 H^{2s} . Unfortunately, this is not the case in general, as \mathcal{H}_{k_*} might be a smaller space that involves
 1080 additional boundary conditions (Schaback, 2018). However, perhaps since the sphere is a manifold
 1081 without boundary, the desired characterization of \mathcal{H}_{k_*} holds for dot-product kernels on the sphere:

1082 **Lemma F.12** (RKHS of convolution kernels). *Let k be a dot-product kernel on \mathbb{S}^d such that its RKHS*
 1083 *\mathcal{H}_k is equivalent to a Sobolev space $H^s(\mathbb{S}^d)$ with smoothness $s > d/2$, and let P_X be a distribution*
 1084 *on \mathbb{S}^d with lower and upper bounded density. Then, the RKHS \mathcal{H}_{k_*} of the kernel*

$$k_* : \mathbb{S}^d \times \mathbb{S}^d \rightarrow \mathbb{R}, k_*(\mathbf{x}, \mathbf{x}') := \int k(\mathbf{x}, \mathbf{x}'')k(\mathbf{x}'', \mathbf{x}') dP_X(\mathbf{x}'')$$

1085 *is equivalent to the Sobolev space $H^{2s}(\mathbb{S}^d)$.*

Proof. Define

$$k_{*,\text{unif}}(\mathbf{x}, \mathbf{x}') = \int k(\mathbf{x}, \mathbf{x}'')k(\mathbf{x}'', \mathbf{x}') d\mathcal{U}(\mathbb{S}^d)(\mathbf{x}'').$$

For the corresponding integral operator, we have

$$T_{k_{*,\text{unif}}, \mathcal{U}(\mathbb{S}^d)} = T_{k, \mathcal{U}(\mathbb{S}^d)}^2.$$

This means that the corresponding eigenvalues are the squares of the eigenvalues of the corresponding integral operator of k . Especially, we obtain the Mercer representations

$$k(\mathbf{x}, \mathbf{x}') = \sum_{l=0}^{\infty} \mu_l \sum_{i=1}^{N_{l,d}} Y_{l,i}(\mathbf{x}) Y_{l,i}(\mathbf{x}'),$$

$$k_{*,\text{unif}}(\mathbf{x}, \mathbf{x}') = \sum_{l=0}^{\infty} \mu_l^2 \sum_{i=1}^{N_{l,d}} Y_{l,i}(\mathbf{x}) Y_{l,i}(\mathbf{x}'),$$

1086 where Lemma B.1 yields $\mu_l = \Theta((l+1)^{-2s})$, hence $\mu_l^2 = \Theta((l+1)^{-4s})$ and hence $\mathcal{H}_{k_{*,\text{unif}}} \cong$
1087 $H^{2s}(\mathbb{S}^d)$.

1088 Next, we show the equality of the ranges of the integral operators:

$$R(T_{k,\mathcal{U}(\mathbb{S}^d)}) = R(T_{k,P_X}).$$

1089 Let p_X be a density of P_X w.r.t. the uniform distribution $\mathcal{U}(\mathbb{S}^d)$. If $f \in R(T_{k,\mathcal{U}(\mathbb{S}^d)})$, there exists
1090 $g \in L_2(\mathcal{U}(\mathbb{S}^d))$ with $f = T_{k,\mathcal{U}(\mathbb{S}^d)}g$. But then, since p_X is lower bounded, we have $g/p_X \in L_2(P_X)$
1091 and therefore

$$f = T_{k,P_X}(g/p_X) \in R(T_{k,P_X}).$$

1092 An analogous argument shows that $R(T_{k,P_X}) \subseteq R(T_{k,\mathcal{U}(\mathbb{S}^d)})$ since p_X is upper bounded.

The equality of the ranges yields for the RKHSs (as sets)

$$\mathcal{H}_{k_{*,\text{unif}}} = R(T_{k,\mathcal{U}(\mathbb{S}^d)}) = R(T_{k,P_X}) = \mathcal{H}_{k_*},$$

1093 Applying Lemma F.9 twice then shows $\mathcal{H}_{k_*} \cong H^{2s}(\mathbb{S}^d)$. □

1094 **Theorem 5 (Inconsistency for Sobolev dot-product kernels on the sphere).** *Let k be a dot-product*
1095 *kernel on \mathbb{S}^d , i.e., a kernel of the form $k(\mathbf{x}, \mathbf{x}') = \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle)$, such that its RKHS \mathcal{H}_k is equivalent*
1096 *to a Sobolev space $H^s(\mathbb{S}^d)$, $s > d/2$. Moreover, let P be a distribution on $\mathbb{S}^d \times \mathbb{R}$ such that*
1097 *P_X has a lower and upper bounded density w.r.t. the uniform distribution $\mathcal{U}(\mathbb{S}^d)$, and such that*
1098 *$\text{Var}(y|\mathbf{x}) \geq \sigma^2 > 0$ for P_X -almost all $\mathbf{x} \in \mathbb{S}^d$. Then, for every $C > 0$, there exists $c > 0$ independent*
1099 *of σ^2 such that for all $n \geq 1$, $t \in (C^{-1}n^{2s/d}, \infty]$, and $\rho \in [0, Cn^{-2s/d})$, the expected excess risk*
1100 *satisfies*

$$\mathbb{E}_D \mathcal{R}_P(f_{t,\rho}) - \mathcal{R}_P^* \geq c\sigma^2 > 0.$$

1101 *Proof. Step 0: Preparation.* Since the Sobolev space $H^{2s}(\mathbb{S}^d)$ is dense in the space of continuous
1102 functions $\mathbb{S}^d \rightarrow \mathbb{R}$, the kernel k is universal. Applying (Steinwart and Christmann, 2008, Corollary
1103 5.29 and Corollary 5.34) for the least squares loss thus shows that k is strictly positive definite. If
1104 we have mutually distinct $\mathbf{x}_1, \dots, \mathbf{x}_n$, the corresponding Gram matrix $k((\mathbf{x}_i, \mathbf{x}_j))_{i,j=1}^n$ is therefore
1105 invertible. Now, our assumptions on P guarantee that \mathbf{X} consists almost surely of mutually distinct
1106 observations, and therefore $k(\mathbf{X}, \mathbf{X})$ is almost surely invertible.

By Proposition 4, we know that

$$\begin{aligned} \mathbb{E}_D \mathcal{R}_P(f_{\mathbf{X}, \mathbf{y}, t, \rho}) - \mathcal{R}_P^* &\geq \frac{\sigma^2}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{X}} \frac{\lambda_i(k_*(\mathbf{X}, \mathbf{X})/n) \left(1 - e^{-2t(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + \rho)}\right)^2}{(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + \rho)^2} \\ &\geq \frac{\sigma^2}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{X}} \frac{\lambda_i(k_*(\mathbf{X}, \mathbf{X})/n) \left(1 - e^{-2C^{-1}n^{2s/d}(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + 0)}\right)^2}{(\lambda_i(k(\mathbf{X}, \mathbf{X})/n) + Cn^{-2s/d})^2} \\ &\geq c_n \sigma^2 \end{aligned}$$

1107 for a suitable constant $c_n > 0$ depending on n but not on σ^2, t, ρ , since the kernel matrix eigenvalues
1108 are nonzero almost surely. It is therefore sufficient to show the desired statement (with c independent
1109 of n, σ^2, t, ρ) for sufficiently large n .

1110 In the following, we assume $n \geq 40$ and set $\varepsilon := 1/100$.

Step 1: Eigenvalue decay for the integral operator. From Lemma B.1, we know that

$$\lambda_i(T_{k,\mathcal{U}(\mathbb{S}^d)}) = \Theta(i^{-2s/d}).$$

Therefore, by Lemma F.8, we know that

$$\lambda_i(T_{k,P_X}) = \Theta(i^{-2s/d}).$$

1111 **Step 2: Eigenvalue upper bound.** Next, we want to upper-bound suitable eigenvalues of the form
 1112 $\lambda_i(k(\mathbf{X}, \mathbf{X})/n)$ using Proposition F.10. Using Step 1, we derive

$$\sum_{i=m}^{\infty} \lambda_i(T_{k,P_X}) \leq C_1 \sum_{i=m}^{\infty} i^{-2s/d} \leq C_2 \int_m^{\infty} x^{-2s/d} dx = C_3 m^{1-2s/d}$$

1113 with constants independent of $m \geq 1$. For sufficiently large n , we can choose $m \in \mathbb{N}_{\geq 1}$ such that
 1114 $\varepsilon n \leq m \leq 2\varepsilon n$. Then, Proposition F.10 yields

$$\mathbb{E}_{\mathbf{X}} \sum_{i=m}^n \lambda_i(k(\mathbf{X}, \mathbf{X})/n) \leq \sum_{i=m}^{\infty} \lambda_i(T_k) \leq C_3 m^{1-2s/d} \leq C_4 n^{1-2s/d}.$$

1115 Since $\mathbb{E}_{\mathbf{X}} \lambda_i(k(\mathbf{X}, \mathbf{X})/n)$ is decreasing with i , we have for $i \geq 4\varepsilon n \geq 2m$:

$$\mathbb{E}_{\mathbf{X}} \lambda_i(k(\mathbf{X}, \mathbf{X})/n) \leq C_4 n^{1-2s/d}/m \leq C_5 n^{-2s/d} \leq C_6 \lambda_i(T_{k,P_X}).$$

Step 3: Eigenvalue lower bounds. From Lemma F.12, we know that $\mathcal{H}_{k_*} \cong H^{2s}(\mathbb{S}^d)$. Therefore, we can apply Lemma F.12 to both k and k_* and obtain for sufficiently large n and suitable constants $c_1, c_2 > 0$ that

$$\begin{aligned} \lambda_i(k(\mathbf{X}, \mathbf{X})/n) &\geq c_1 n^{-2s/d} \\ \lambda_i(k_*(\mathbf{X}, \mathbf{X})/n) &\geq c_2 n^{-4s/d} \end{aligned}$$

1116 individually hold with probability $\geq 4/5$ for all $i \in \mathbb{N}$ with $1 \leq i \leq (1 - 11\varepsilon)n$. By the union bound,
 1117 both bounds hold at the same time with probability $\geq 3/5$.

