On the Sample Complexity of Learning with Geometric Stability

Alberto Bietti[a], Luca Venturi[b], and Joan Bruna[a,b]

[a]Center for Data Science, New York University
[b]Courant Institute of Mathematical Sciences, New York University

June 15, 2021

Abstract

Many supervised learning problems involve high-dimensional data such as images, text, or graphs. In order to make efficient use of data, it is often useful to leverage certain geometric priors in the problem at hand, such as invariance to translations, permutation subgroups, or stability to small deformations. We study the sample complexity of learning problems where the target function presents such invariance and stability properties, by considering spherical harmonic decompositions of such functions on the sphere. We provide non-parametric rates of convergence for kernel methods, and show improvements in sample complexity by a factor equal to the size of the group when using an invariant kernel over the group, compared to the corresponding non-invariant kernel. These improvements are valid when the sample size is large enough, with an asymptotic behavior that depends on spectral properties of the group. Finally, these gains are extended beyond invariance groups to also cover geometric stability to small deformations, modeled here as subsets (not necessarily subgroups) of permutations.

1 Introduction

Learning from high-dimensional data is known to be statistically intractable without strong assumptions on the problem. A canonical example is learning Lipschitz functions, which generally requires a number of samples exponential in the dimension due to the curse of dimensionality (*e.g.*, [30]). Many high-dimensional machine learning problems involve highly structured data such as images, text, or graphs, and may exhibit invariance to certain transformations of the input data, such as permutations, translations or rotations, and near invariance to small deformations. More precisely, if $\mathcal{X}$ is the high-dimensional data domain, and G is a set of transformations $\sigma : \mathcal{X} \to \mathcal{X}$, the learning task can be alleviated if one knows in advance that the target function f varies smoothly under transformations in G: $|f(\sigma \cdot x) - f(x)|$ is uniformly small over $x \in \mathcal{X}$ for $\sigma \in G$.

To further motivate this property, it is useful to view the data domain $\mathcal{X}$ as a space of signals $\mathcal{X} = L^2(\Omega; \mathbb{R})$ defined over a geometric domain Ω, such as a 2d grid. The set of transformations G can then be articulated in terms of Ω rather than $\mathcal{X}$, a much simpler geometric object, and then *lifted* into $\mathcal{X}$ by composition: if $\sigma : \Omega \to \Omega$, and $x \in \mathcal{X}$ then $(\sigma \cdot x)(u) := x(\sigma^{-1}(u))$ for every $u \in \Omega$. The smoothness to transformations can thus be interpreted as a form of *geometric stability*.

In this paper, we quantify the sample complexity gains brought by geometric stability. Concretely, we consider target functions f defined on the sphere $\mathcal{X} = \mathbb{S}^{d-1}$ in d dimensions with finite $L^2(\mathbb{S}^{d-1})$ norm. In this case, we view the geometric domain as the discrete 1d grid $\Omega = [1, \dots, d]$, and consider geometric transformations G as subsets of the symmetric group of permutations of d elements. Given a set G (not necessarily a group), we consider the *smoothing* operator given by $S_G f(x) = \frac{1}{|G|} \sum_{\sigma \in G} f(\sigma \cdot x)$ for $f \in L^2(\mathbb{S}^{d-1})$, and assume that our target function f is geometrically stable, in the sense that $f = S_G g$ for some $g \in L^2(\mathbb{S}^{d-1})$. In words, the smoothing operator S_G replaces the prediction $f(x)$ by the average over

transformations of x. In particular, functions which are invariant under the action of $\sigma \in G$, namely

$$f(\sigma \cdot x) = f(x), \quad \sigma \in G, x \in \mathbb{S}^{d-1}, \tag{1}$$

are also stable, with $f = S_G f$.

Building on the recent work [23], we proceed by studying harmonic decompositions of such functions using spherical harmonics [14], which generalize Fourier series on the circle to higher dimensions. This allows us to obtain rates of approximation for invariant and geometrically stable functions with varying levels of smoothness, and to study the generalization properties of invariant kernel methods using kernels defined on the sphere. Specifically, our main contributions are:

- By comparing spectral properties of usual kernels on the sphere with invariant ones, we find that the latter provide improvements in sample complexity by a factor of the order of the size of the group when the sample size is large enough (Section 3).
- We study how this improvement factor varies with sample size, in terms of the structure of the group and on spectral properties of the permutation matrices it contains (Section 4).
- We extend the invariance analysis to geometrically stable functions, establishing similar gains in sample complexity that depend on the size of the transformation subset (Section 5).