Step 4: Final result. Now, using the value of m from Step 2, consider an index i with $2m \leq i \leq (1 - 11\varepsilon)n$. Since $2m \leq 4\varepsilon n$ and $\varepsilon = 1/100$, there are at least $n/2$ such indices. By combining Step 3 and Step 1, we have

$$\begin{aligned} \lambda_i(k(\mathbf{X}, \mathbf{X})/n) &\geq c_3 \lambda_i(T_{k,P_X}) \\ \lambda_i(k_*(\mathbf{X}, \mathbf{X})/n) &\geq c_4 \lambda_i(T_{k,P_X})^2 \end{aligned}$$

1118 with probability $\geq 3/5$. By applying Markov's inequality to Step 2, we obtain

$$\lambda_i(k(\mathbf{X}, \mathbf{X})/n) \leq 10C_6 \lambda_i(T_{k,P_X})$$

with probability $\geq 9/10$. Therefore, by the union bound, all three inequalities hold simultaneously with probability $\geq 1/2$. Moreover, for $q = 2s/d$, we have $\lambda_i(T_{k,P_X}) \geq c_5 i^{-q}$ by Step 1. We can thus apply the first lower bound from Theorem F.2 to obtain

$$\begin{aligned} \mathbb{E}_D R_P(f_{\mathbf{X}, \mathbf{y}, t, \rho}) - R_P^* &\geq c' \sigma^2 \frac{1}{1 + (\rho + t^{-1})n^{2s/d}} \cdot \frac{|\mathcal{I}(n)|}{n} \\ &\geq c' \sigma^2 \frac{1}{1 + (Cn^{-2s/d} + Cn^{-2s/d})n^{2s/d}} \cdot \frac{n/2}{n} \\ &= \frac{c'}{2 + 2C} \sigma^2. \quad \square \end{aligned}$$

1119 G Proof of Theorem 7

1120 Here we denote the solution of kernel ridge regression on D with the kernel function k and regular-
 1121 ization parameter $\rho > 0$ as

$$\hat{f}_\rho^k(\mathbf{x}) = k(\mathbf{x}, \mathbf{X}) (k(\mathbf{X}, \mathbf{X}) + \rho \mathbf{I})^{-1} \mathbf{y},$$

1122 and write $\hat{f}_0^k(\mathbf{x}) = k(\mathbf{x}, \mathbf{X}) k(\mathbf{X}, \mathbf{X})^+ \mathbf{y}$ for the minimum-norm interpolant in the RKHS of k .

1123 While Theorem 1 states that overfitting kernel ridge regression using Sobolev kernels is always
 1124 inconsistent as long as the derivatives remain bounded by the derivatives of the minimum-norm

1125 interpolant of the fixed kernel (Assumption (N)), here we show that consistency over a large class
 1126 of distributions is achievable by designing a kernel sequence, which can have Sobolev RKHS, that
 1127 consists of a smooth component for generalization and a spiky component for interpolation.

1128 Recall that \tilde{k} denotes any universal kernel function for the smooth component, and \check{k}_γ denotes
 1129 the kernel function of the spiky component with bandwidth γ . Then we define the ρ -regularized
 1130 *spiky-smooth kernel with spike bandwidth* γ as

$$k_{\rho,\gamma}(\mathbf{x}, \mathbf{x}') = \tilde{k}(\mathbf{x}, \mathbf{x}') + \rho \cdot \check{k}_\gamma(\mathbf{x}, \mathbf{x}').$$

1131 Let $B_t(\mathbf{x}) := \{\mathbf{y} \in \mathbb{R}^d \mid \|\mathbf{x} - \mathbf{y}\| \leq t\}$ denote the Euclidean ball of radius $t \geq 0$ around $\mathbf{x} \in \mathbb{R}^d$.

1132 (D2) There exists a constant $\beta_X > 0$ and a continuous function $\phi : [0, \infty) \rightarrow [0, 1]$ with $\phi(0) = 0$
 1133 such that $P_X(B_t(\mathbf{x})) \leq \phi(t) = O(t^{\beta_X})$ for all $\mathbf{x} \in \Omega$ and all $t \geq 0$.

1134 The kernel \check{k}_γ of the spiky component should fulfill the following weak assumption on its decay
 1135 behaviour. For example, Laplace, Matérn, and Gaussian kernels all fulfill Assumption (SK).

1136 (SK) There exists a function $\varepsilon : (0, \infty) \times [0, \infty) \rightarrow [0, 1]$ such that for any bandwidth $\gamma > 0$ and
 1137 any $\delta > 0$ it holds that

- 1138 (i) $\varepsilon(\gamma, 0) = 1$,
- 1139 (ii) $\varepsilon(\gamma, \delta)$ is monotonically increasing in γ ,
- 1140 (iii) For all $\mathbf{x}, \mathbf{y} \in \Omega$, if $\|\mathbf{x} - \mathbf{y}\| \geq \delta$ then $|\check{k}_\gamma(\mathbf{x}, \mathbf{y})| \leq \varepsilon(\gamma, \delta)$,
- 1141 (iv) For any rates $\beta_X, \beta_k > 0$ there exists a rate $\beta_\gamma > 0$ such that, if $\delta_n = \Omega(n^{-\beta_X})$ and
 1142 $\gamma_n = O(n^{-\beta_\gamma})$, then $\varepsilon(\gamma_n, \delta_n) = O(n^{-\beta_k})$.

1143 **Theorem G.1 (Consistency of spiky-smooth ridgeless kernel regression).** Assume that the
 1144 training set D consists of n i.i.d. pairs $(\mathbf{x}, y) \sim P$ such that the marginal P_X fulfills (D2) and
 1145 $\mathbb{E}y^2 < \infty$. Let the kernel components satisfy:

- 1146 • \tilde{k} denotes an arbitrary universal kernel, and $\rho_n \rightarrow 0$ and $n\rho_n^4 \rightarrow \infty$.
- 1147 • \check{k}_{γ_n} denotes a kernel function that fulfills Assumption (SK) with a sequence of positive
 1148 bandwidths (γ_n) fulfilling $\gamma_n = O(\exp(-\beta n))$ for some arbitrary $\beta > 0$.

1149 Then the minimum-norm interpolant of the ρ_n -regularized spiky-smooth kernel sequence $k_n := k_{\rho_n, \gamma_n}$
 1150 is consistent for P .

1151 **Remark G.2 (Spike bandwidth scaling).** Under stronger assumptions on ϕ and ε in assumptions
 1152 (D2) and (SK), the spike bandwidths γ_n can be chosen to converge to 0 at a much slower rate. For
 1153 example, if we choose \check{k}_γ to be the Laplace kernel, choosing bandwidths $0 < \gamma_n \leq \frac{\delta}{\beta \ln n}$ yields, for
 1154 separated points $\|\mathbf{x} - \mathbf{y}\| \geq \delta$,

$$\check{k}_{\gamma_n}(\mathbf{x}, \mathbf{y}) \leq \exp\left(-\frac{\delta}{\gamma_n}\right) \leq n^{-\beta}.$$

1155 For probability measures with upper bounded Lebesgue density, we can choose $\delta_n = n^{-\frac{2+\alpha}{d}}$ and
 1156 $\beta = \frac{9}{4} + \frac{\alpha}{2}$, for any fixed $\alpha > 0$, in the proof of Theorem 7. Hence the Laplace kernel only requires a
 1157 slow bandwidth decay rate of $\gamma_n = \Omega\left(\frac{n^{-\frac{2+\alpha}{d}}}{\alpha \ln(n)}\right)$, where $\alpha > 0$ arbitrary. For the Gaussian kernel an
 1158 analogous argument yields $\gamma_n = \Omega\left(\frac{n^{-\frac{4+2\alpha}{d}}}{\alpha \ln(n)}\right)$. The larger the dimension d , the slower the required
 1159 bandwidth decay. ◀

1160 **Remark G.3 (Generalizations).** If one does not care about continuous kernels, one could simply
 1161 take a Dirac kernel as the spike and then obtain consistency for all atom-free P_X . However, we
 1162 need a continuous kernel to be able to translate it to an activation function for the NTK. Beyond
 1163 kernel regression, the spike component \check{k}_γ does not even need to be a kernel, it just needs to fulfill
 1164 Assumption (SK) or a similar decay criterion. Then one could still use the 'quasi minimum-norm
 1165 estimator' $\mathbf{x} \mapsto (\tilde{k} + \rho_n \check{k}_{\gamma_n})(\mathbf{x}, \mathbf{X}) \cdot (\tilde{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^+ \mathbf{y}$. ◀

1166 **Remark G.4 (Consistency with a single kernel function).** Without resorting to kernel sequences as
 1167 we do, there seems to be no rigorous proof showing that ridgeless kernel regression can be consistent
 1168 in fixed dimension. In future work, can an analytical expression of such a kernel be found? According

1169 to the semi-rigorous results in Mallinar et al. (2022) a spectral decay like $\lambda_k = \Theta(k^{-1} \cdot \log^\alpha(k))$,
 1170 $\alpha > 1$ could lead to such a kernel. ◀

1171 *Proof of Theorem G.1.* Given any universal kernel, (Steinwart, 2001, Theorem 3.11 or Example 4.6)
 1172 implies universal consistency of kernel ridge regression if $\rho_n \rightarrow 0$ and $n\rho_n^4 \rightarrow \infty$. Hence, for any
 1173 $\varepsilon > 0$ it holds that

$$\lim_{n \rightarrow \infty} P^n \left(D \in (\mathbb{R}^d \times \mathbb{R})^n \mid R_P(\hat{f}_{\rho_n}^{\tilde{k}}) - R_P(f_P^*) = \mathbb{E}_{\mathbf{x}}(\hat{f}_{\rho_n}^{\tilde{k}}(\mathbf{x}) - f_P^*(\mathbf{x}))^2 \geq (\varepsilon/2)^2 \right) = 0.$$

Due to the triangle inequality in $L_2(P_X)$, we know

$$\begin{aligned} R_P(\hat{f}_0^{k_n}) - R_P(f_P^*) &= \mathbb{E}_{\mathbf{x}}(\hat{f}_0^{k_n}(\mathbf{x}) - f_P^*(\mathbf{x}))^2 \\ &\leq \left(\left(\mathbb{E}_{\mathbf{x}}(\hat{f}_0^{k_n}(\mathbf{x}) - \hat{f}_{\rho_n}^{\tilde{k}}(\mathbf{x}))^2 \right)^{1/2} + \left(\mathbb{E}_{\mathbf{x}}(\hat{f}_{\rho_n}^{\tilde{k}}(\mathbf{x}) - f_P^*(\mathbf{x}))^2 \right)^{1/2} \right)^2. \end{aligned}$$

1174 It is left to show that k_n fulfills

$$\lim_{n \rightarrow \infty} P^n \left(D \in (\mathbb{R}^d \times \mathbb{R})^n \mid \mathbb{E}_{\mathbf{x}}(\hat{f}_0^{k_n}(\mathbf{x}) - \hat{f}_{\rho_n}^{\tilde{k}}(\mathbf{x}))^2 \geq (\varepsilon/2)^2 \right) = 0.$$

For this purpose we decompose the above difference into the difference of $\check{\mathbf{K}}_{\gamma_n} := \check{k}_{\gamma_n}(\mathbf{X}, \mathbf{X})$ and \mathbf{I}_n and a remainder term depending on \check{k}_{γ_n} . We denote the 2-operator norm by $\|\cdot\|$ and the Euclidean norm in \mathbb{R}^n by $|\cdot|$. For any $\mathbf{x} \in \mathbb{R}^d$ it holds that

$$\begin{aligned} |\hat{f}_0^k(\mathbf{x}) - \hat{f}_{\rho_n}^{\tilde{k}}(\mathbf{x})| &\leq \left| (\tilde{k} + \rho_n \check{k}_{\gamma_n})(\mathbf{x}, \mathbf{X}) \cdot (\check{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} \mathbf{y} - \tilde{k}(\mathbf{x}, \mathbf{X}) \cdot (\check{\mathbf{K}} + \rho_n \mathbf{I}_n)^{-1} \mathbf{y} \right| \\ &\leq \left| \tilde{k}(\mathbf{x}, \mathbf{X}) \left((\check{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} - (\check{\mathbf{K}} + \rho_n \mathbf{I}_n)^{-1} \right) \mathbf{y} \right| \\ &\quad + \rho_n \cdot \left| \check{k}_{\gamma_n}(\mathbf{x}, \mathbf{X}) (\check{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} \mathbf{y} \right| \\ &\leq \|\tilde{k}(\mathbf{x}, \mathbf{X})\| \cdot \left\| (\check{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} - (\check{\mathbf{K}} + \rho_n \mathbf{I}_n)^{-1} \right\| \cdot |\mathbf{y}| \\ &\quad + \rho_n \cdot \|\check{k}_{\gamma_n}(\mathbf{x}, \mathbf{X})\| \cdot \left\| (\check{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} \right\| \cdot |\mathbf{y}|. \end{aligned}$$

Consequently we get

$$\mathbb{E}_{\mathbf{x}}(\hat{f}_0^k(\mathbf{x}) - \hat{f}_{\rho_n}^{\tilde{k}}(\mathbf{x}))^2 \leq 2 \mathbb{E}_{\mathbf{x}} \|\tilde{k}(\mathbf{x}, \mathbf{X})\|^2 \cdot \left\| (\check{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} - (\check{\mathbf{K}} + \rho_n \mathbf{I}_n)^{-1} \right\|^2 \cdot |\mathbf{y}|^2 \quad (\text{G.1})$$

$$+ 2 \rho_n^2 \cdot \mathbb{E}_{\mathbf{x}} \|\check{k}_{\gamma_n}(\mathbf{x}, \mathbf{X})\|^2 \cdot \left\| (\check{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} \right\|^2 \cdot |\mathbf{y}|^2. \quad (\text{G.2})$$

1175 We now bound the individual terms in Eq. (G.1) and (G.2). To this end, fix any $\alpha > 0$.

1176 **Bounding Eq. (G.1):**

1177 Since we assumed y_i i.i.d. and $\mathbb{E}y_1^2 < \infty$, the Markov inequality implies, with $b_n = \mathbb{E}y_1^2 \cdot n^\alpha$,

$$P(|\mathbf{y}|^2 \geq b_n n) \leq \frac{\mathbb{E}y_1^2}{b_n} = n^{-\alpha}.$$

1178 Stated differently, with probability at least $1 - n^{-\alpha}$ it holds that $|\mathbf{y}|^2 \leq \mathbb{E}y_1^2 \cdot n^{1+\alpha}$.

1179 In order to bound the spectrum of $\check{\mathbf{K}}_{\gamma_n}$, Lemma G.5 implies that there exists a positive sequence

1180 $\delta_\alpha(n) = n^{-\frac{2+\alpha}{\beta_X}}$ such that with probability at least $1 - O(n^{-\alpha})$ it holds that

$$\min_{i,j \in [n]: i \neq j} |\mathbf{x}_i - \mathbf{x}_j| \geq \delta_\alpha(n).$$

1181 Since (γ_n) fulfills $\gamma_n = O(n^{-\beta_\gamma})$ for any $\beta_\gamma > 0$, by Assumption (SK) there exists a sequence

1182 $\varepsilon_n = o(\rho_n n^{-2-\frac{\alpha}{\beta}})$ such that $\varepsilon(\gamma_n, \delta_\alpha(n)) \leq \varepsilon_n$. Assumption (SK) further implies that whenever

1183 $\min_{i,j \in [n]: i \neq j} |\mathbf{x}_i - \mathbf{x}_j| \geq \delta_\alpha(n)$ it holds that $(\check{\mathbf{K}}_{\gamma_n})_{ii} = 1$ and $0 \leq (\check{\mathbf{K}}_{\gamma_n})_{ij} \leq \varepsilon(\gamma_n, \delta_\alpha(n)) \leq \varepsilon_n$

1184 for $i \neq j$. Then Gershgorin's theorem (Gerschgorin, 1931) implies that for all eigenvalues of $\check{\mathbf{K}}_{\gamma_n}$

$$|\lambda_i(\check{\mathbf{K}}_{\gamma_n}) - 1| \leq (n-1)\varepsilon_n \text{ for all } i \in [n].$$

1185 This in turn implies

$$\|\check{\mathbf{K}}_{\gamma_n} - \mathbf{I}_n\| \leq (n-1)\varepsilon_n, \quad \lambda_{\max}(\check{\mathbf{K}}_{\gamma_n}) \leq 1 + (n-1)\varepsilon_n, \quad \lambda_{\min}(\check{\mathbf{K}}_{\gamma_n}) \geq 1 - (n-1)\varepsilon_n.$$

Using $\|(\tilde{\mathbf{K}} + \rho_n \mathbf{I}_n)^{-1}\| \leq \frac{1}{\lambda_{\min}(\tilde{\mathbf{K}}) + \rho_n} \leq \rho_n^{-1}$ and $\|\check{\mathbf{K}}_{\gamma_n} - \mathbf{I}_n\| \leq (n-1)\varepsilon_n$, Lemma G.6 implies

$$\begin{aligned} \left\| (\tilde{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} - (\tilde{\mathbf{K}} + \rho_n \mathbf{I}_n)^{-1} \right\| &\leq \frac{\|(\tilde{\mathbf{K}} + \rho_n \mathbf{I}_n)^{-1}\|^2 \cdot \rho_n \|\check{\mathbf{K}}_{\gamma_n} - \mathbf{I}_n\|}{1 - \|(\tilde{\mathbf{K}} + \rho_n \mathbf{I}_n)^{-1}\| \cdot \rho_n \|\check{\mathbf{K}}_{\gamma_n} - \mathbf{I}_n\|} \\ &\leq \frac{\rho_n^{-1}(n-1)\varepsilon_n}{1 - (n-1)\varepsilon_n}. \end{aligned}$$

1186 Using $|\check{k}(\mathbf{x}, \mathbf{X}_i)| \leq 1$ for all $i \in [n]$ yields the naive bound $\|\check{k}(\mathbf{x}, \mathbf{X})\|^2 \leq n$.