Our proofs rely on comparing the dimension of invariant and non-invariant spherical harmonics at a given degree, and showing that their ratio decays to the inverse group size as the degree tends to infinity. In contrast to [23], we consider the dimension to be fixed and study non-parametric rates of convergence for potentially non-smooth target functions and general groups of permutations, while they consider high-dimensional limits where only polynomials can be learned, and focus on invariance to translation groups.

Related work. Invariance and deformation stability have been analysed using convolutional neural network-type architectures such as the scattering transform [22, 7], or convolutional kernels [6, 20]. While these works characterise the stability in terms of the dyadic structure of convolutional filters (such as wavelets), they do not cover a statistical analysis of sample complexity. Similarly, models of compositional functions such as those in [11, 24, 27] study the benefit of hierarchical representations with local connectivity for approximation, while [19, 21] study benefits of local connectivity with optimization-based algorithms; yet these works do not consider invariance or stability. [23] studies similar benefits but focus on polynomial target functions, and consider high-dimensional limits, while we consider a fixed dimension and obtain non-parametric rates of convergence. While their framework is potentially applicable to general groups, their results only cover the case of translation groups, leading to improvements in sample complexity of at most a factor equal to the dimension. [15] also sudies benefits of group invariance, but focuses on linear models, and only considers interpolating estimators. [10] study benefits of equivariant kernels in structured prediction problems. [13] studies generalisation advantages of CNNs over fully-connected models, while our focus is on non-parametrics.

2 Preliminaries

In this section, we describe our setup and provide some background on harmonic decompositions on the sphere, and how these are affected by invariance.

Statistical learning setup. We consider a supervised learning problem where the data distribution ρ on input-label pairs (x, y) is such that $x \in \mathbb{S}^{d-1}$ and $\mathbb{E}[y|x] = f^*(x)$ for some target function f^*in $L^2(\mathbb{S}^{d-1})$. For simplicity, we will assume that x is uniformly distributed on the sphere, and denote the uniform measure on $\mathbb{S}^{d-1}$ by $d\tau$. We consider a regression setup with L^2 risk given by

$$R(f) = \mathbb{E}_{(x,y)\sim\rho}\left[(f(x) - y)^2\right].$$

For a given estimator $\hat{f}_n$ based on n samples from ρ, the goal is then to obtain generalization bounds as a function of n on the excess risk

$$\mathbb{E}[R(\hat{f}_n)] - R(f^*) = \mathbb{E}[\|\hat{f}_n - f^*\|^2_{L^2(d\tau)}], \tag{2}$$

where the expectation is over the n samples. Such bounds are well-studied for various classes of target functions f^* such as smoothness classes, and estimators such as kernel ridge regression. These are typically studied through harmonic decompositions of f^* and of a kernel function in appropriate L^2 bases, which then relate function regularity and decays of Fourier coefficients.

Harmonic analysis on the sphere. When considering functions in $L^2(d\tau)$, an appropriate choice of orthonormal basis is that of spherical harmonic polynomials [1, 14]. More precisely, denote $\{Y_{k,j}\}_{j=1}^{N(d,k)}$ denote an orthonormal basis of the space $V_{d,k}$ of spherical harmonics of degree k, *i.e.*, homogeneous harmonic polynomials of degree k, where $N(d,k) = \frac{2k+d-2}{k}\binom{k+d-3}{d-2}$. Then, the collection $\{Y_{k,j} : k \geq 0, j = 1, \dots, N(d,k)\}$ forms an orthonormal basis of $L^2(d\tau)$, so that any function $f \in L^2(d\tau)$ may be written

$$f(x) = \sum_{k \geq 0} \sum_{j=1}^{N(d,k)} a_{k,j} Y_{k,j}(x), \tag{3}$$

with $\sum_k \sum_{j=1}^{N(d,k)} a_{k,j}^2 < \infty$. Similarly, any dot-product kernel $K(x,x') = \kappa(\langle x, x' \rangle)$ on the sphere may be written

$$\kappa(\langle x, x' \rangle) = \sum_{k \geq 0} \mu_k \sum_{j=1}^{N(d,k)} Y_{k,j}(x) Y_{k,j}(x'), \tag{4}$$

where μ_k is given by

$$\mu_k = \frac{\omega_{d-2}}{\omega_{d-1}} \int_{-1}^{1} \kappa(t) P_{d,k}(t) (1-t^2)^{(d-3)/2} dt. \tag{5}$$

Here, ω_{p-1} is the surface measure of the sphere in p dimensions, and $P_{d,k}$ are Legendre or Gegenbauer polynomials of degree k in d dimensions (normalized with $P_{d,k}(1) = 1$), which form an orthogonal basis of $L^2([-1,1], dw)$, with $dw(t) = (1-t^2)^{(d-3)/2} dt$. When the kernel K is positive definite and is used in the context of kernel ridge regression with data uniformly distributed on the sphere, then the μ_k also correspond to the eigenvalues of the covariance operator. These eigenvalues and their decay then control the statistical properties of the kernel ridge regression estimator [8].