1187 Combining all terms in Eq. (G.1) yields its convergence to 0 as the product satisfies the rate
1188 $O(n^{4+\alpha} \rho_n^{-2} \varepsilon_n^2) = o(1)$ with probability at least $1 - 2n^{-\alpha}$.

1189 **Bounding Eq. (G.2):**

1190 The analysis below is restricted to the event of probability at least $1 - 2n^{-\alpha}$, on which the bound on
1191 Eq. (G.1) holds.

1192 Since $(n-1)\varepsilon_n \rightarrow 0$, for any $C > 1$ it holds for n large enough,

$$\rho_n \cdot \|(\tilde{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1}\| \leq \frac{\rho_n}{\lambda_{\min}(\tilde{\mathbf{K}}) + \rho_n(1 - (n-1)\varepsilon_n)} \leq \frac{1}{(1 - (n-1)\varepsilon_n)} \leq C.$$

1193 Finally we show $\sup_{\mathbf{x}' \in \mathbb{R}^d} \mathbb{E}_{\mathbf{x}} \check{k}_{\gamma_n}(\mathbf{x}, \mathbf{x}')^2 \leq 2n^{-(2+\alpha)}$ for n large enough.

Fix an arbitrary $\mathbf{x}' \in \mathbb{R}^d$. Then by construction of $\delta_\alpha(n)$ and ε_n it holds that

$$\begin{aligned} \mathbb{E}_{\mathbf{x}} \check{k}_{\gamma_n}(\mathbf{x}, \mathbf{x}')^2 &\leq 1 \cdot P_X(\{\mathbf{x} \in \mathbb{R}^d : \check{k}_{\gamma_n}(\mathbf{x}, \mathbf{x}')^2 \geq \varepsilon_n^2\}) + \varepsilon_n^2 \\ &\leq P_X(\{\mathbf{x} \in \mathbb{R}^d : |\mathbf{x} - \mathbf{x}'| < \delta_\alpha(n)\}) + \varepsilon_n^2 \\ &\leq \phi(\delta_\alpha(n)) + \varepsilon_n^2 \leq n^{-(2+\alpha)} + \varepsilon_n^2. \end{aligned}$$

1194 Since $\varepsilon_n^2 = o(\rho_n^2 n^{-4-\alpha})$, we get $\mathbb{E}_{\mathbf{x}} \check{k}_{\gamma_n}(\mathbf{x}, \mathbf{x}')^2 \leq 2n^{-(2+\alpha)}$ for n large enough.

1195 Combining all terms in Eq. (G.2) yields its convergence to 0 with the rate $O(n^{-(2+\alpha)} \cdot 1 \cdot n^{1+\alpha}) =$
1196 $O(n^{-1})$ with probability at least $1 - 2n^{-\alpha}$, which concludes the proof. \square

1197 G.1 Auxiliary results for the proof of Theorem 7

1198 The distributional Assumption (D2) immediately implies that the training points are separated with
1199 high probability.

1200 **Lemma G.5.** Assume (D2) is fulfilled with $\beta_X > 0$. Then with probability at least $1 - O(n^{-\alpha})$,

$$\min_{i,j \in [n]: i \neq j} |\mathbf{x}_i - \mathbf{x}_j| \geq n^{-\frac{2+\alpha}{\beta_X}}.$$

1201 *Proof.* For any $i \in [n]$, the union bound implies

$$P\left(\min_{j \in [n]: i \neq j} |\mathbf{x}_i - \mathbf{x}_j| \leq \delta\right) = P\left(\bigcup_{j \in [n]: j \neq i} \{\mathbf{x}_j \in B_\delta(\mathbf{x}_i)\}\right) \leq (n-1)\phi(\delta).$$

1202 Another union bound yields

$$P\left(\min_{i,j \in [n]: i \neq j} |\mathbf{x}_i - \mathbf{x}_j| \leq \delta\right) \leq n(n-1)\phi(\delta).$$

1203 Choosing $\delta_\alpha(n) = n^{-\frac{2+\alpha}{\beta_X}}$ yields $\phi(\delta_\alpha(n)) = O(\frac{1}{n^{2+\alpha}})$, which concludes the proof. \square

1204 The following lemma bounds $\|\mathbf{A}^{-1} - \mathbf{B}^{-1}\|$ via $\|\mathbf{A}^{-1}\|$ and $\|\mathbf{A} - \mathbf{B}\|$. Similar results can for
1205 example be found in (Horn and Johnson, 2013, Section 5.8).

1206 **Lemma G.6.** Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ be invertible matrices and let $\|\cdot\|$ be a submultiplicative matrix
 1207 norm with $\|\mathbf{I}_n\| = 1$. If \mathbf{A} and \mathbf{B} fulfill $\|\mathbf{A}^{-1}\|\|\mathbf{A} - \mathbf{B}\| < 1$, then it holds that

$$\|\mathbf{B}^{-1} - \mathbf{A}^{-1}\| \leq \frac{\|\mathbf{A}^{-1}\|^2 \cdot \|\mathbf{A} - \mathbf{B}\|}{1 - \|\mathbf{A}^{-1}\|\|\mathbf{A} - \mathbf{B}\|}.$$

1208 *Proof.* Because of $\|\mathbf{A}^{-1}(\mathbf{A} - \mathbf{B})\| \leq \|\mathbf{A}^{-1}\|\|\mathbf{A} - \mathbf{B}\| < 1$ we get

$$\|\mathbf{I} - \mathbf{A}^{-1}(\mathbf{A} - \mathbf{B})\| \geq 1 - \|\mathbf{A}^{-1}\|\|\mathbf{A} - \mathbf{B}\|.$$

1209 Writing $\mathbf{B} = \mathbf{A}(\mathbf{I} - \mathbf{A}^{-1}(\mathbf{A} - \mathbf{B}))$ yields $\mathbf{B}^{-1} = (\mathbf{I} - \mathbf{A}^{-1}(\mathbf{A} - \mathbf{B}))^{-1}\mathbf{A}^{-1}$ which implies

$$\|\mathbf{B}^{-1}\| \leq \frac{\|\mathbf{A}^{-1}\|}{1 - \|\mathbf{A}^{-1}\|\|\mathbf{A} - \mathbf{B}\|}.$$

1210 Now write $\mathbf{B}^{-1} - \mathbf{A}^{-1} = \mathbf{A}^{-1}(\mathbf{A} - \mathbf{B})\mathbf{B}^{-1}$ to get

$$\|\mathbf{B}^{-1} - \mathbf{A}^{-1}\| \leq \|\mathbf{A}^{-1}\|\|\mathbf{A} - \mathbf{B}\|\|\mathbf{B}^{-1}\|.$$

1211 Combining the last two inequalities concludes the proof. \square

1212 G.2 RKHS norm bounds

1213 Here we show that if \tilde{k} and \tilde{k}_γ have RKHS equivalent to some Sobolev space H^s , $s > d/2$, then
 1214 the RKHS of the spiky-smooth kernel $k_{\rho,\gamma}$ is also equivalent to H^s , for any fixed $\rho, \gamma > 0$. Hence
 1215 all members of the spiky-smooth kernel sequence may have RKHS equivalent to a Sobolev space
 1216 H^s and are individually inconsistent due to Theorem 1; yet the sequence is consistent. This shows
 1217 that when arguing about generalization properties based on RKHS equivalence, the constants matter
 1218 and the narrative that depth does not matter in the NTK regime as in Bietti and Bach (2021) is too
 1219 simplified.

1220 The following proposition states that the sum of kernels with equivalent RKHS yields an RKHS that
 1221 is equivalent to the RKHS of the summands. For example, the spiky-smooth kernel with Laplace
 1222 components possesses an RKHS equivalent to the RKHS of the Laplace kernel.

1223 **Proposition G.7.** Let \mathcal{H}_1 and \mathcal{H}_2 denote the RKHS of k_1 and k_2 respectively. If $\mathcal{H}_1 = \mathcal{H}_2$ then
 1224 the RKHS \mathcal{H} of $k = k_1 + k_2$ fulfills $\mathcal{H} = \mathcal{H}_1$. Moreover, if $C \geq 1$ is a constant with $\frac{1}{C}\|f\|_{\mathcal{H}_2} \leq$
 1225 $\|f\|_{\mathcal{H}_1} \leq C\|f\|_{\mathcal{H}_2}$, then we have $\frac{1}{\sqrt{2C}}\|f\|_{\mathcal{H}_1} \leq \|f\|_{\mathcal{H}} \leq \|f\|_{\mathcal{H}_1}$.

Proof. The RKHS of $k = k_1 + k_2$ is given by $\mathcal{H} = \mathcal{H}_1 + \mathcal{H}_2$ with norm

$$\|f\|_{\mathcal{H}}^2 = \min\{\|f_1\|_{\mathcal{H}_1}^2 + \|f_2\|_{\mathcal{H}_2}^2 : f = f_1 + f_2, f_1 \in \mathcal{H}_1, f_2 \in \mathcal{H}_2\}.$$

1226 To see this we consider the map $\Phi : X \rightarrow \mathcal{H}_1 \times \mathcal{H}_2$ defined by $\Phi(\mathbf{x}) := (\Phi_1(\mathbf{x}, \cdot), \Phi_2(\mathbf{x}, \cdot))$ for all
 1227 $\mathbf{x} \in X$, where X is the set, the spaces \mathcal{H}_i live on and $\Phi_i(\mathbf{x}) := k_i(\mathbf{x}, \cdot)$. The reproducing property
 1228 of k_1 and k_2 immediately ensures that Φ is a feature map of $k_1 + k_2$ and Theorem E.3 then shows

$$\begin{aligned} \mathcal{H} &= \{\langle w, \Phi(\cdot) \rangle_{\mathcal{H}_1 \times \mathcal{H}_2} : w \in \mathcal{H}_1 \times \mathcal{H}_2\} \\ &= \{\langle w_1, \Phi_1(\cdot) \rangle_{\mathcal{H}_1} + \langle w_2, \Phi_2(\cdot) \rangle_{\mathcal{H}_2} : w_1 \in \mathcal{H}_1, w_2 \in \mathcal{H}_2\} = \mathcal{H}_1 + \mathcal{H}_2 \end{aligned}$$

1229 as well as the formula for the norm on \mathcal{H} . Now let $f \in \mathcal{H}$. Considering the decomposition $f = f_1 + 0$
 1230 then gives $\|f\|_{\mathcal{H}} \leq \|f\|_{\mathcal{H}_1}$. Moreover, for $f = f_1 + f_2$ with $f_i \in \mathcal{H}_i$ we have

$$\|f\|_{\mathcal{H}_1} \leq \|f_1\|_{\mathcal{H}_1} + \|f_2\|_{\mathcal{H}_1} \leq \|f_1\|_{\mathcal{H}_1} + C\|f_2\|_{\mathcal{H}_2} \leq \sqrt{2C}(\|f_1\|_{\mathcal{H}_1}^2 + \|f_2\|_{\mathcal{H}_2}^2)^{1/2}.$$

1231 Taking the infimum over all decomposition then yields the estimate $\|f\|_{\mathcal{H}_1} \leq \sqrt{2C}\|f\|_{\mathcal{H}}$. \square

1232 H Spiky-smooth activation functions induced by Gaussian components

1233 Here we explore the properties of the NNGP and NTK activation functions induced by spiky-smooth
1234 kernels with Gaussian components.

1235 To offer some more background, it is well-known that NNGPs and NTKs on the sphere \mathbb{S}^d are
1236 dot-product kernels, i.e., kernels of the form $k_d(\mathbf{x}, \mathbf{x}') = \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle)$, where the function κ has a
1237 series representation $\kappa(t) = \sum_{i=0}^{\infty} b_i t^i$ with $b_i \geq 0$ and $\sum_{i=0}^{\infty} b_i < \infty$. The function κ is independent
1238 of the dimension d of the sphere. Conversely, all such kernels can be realized as NNGPs or NTKs
1239 (Simon et al., 2022, Theorem 3.1).

1240 As dot-product kernel $k(\mathbf{x}, \mathbf{y}) = \kappa(\langle \mathbf{x}, \mathbf{y} \rangle)$ on the sphere, the Gaussian kernel has the simple analytic
1241 expression,

$$\kappa_{\gamma}^{Gauss}(z) = \exp\left(\frac{2(z-1)}{\gamma}\right),$$

1242 with Taylor expansion

$$\kappa_{\gamma}^{Gauss}(z) = \sum_{i=0}^{\infty} \underbrace{\frac{2^i}{\gamma^i i!}}_{b_i^{Gauss}} \exp(-2/\gamma) z^i.$$

For spiky-smooth kernels $k = \tilde{k} + \rho \check{k}_{\gamma}$ with Gaussian components \tilde{k} and \check{k}_{γ} of width $\tilde{\gamma}$ and γ
respectively, we get Taylor series coefficients

$$b_i = \frac{\exp(-2/\tilde{\gamma})}{i!} \left(\frac{2}{\tilde{\gamma}}\right)^i + \rho \frac{\exp(-2/\gamma)}{i!} \left(\frac{2}{\gamma}\right)^i. \quad (\text{H.1})$$

1243 Now Theorem 8 states that as soon as κ induces a dot-product kernel for every input dimension d ,
1244 then the dot-product kernels can be written as the NNGP kernel of a 2-layer fully-connected network
1245 without biases and with the induced activation function

$$\phi_{NNGP}^{\kappa}(x) = \sum_{i=0}^{\infty} s_i b_i^{1/2} h_i(x),$$

1246 or as the NTK of a 2-layer fully-connected network without biases and with the induced activation
1247 function

$$\phi_{NTK}^{\kappa}(x) = \sum_{i=0}^{\infty} s_i \left(\frac{b_i}{i+1}\right)^{1/2} h_i(x),$$

1248 where h_i denotes the i -th Probabilist's Hermite polynomial normalized such that $\|h_i\|_{L_2(\mathcal{N}(0,1))} = 1$
1249 and $s_i \in \{-1, +1\}$ are arbitrarily chosen for all $i \in \mathbb{N}$.

1250 Now we can study the induced activation functions if we know the kernel's Taylor coefficients
1251 $(b_i)_{i \in \mathbb{N}_0}$. If infinitely many $b_i > 0$, then infinitely many activation functions induce the same dot-
1252 product kernel, with different choices of the signs s_i . For alternating signs $s_i = -1$ if i is odd and
1253 $s_i = 1$ if i is even, the symmetry property $h_i(-x) = (-1)^i h_i(x)$ of the Hermite polynomials implies

$$\phi_{NNGP,+,-}(x) = \phi_{NNGP,+}(-x), \quad \phi_{NTK,+,-}(x) = \phi_{NTK,+}(-x).$$

To form an orthonormal basis of $L_2(\mathcal{N}(0,1))$ the unnormalized Probabilist's Hermite polynomials
 He_i have to be normalized by $h_i(x) = \frac{1}{\sqrt{i!}} He_i(x)$. We can use the identity $\exp(xt - \frac{t^2}{2}) =$
 $\sum_{i=0}^{\infty} He_i(x) \frac{t^i}{i!}$ with $t = \sqrt{2/\gamma}$ to analytically express the NNGP activation of the Gaussian kernel
with all $s_i = +1$ as the exponential function

$$\phi_{NNGP}^{Gauss}(x) = \exp(-1/\gamma) \sum_{i=0}^{\infty} \frac{1}{i!} \left(\frac{2}{\gamma}\right)^{\frac{i}{2}} h_i(x) = \exp\left(\left(\frac{2}{\gamma}\right)^{\frac{1}{2}} \cdot x - \frac{2}{\gamma}\right). \quad (\text{H.2})$$

For $\phi_{NNGP}(x) = \sum_{i=0}^{\infty} s_i \sqrt{b_i} h_i(x)$, we get $\|\phi\|_{L_2(\mathcal{N}(0,1))}^2 = \sum_{i=0}^{\infty} b_i$ invariant to the choice
 $\{s_i\}_{i \in \mathbb{N}}$, which yields

$$\|\phi_{NNGP}^{Gauss}\|_{L_2(\mathcal{N}(0,1))}^2 = \exp(-2/\gamma) \sum_{i=0}^{\infty} \frac{1}{i!} \left(\frac{2}{\gamma}\right)^i = 1, \quad (\text{H.3})$$

because $\{h_i\}_{i \in \mathbb{N}_0}$ is an ONB of $L_2(\mathcal{N}(0, 1))$. Analogously,

$$\begin{aligned} \|\phi_{NTK}^{Gauss}\|_{L_2(\mathcal{N}(0,1))}^2 &= \exp(-2/\gamma) \sum_{i=0}^{\infty} \frac{1}{(i+1)!} \left(\frac{2}{\gamma}\right)^i \\ &= \exp(-2/\gamma) \frac{\gamma}{2} \sum_{i=1}^{\infty} \frac{1}{i!} \left(\frac{2}{\gamma}\right)^i = \frac{\gamma}{2} \left(1 - \exp\left(-\frac{2}{\gamma}\right)\right). \end{aligned} \quad (\text{H.4})$$

1254 This implies that the average amplitude of NNGP activation functions does not depend on γ , while
1255 the average amplitude of NTK activation functions decays with $\gamma \rightarrow 0$.

1256 By the fact $h'_n(x) = \sqrt{n}h_{n-1}(x)$, we know that any activation function $\phi(x) = \sum_{i=0}^{\infty} s_i a_i h_i(x)$ has
1257 the derivative $\phi'(x) = \sum_{i=0}^{\infty} s_i a_{i+1} \sqrt{i+1} \cdot h_i(x)$ as long as $\sum_{i=0}^{\infty} |a_{i+1} \sqrt{i+1}| < \infty$.

1258 The following proposition formalizes the additive approximation $\phi^k \approx \phi^{\tilde{k}} + \rho^{1/2} \phi^{\tilde{k}_\gamma}$, and quantifies
1259 the necessary scaling of γ for any demanded precision of the approximation.

Proposition H.1. Fix $\tilde{\gamma}, \rho > 0$ arbitrary. Let $k = \tilde{k} + \rho \tilde{k}_\gamma$ denote the spiky-smooth kernel where \tilde{k} and \tilde{k}_γ are Gaussian kernels of bandwidth $\tilde{\gamma}$ and γ , respectively. Assume that we choose the activation functions $\phi_{NTK}^k, \phi_{NTK}^{\tilde{k}}$ and $\phi_{NTK}^{\tilde{k}_\gamma}$ as in Theorem 8 with same signs $\{s_i\}_{i \in \mathbb{N}}$. Then, for $\gamma > 0$ small enough, it holds that

$$\begin{aligned} \|\phi_{NTK}^k - (\phi_{NTK}^{\tilde{k}} + \sqrt{\rho} \cdot \phi_{NTK}^{\tilde{k}_\gamma})\|_{L_2(\mathcal{N}(0,1))}^2 &\leq 2^{1/2} \rho \gamma^{3/2} \exp\left(-\frac{1}{\gamma}\right) + \frac{4\pi\gamma(1+\tilde{\gamma})}{\tilde{\gamma}}, \\ \|\phi_{NNGP}^k - (\phi_{NNGP}^{\tilde{k}} + \sqrt{\rho} \cdot \phi_{NNGP}^{\tilde{k}_\gamma})\|_{L_2(\mathcal{N}(0,1))}^2 &\leq 2^{3/2} \rho \gamma^{1/2} \exp\left(-\frac{1}{\gamma}\right) + \frac{8\pi\gamma(1+\tilde{\gamma})}{\tilde{\gamma}^2}. \end{aligned}$$

Proof. Let $b_{i,\gamma} = \frac{2^i}{\gamma^i i!} \exp(-2/\gamma)$ denote the Taylor coefficients of the Gaussian kernel. All considered infinite series converge absolutely.

$$\begin{aligned} &\|\phi_{NNGP}^{\tilde{\gamma}; \gamma; \rho} - (\phi_{NNGP}^{\tilde{\gamma}} + \sqrt{\rho} \cdot \phi_{NNGP}^{\gamma})\|_{L_2(\mathcal{N}(0,1))}^2 \\ &= \left\| \sum_{i=0}^{\infty} s_i \sqrt{b_{i,\tilde{\gamma}} + \rho b_{i,\gamma}} h_i(x) - \sum_{i=0}^{\infty} s_i (\sqrt{b_{i,\tilde{\gamma}}} + \sqrt{\rho b_{i,\gamma}}) h_i(x) \right\|_{L_2(\mathcal{N}(0,1))}^2 \\ &= \sum_{i=0}^{\infty} \left(\sqrt{b_{i,\tilde{\gamma}} + \rho b_{i,\gamma}} - (\sqrt{b_{i,\tilde{\gamma}}} + \sqrt{\rho b_{i,\gamma}}) \right)^2 \\ &\leq 2 \underbrace{\sum_{i=0}^I (\sqrt{b_{i,\tilde{\gamma}} + \rho b_{i,\gamma}} - b_{i,\tilde{\gamma}}^{1/2})^2}_{(I)} + 2\rho \underbrace{\sum_{i=0}^I b_{i,\gamma}}_{(II)} + 2 \underbrace{\sum_{i=I+1}^{\infty} (\sqrt{b_{i,\tilde{\gamma}} + \rho b_{i,\gamma}} - \rho^{1/2} b_{i,\gamma}^{1/2})^2}_{(III)} + 2 \underbrace{\sum_{i=I+1}^{\infty} b_{i,\tilde{\gamma}}}_{(IV)}, \end{aligned}$$

for any $I \in \mathbb{N}$. To bound (I) observe

$$\sum_{i=0}^I (\sqrt{b_{i,\tilde{\gamma}} + \rho b_{i,\gamma}} - b_{i,\tilde{\gamma}}^{1/2})^2 = \sum_{i=0}^I \left(\rho b_{i,\gamma} + 2b_{i,\tilde{\gamma}} \left(1 - \sqrt{1 + \frac{\rho b_{i,\gamma}}{b_{i,\tilde{\gamma}}}}\right) \right) \leq \rho \sum_{i=0}^I b_{i,\gamma}.$$

1260 An analogous calculation for (III) yields

$$\sum_{i=I+1}^{\infty} (\sqrt{b_{i,\tilde{\gamma}} + \rho b_{i,\gamma}} - \rho^{1/2} b_{i,\gamma}^{1/2})^2 \leq \sum_{i=I+1}^{\infty} b_{i,\tilde{\gamma}}.$$

So overall we get the bound

$$\|\phi_{NNGP}^{\tilde{\gamma}; \gamma; \rho} - (\phi_{NNGP}^{\tilde{\gamma}} + \sqrt{\rho} \cdot \phi_{NNGP}^{\gamma})\|_{L_2(\mathcal{N}(0,1))}^2 \leq 4\rho \sum_{i=0}^I b_{i,\gamma} + 4 \sum_{i=I+1}^{\infty} b_{i,\tilde{\gamma}}. \quad (\text{H.5})$$

1261 Now, defining $c := 2/\gamma$,

$$\sum_{i=0}^I b_{i,\gamma} = \exp(-c) \sum_{i=0}^I \frac{1}{i!} c^i = \frac{\Gamma(I+1, c)}{I!},$$

where $\Gamma(k+1, c)$ denotes the upper incomplete Gamma function. Choosing $I = \lfloor \frac{c}{2\pi} \rfloor$, (Pinelis, 2020, Theorem 1.1) yields, for $c \geq 121$,

$$\begin{aligned} \frac{\Gamma(I+1, c)}{I!} &\leq \exp(-c) \frac{(c + (I+1)!^{1/I})^{I+1}}{(I+1)! \cdot (I+1)!^{1/I}} \leq \frac{\exp(-c)(c+I)^{I+1}}{(I+1)!^{(I+1)/I}} \\ &\leq \frac{\exp(-c)(c+I)^{I+1}}{(2\pi(I+1))^{1/2} \left(\frac{I+1}{e}\right)^{(I+1)^2/I}} \\ &\leq \frac{1}{\sqrt{2\pi(I+1)}} \exp\left(-c + (I+1)(\ln(c+I) - \ln(I+1) + 1)\right) \\ &\leq \frac{1}{\sqrt{c}} \exp\left(-c + \left(\frac{c}{2\pi} + 1\right)(\ln\left(\frac{2\pi+1}{2\pi}c\right) - \ln\left(\frac{c}{2\pi}\right) + 1)\right) \\ &\leq \frac{1}{\sqrt{c}} \exp\left(-c + \left(\frac{c}{2\pi} + 1\right)(\ln(2\pi+1) + 1)\right) \leq \frac{\exp\left(-\frac{c}{2}\right)}{\sqrt{c}}, \end{aligned} \quad (\text{H.6})$$

1262 where we used $(I+1)!^{1/I} \leq I$ for $I \geq 3$ in the first line, Stirling's approximation in the second line,
1263 and $(\frac{c}{2\pi} + 1)(\ln(2\pi+1) + 1) \leq c/2$ for $c \geq 121$ in the last line.

1264 It is obvious that

$$\sum_{i=I+1}^{\infty} b_{i, \tilde{\gamma}} \rightarrow 0, \quad \text{for } I = \lfloor \frac{c}{2\pi} \rfloor \rightarrow \infty.$$

To quantify the rate of convergence, we use the bound $\Gamma(I+1, c_0) \geq e^{-c_0} I! (1 + c_0/(I+1))^I$, which follows from applying Jensen's inequality to $\Gamma(I+1, c_0) = e^{-c_0} I! \mathbb{E}(1 + c_0/G)^I$, where $G \sim \Gamma(I+1, 1)$ and $\mathbb{E}G = I+1$. Defining $c_0 = 2/\tilde{\gamma}$, it holds that

$$\sum_{i=I+1}^{\infty} b_{i, \tilde{\gamma}} \leq 1 - \frac{\Gamma(I+1, c_0)}{I!} \leq 1 - e^{-c_0} \left(1 + \frac{c_0}{I+1}\right)^I \leq 1 - e^{-c_0} \left(1 + \frac{c_0}{I+1}\right)^{I+1}.$$

1265 Taking the first two terms of the Laurent series expansion of $n \mapsto (1 + \frac{c_0}{n})^n$ about $n = \infty$ yields

1266 $\left(1 + \frac{c_0}{I+1}\right)^{I+1} > e^{c_0} (1 - \frac{c_0^2}{2(I+1)})$ for I large enough (where we demand $\gamma \in o(\tilde{\gamma}^2)$), thus

$$\begin{aligned} \sum_{i=I+1}^{\infty} b_{i, \tilde{\gamma}} &\leq 1 - e^{-c_0} \left(1 + \frac{c_0}{I+1}\right)^{I+1} \cdot \left(1 + \frac{c_0}{I+1}\right)^{-1} \\ &\leq \frac{c_0/(I+1) + c_0^2/(2(I+1))}{1 + c_0/(I+1)} \leq \frac{c_0}{I+1} + \frac{c_0^2}{2(I+1)} \leq \frac{4\pi}{\tilde{\gamma}c} + \frac{4\pi}{\tilde{\gamma}^2 c}. \end{aligned} \quad (\text{H.7})$$

1267 Plugging (H.6) and (H.7) into (H.5) yields, for $\gamma \leq 1/61$,

$$\|\phi_{NNGP}^{\tilde{\gamma}, \gamma, \rho} - (\phi_{NNGP}^{\tilde{\gamma}} + \sqrt{\rho} \cdot \phi_{NNGP}^{\gamma})\|_{L_2(\mathcal{N}(0,1))}^2 \leq 2^{3/2} \rho \gamma^{1/2} \exp\left(-\frac{1}{\gamma}\right) + \frac{8\pi\gamma(1+\tilde{\gamma})}{\tilde{\gamma}^2}.$$

For the NTK we get

$$\begin{aligned} &\|\phi_{NTK}^{\tilde{\gamma}, \gamma, \rho} - (\phi_{NTK}^{\tilde{\gamma}} + \sqrt{\rho} \cdot \phi_{NTK}^{\gamma})\|_{L_2(\mathcal{N}(0,1))}^2 \\ &= \left\| \sum_{i=0}^{\infty} s_i \sqrt{\frac{b_{i, \tilde{\gamma}} + \rho b_{i, \gamma}}{i+1}} h_i(x) - \sum_{i=0}^{\infty} s_i \left(\sqrt{\frac{b_{i, \tilde{\gamma}}}{i+1}} + \sqrt{\frac{\rho b_{i, \gamma}}{i+1}} \right) h_i(x) \right\|_{L_2(\mathcal{N}(0,1))}^2 \\ &= \sum_{i=0}^{\infty} \frac{1}{i+1} \left(\sqrt{b_{i, \tilde{\gamma}} + \rho b_{i, \gamma}} - (\sqrt{b_{i, \tilde{\gamma}}} + \sqrt{\rho b_{i, \gamma}}) \right)^2. \end{aligned}$$

1268 We can proceed exactly as for the NNGP, but choose $I = \lfloor \frac{c}{2\pi} \rfloor - 1$ to get

$$\sum_{i=0}^I \frac{b_{i, \gamma}}{i+1} = \exp(-c) \sum_{i=0}^I \frac{c^i}{(i+1)!} = \frac{\exp(-c)}{c} \left(\sum_{i=0}^{I+1} \frac{c^i}{i!} - 1 \right) \leq \frac{\exp(-c/2)}{c^{3/2}} - \frac{\exp(-c)}{c},$$

and replace (H.7) with

$$\sum_{i=I+1}^{\infty} \frac{b_{i, \tilde{\gamma}}}{i+1} = \frac{\exp(-c_0)}{c_0} \sum_{i=I+2}^{\infty} \frac{c_0^i}{i!} \leq \frac{1}{I+2} + \frac{c_0}{2(I+2)} \leq \frac{\pi\gamma(1+\tilde{\gamma})}{\tilde{\gamma}}. \quad \square$$

1269 I Additional experimental results

1270 The code to reproduce all our experiments is provided in the supplementary material. Our imple-
1271 mentations rely on PyTorch (Paszke et al., 2019) for neural networks and mpmath (Johansson et al.,
1272 2023) for high-precision calculations.

1273 I.1 Experimental details of Figure 1

1274 For the kernel experiment (Figure 1a), we used the Laplace kernel with bandwidth 0.4 and the
1275 spiky-smooth kernel (4) with Laplace components with $\rho = 1, \tilde{\gamma} = 1, \gamma = 0.01$.

1276 For the neural network experiment (Figure 1b,c) we initialize 2-layer networks with NTK parametriza-
1277 tion (Jacot et al., 2018) and He initialization (He et al., 2015). Using the antisymmetric initialization
1278 trick from Zhang et al. (2020) doubles the network width from 10000 to 20000 and helps to prevent
1279 errors induced by the random initialization function. We train the network with stochastic gradient
1280 descent of batch size 1 over the 15 training samples with learning rate 0.04 for 2500 epochs. Train-
1281 ing with gradient descent and learning rate 0.4 produces similar results. We use the spiky-smooth
1282 activation function given by $x \mapsto ReLU(x) + 0.01 \cdot (\sin(100x) + \cos(100x))$, which corresponds
1283 to $x \mapsto ReLU(x) + \omega_{NTK}(x, \frac{1}{5000})$, including both even and uneven Hermite coefficients.

1284 I.2 Disentangling signal from noise in neural networks with spiky-smooth activation functions

Since our spiky-smooth activation function has the additive form $\sigma_{spsm}(x) = ReLU(x) + \omega_{NTK}(x; \frac{1}{5000})$, we can dissect the learned neural network

$$f_{spsm}(\mathbf{x}) = \mathbf{W}_2 \cdot \sigma_{spsm}(\mathbf{W}_1 \cdot \mathbf{x} + \mathbf{b}_1) + b_2 = f_{ReLU}(\mathbf{x}) + f_{spikes}(\mathbf{x}) \quad (\text{I.1})$$

1285 into its *ReLU*-component

$$f_{ReLU}(\mathbf{x}) = \mathbf{W}_2 \cdot ReLU(\mathbf{W}_1 \cdot \mathbf{x} + \mathbf{b}_1) + b_2,$$

1286 and its spike component

$$f_{spikes}(\mathbf{x}) = \mathbf{W}_2 \cdot \omega_{NTK}(\mathbf{W}_1 \cdot \mathbf{x} + \mathbf{b}_1; \frac{1}{5000}).$$

1287 If the analogy to the spiky-smooth kernel holds and f_{spikes} fits the noise in the labels while having a
1288 small L_2 -norm, then f_{ReLU} would have learned the signal in the data. Indeed Figure I.1 demonstrates
1289 that this simple decomposition is useful to disentangle the learned signal from the spike component
1290 in our setting. The figure also suggests that the oscillations in the activations of the hidden layer
1291 constructively interfere to interpolate the training points, while the differing frequencies and phases
1292 approximately destructively interfere on most of the remaining covariate support. Figure I.2 shows
1293 some of the functions generated by the hidden layer neurons of the spike component f_{spikes} . Both
1294 the phases and frequencies vary. Destructive interference in sums of many oscillations occurs, for
1295 example, under a uniform phase distribution.

1296 An exciting direction of future work will be to understand when and why the neural networks with
1297 spiky-smooth activation functions learn the target function well, and when the decomposition into
1298 *ReLU*- and spike component succeeds to disentangle the noise from the signal. Particular challenges
1299 will be to design architectures and learning algorithms that provably work on complex data sets and
1300 to determine their statistical convergence rates. A different line of work could evaluate whether there
1301 exist useful spike components for deep and narrow networks beyond the pure infinite-width limit.
1302 Maybe for deep architectures it suffices to apply spiky-smooth activation functions only between the
1303 penultimate and the last layer.




Figure I.1: **a.** The $ReLU$ -component f_{ReLU} (blue) and the full spiky-smooth network f_{spsm} (orange) of the learned neural network from Figure 1. **b.** The spike component f_{spikes} of the learned neural network from Figure 1 against the label noise in the training set, derived by subtracting the signal from the training points. Observe that the $ReLU$ -component has learned the signal, while the spike component has fitted the noise in the data while regressing to 0 between data points.




Figure I.2: Here we plot the functions learned by 12 random hidden layer neurons of the spike component network f_{spikes} corresponding to Figure 1.

Of course an analogous additive decomposition exists for the minimum-norm interpolant \hat{f}_0^k of the spiky-smooth kernel,

$$\hat{f}_0^k(\mathbf{x}) = (\tilde{k} + \rho_n \check{k}_{\gamma_n})(\mathbf{x}, \mathbf{X}) \cdot (\tilde{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} \mathbf{y} = f_{signal}(\mathbf{x}) + f_{spikes}(\mathbf{x}), \quad (I.2)$$

1304 where

$$f_{signal}(\mathbf{x}) = \tilde{k}(\mathbf{x}, \mathbf{X}) \cdot (\tilde{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} \mathbf{y}, \quad f_{spikes}(\mathbf{x}) = \rho_n \check{k}_{\gamma_n}(\mathbf{x}, \mathbf{X}) \cdot (\tilde{\mathbf{K}} + \rho_n \check{\mathbf{K}}_{\gamma_n})^{-1} \mathbf{y}.$$

1305 We plot the results in Figure I.3. Observe that the spikes f_{spikes} regress to 0 more reliably than in the
1306 neural network.

1307 Although spiky-smooth estimators can be consistent, any method that interpolates noise cannot
1308 be adversarially robust. The signal component f_{signal} may be a simple correction towards robust
1309 estimators. Figure I.4 suggests that the signal components of spiky-smooth estimators behave
1310 more robustly than ReLU networks or minimum-norm interpolants of Laplace kernels in terms of
1311 finite-sample variance.




Figure I.3: **a.** The signal component f_{signal} (blue) and the full minimum-norm interpolant \hat{f}_0^k (orange) of the spiky-smooth kernel from Figure 1. **b.** The spike component f_{spikes} of the spiky-smooth kernel from Figure 1 against the label noise in the training set, derived by subtracting the signal from the training points.

1312 I.3 Repeating the finite-sample experiments

1313 We repeat the experiment from Figure 1 100 times, both randomizing with respect to the training set
 1314 and with respect to neural network initialization.

1315 For the kernels (Figure I.4a), observe that all minimum-norm kernel interpolants are biased towards
 1316 0. While the Laplace kernel and the signal component (I.2) of the spiky-smooth kernel have similar
 1317 averages, the spiky-smooth kernel has a slightly larger bias. However, both the spiky-smooth kernel
 1318 as well as its signal component produces lower variance estimates than the Laplace kernel.

1319 Considering the trained neural networks (Figure I.4b), the ReLU networks are approximately unbiased,
 1320 but have large variance. The neural networks with spiky-smooth activation function as well as the
 1321 extracted signal network (I.1) are similar on average: They are slightly biased towards 0, but have
 1322 much smaller variance than the ReLU networks.




Figure I.4: We repeat the experiment from Figure 1 100 times and report the mean values (lines). Confidence bands denote the interval between the empirical 2.5%- and 97.5%-quantiles from the 100 independent runs.

1323 The training curves (Figure I.4c) offer similar conclusions as Figure 1: While the ReLU networks
 1324 harmfully overfit over the course of training, the neural networks with spiky-smooth activation
 1325 function quickly overfit to 0 training error with monotonically decreasing test error, which on average
 1326 is almost optimal, already with only 15 training points. The spiky-smooth networks have smaller
 1327 confidence bands, indicating increased robustness compared to the ReLU networks. If the ReLU
 1328 networks would be early-stopped with perfect timing, they would generalize similarly well as the
 1329 networks with spiky-smooth activation function.

1330 I.4 Spiky-smooth activation functions

1331 In Figures I.5 and I.6 we plot the 2-layer NTK activation functions induced by spiky-smooth kernels
 1332 with Gaussian components, where \tilde{k} has bandwidth 1, and in the first figure $\rho = 1$ while in the second
 1333 figure $\rho = 0.1$.




Figure I.5: The 2-layer NTK activation functions for Gaussian-Gaussian spiky-smooth kernels with varying γ (columns) with $k_{max} = 1000$, \tilde{k} has bandwidth 1, $\rho = 1$. Top: all $s_i = +1$, middle: +, +, -, -, +, +, ..., bottom: Random +1 and -1. Although the activation function induced by the spiky-smooth kernel is not exactly the sum of the activation functions induced by its components, this approximation is accurate because the spike components approximately live in a subspace of higher frequency in the Hermite basis orthogonal to the low-frequency subspace of the smooth component.




Figure I.6: Same as above but $\rho = 0.1$. The high-frequency fluctuations are much smaller compared to Figure I.5.

1334 In Figure I.7 we plot the corresponding 2-layer NNGP activation functions with $\rho = 1$. In contrast to
 1335 the NTK activation functions the amplitudes of the fluctuations only depends on ρ and not on γ . Our
 1336 intuition is the following: Since the first layer weights are not learned in case of the NNGP, the first
 1337 layer cannot learn constructive interference, so that the oscillations in the activation function need to
 1338 be larger.

1339 The additive approximation $\phi^k \approx \phi^{\tilde{k}} + \rho^{1/2} \phi^{\tilde{k}\gamma}$ remains accurate in all considered cases (Ap-
 1340 pendix I.6).




Figure I.7: Same as above but NNGP and $\rho = 1$. As expected from the isolated spike plots: Spikes essentially add fluctuations that increase in frequency and stay constant in amplitude for $\gamma \rightarrow 0$, ρ regulates the amplitude.

1341 **I.5 Isolated spike activation functions**

1342 Figure I.8 is the equivalent of Figure 3 for the NNGP.

1343 By plotting the NTK activation components corresponding to Gaussian spikes $\phi^{\tilde{k}\gamma}$ with varying
 1344 choices of the signs s_i in Figure I.9, we observe the following properties:

- 1345 1. All $s_i = +1$ leads to exponentially exploding activation functions, cf. Eq. (H.2).
 1346 2. If the signs s_i alternate every second i , i.e. $s_i = +1$ iff $\lfloor \frac{i}{2} \rfloor$ even, $\phi^{\tilde{k}\gamma}$ is approximately a
 1347 single shifted sin-curve with increasing frequency and decreasing/constant amplitude for
 1348 NTK/NNGP activation functions, cf. Eq. (6).
 1349 3. If s_i is chosen uniformly at random, with high probability, $\phi^{\tilde{k}\gamma}$ both oscillates at a high
 1350 frequency around 0 and explodes for $|x| \rightarrow \infty$.




Figure I.8: Same as Figure 3 but for the NNGP. In contrast to the NTK, the amplitudes of the oscillations in a. do not shrink with $\gamma \rightarrow 0$. Otherwise the behaviour is analogous. For example, the Hermite coefficients peak at $2/\gamma$. The squared coefficients sum to 1 (Eq. (6)).




Figure I.9: The spike activation components of 2-layer NTK for Gaussian spikes with varying γ (columns), $k_{max} = 1000$, top: all $s_i = +1$, middle: signs alternate every second index, bottom: 3 draws from uniformly random signs. With $\gamma \rightarrow 0$, the amplitude shrinks, while the frequency increases.

1351 Figure I.10 visualizes NNGP activation functions induced by Gaussian spikes with varying bandwidth
 1352 γ . Observe similar behaviour as for the NTK but amplitudes invariant to γ as predicted by Eq. (6).
 1353 For smaller γ the explosion of (all+) activation functions starts at larger x , but appears sharper as can
 1354 be seen in the analytic expression (H.2).

1355 Figure I.11 resembles Figure I.9 but plotted on a larger range to visualize the exploding behaviour for
 1356 $|x| \rightarrow \infty$.




Figure I.10: Spike activation components as in Figure I.9, but for the NNGP with x between $[-4, 4]$.




Figure I.11: Spike NTK activation components as in Figure I.9 but with x between $[-4, 4]$. The $+++$ NTK activation explodes exponentially. While random sign activations explode as well, $+ + - - -$ activations remain stable sin-fluctuations with slowly decaying amplitude for $|x| \rightarrow \infty$.

1357 I.6 Additive decomposition and sin-fit

1358 Here we quantify the error of the sin-approximation (8) of Gaussian NTK activation components.
 1359 The additive decomposition $\phi^k \approx \phi^{\bar{k}} + \rho^{1/2} \phi^{\bar{k}\gamma}$ quickly becomes accurate in the limit $\gamma \rightarrow 0$
 1360 (Figures I.12 and I.13), the sin-approximation seems to converge pointwise at rate $\Theta(|x|\gamma)$, where
 1361 a good approximation can be expected when $|x| \ll 1/\gamma$. The error at large $|x|$ arises because the
 1362 spike component decays for $|x| \rightarrow \infty$. For $O(1)$ inputs, we conjecture that this inaccuracy does not
 1363 dramatically affect the test error of neural networks when γ is chosen to be small.




Figure I.12: The isolated NTK spike activation function (orange), the difference between spiky-smooth and smooth activation function (blue) and a fitted sin-curve (8) (green). All curves roughly align, in particular for $\gamma \rightarrow 0$.




Figure I.13: The error of the additive decomposition $\phi^k \approx \phi^{\bar{k}} + \rho^{1/2} \phi^{\bar{k}\gamma}$ (blue) and the sin-fit (8) (orange) for the NTK. While the additive decomposition makes errors of order 10^{-3} , 10^{-4} , 10^{-9} , 10^{-15} (from left to right) in the domain $[-4, 4]$, the sin-fit is increasingly inaccurate for $|x| \rightarrow \infty$, and increasingly accurate for $\gamma \rightarrow 0$.

1364 Now we evaluate the approximation quality of the sin-fits (7) and (8) to the isolated spike activation
 1365 components $\phi^{\bar{k}\gamma}$. The NNGP oscillating activation function $\phi^{\bar{k}\gamma}$ of a Gaussian spike component is
 1366 extremely well approximated by Eq. (7). Both for the NNGP and for the NTK, the approximations
 1367 become increasingly accurate with smaller bandwidths $\gamma \rightarrow 0$ (Figure I.14). Again the approximation
 1368 quality suffers for $|x| \rightarrow \infty$, since $\phi^{\bar{k}\gamma}$ slowly decay to 0 for $|x| \rightarrow \infty$.




Figure I.14: Absolute error between the oscillating activation function $\phi^{\tilde{k}, \gamma}$ of a Gaussian spike component and the approximation **a.** Eq. (7) for the NNGP and **b.** Eq. (8) for the NTK with varying bandwidth γ .

1369 **I.7 Spiky-smooth kernel hyper-parameter selection**

1370 To understand the empirical performance of spiky-smooth kernels on finite data sets, we generate
 1371 i.i.d. data where $\mathbf{x} \sim \mathcal{U}(\mathbb{S}^d)$ and

$$y = \mathbf{x}_1 + \mathbf{x}_2^2 + \sin(\mathbf{x}_3) + \prod_{i=1}^{d+1} \mathbf{x}_i + \varepsilon,$$

1372 with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ independent of \mathbf{x} and evaluate the least squares excess risk of the minimum-norm
 1373 interpolant. Figure I.15 shows that

- 1374 • the smaller the spike bandwidth γ , the better. At some point, the improvement saturates,
- 1375 • ρ should be carefully tuned, it has large impact. As with $\gamma \rightarrow 0$ ridgeless regression with
- 1376 the spiky-smooth kernel approximates ridge regression with \tilde{k} and regularization ρ , simply
- 1377 choose the optimal regularization ρ^{opt} of ridge regression.
- 1378 • The spiky-smooth kernel with Gaussian components exhibits catastrophic overfitting, when
- 1379 γ is too large (cf. Mallinar et al. (2022)), the Laplace kernel is more robust with respect to γ .
- 1380 • With sufficiently thin spikes and properly tuned ρ , spiky-smooth kernels with Gaussian
- 1381 components outperform the Laplace counterparts.

1382 We repeat the experiment in Figure I.16 with a slightly more complex generating function and come
 1383 to the same conclusions.




Figure I.15: Least squares excess risk for spiky-smooth kernel ridgeless regression with Laplace components (*left*) and Gaussian components (*right*), with $n = 1000, d = 2$, estimated on 10000 independent test points, $\sigma^2 = 0.5, \tilde{\gamma} = 1$. The smaller the spike bandwidth γ , the better. Properly tuning ρ is important.




Figure I.16: Same as Figure I.15 but with the more complex generating function $y = |\mathbf{x}_1| + \mathbf{x}_2^2 + \sin(2\pi \mathbf{x}_3) + \prod_{i=1}^{d+1} \mathbf{x}_i + \varepsilon$. The errors are larger compared to Figure I.15 and the optimal values of ρ are smaller, but the conceptual conclusions remain the same.