Spherical harmonics and group-invariant functions. In order to describe harmonic decompositions of functions satisfying the group invariance property (1) for a discrete group G, we follow [23] and define the symmetrization operator

$$S_G f(x) = \frac{1}{|G|} \sum_{\sigma \in G} f(\sigma \cdot x). \tag{6}$$

This operator acts as a projection from $L^2(d\tau)$ to a subset thereof which contains invariant functions. It can be shown that the spaces $V_{d,k}$ of spherical harmonics of degree k are stable by S_G [23], and we may then define an orthonormal basis of $\overline{V}_{d,k} := S_G V_{d,k}$ consisting of *invariant* spherical harmonics $\{\overline{Y}_{k,j}\}_{j=1}^{\overline{N}(d,k)}$. We then have the following lemma.

Lemma 1 (Representation of projection [23])**.** *For any $k \geq 0$, we have*

$$\gamma_d(k) := \frac{\overline{N}(d,k)}{N(d,k)} = \frac{1}{|G|} \sum_{\sigma \in G} \mathbb{E}_x \left[P_{d,k}(\langle \sigma \cdot x, x \rangle) \right]. \tag{7}$$

The quantity $\gamma_d(k)$ will play an important role in determining the gains in sample complexity brought by invariance. We will show in Section 4 that $\gamma_d(k)$ converges to $1/|G|$ for large k, with an asymptotic behavior that is governed by spectral properties of the elements of the group.

3 Sample Complexity of Invariant Kernels

We begin our study by focusing on the invariant case. In this section, we study the sample complexity of learning invariant functions, by considering kernel ridge regression estimators and providing non-parametric rates of convergence that illustrate the gains achievable with invariant kernels compared to non-invariant ones.

Kernel ridge regression (KRR) and invariant kernels. For a positive definite kernel K with RKHS $\mathcal{H}_K$, we consider the KRR estimator $\hat{f}_\lambda$ given by

$$\hat{f}_\lambda := \arg\min_{f \in \mathcal{H}_K} \frac{1}{n} \sum_{i=1}^{n} (f(x_i) - y_i)^2 + \lambda \|f\|_{\mathcal{H}_K}^2. \tag{8}$$

We consider the following kernels, which we assume positive definite, given for $x, x' \in \mathbb{S}^{d-1}$ by

$$K(x, x') = \kappa(\langle x, x' \rangle), \qquad K_G(x, x') = \frac{1}{|G|} \sum_{\sigma \in G} \kappa(\langle \sigma \cdot x, x' \rangle), \tag{9}$$

with $\kappa(u) \leq 1$. A common example for κ is the arc-cosine kernel [9], which arises from infinite-width shallow neural networks with ReLU activations. The following integral operator defined on $L^2(d\tau)$ and its eigen decomposition play an important role for the statistical and approximation properties of kernel methods:

$$T_K f(x) = \int K(x, x') f(x') d\tau(x'). \tag{10}$$

We now show that its spectral properties are closely related for K and K_G.

Lemma 2 (Spectral properties of K and K_G.)**.** *There exists a basis of spherical harmonics in which the operators T_K and T_{K_G} are jointly diagonalized. They admit the same eigenvalues μ_k given by* (5)*, with multiplicity $N(d, k)$ for T_K and $\overline{N}(d, k)$ for T_{K_G}.*

The decay of the eigenvalues μ_k controls the smoothness of functions in the RKHS, for instance when μ_k decays polynomially, T_K behaves similarly to powers of the Laplacian on the sphere, leading to functional spaces similar to Sobolev spaces. For the example of the arc-cosine kernel, μ_k decays as k^{-d-2}. leading to an RKHS containing functions with $d/2 + 1$ bounded derivatives [2].

Approximation error. The approximation error of kernel methods is often controlled by the following quantity (*e.g.*, [3, 12]):

$$A_{\mathcal{H}}(\lambda, f^*) = \inf_{f \in \mathcal{H}} \|f - f^*\|_{L^2(d\tau)}^2 + \lambda \|f\|_{\mathcal{H}}^2, \tag{11}$$

where f^* is a target function in $L^2(d\tau)$, and $\mathcal{H}$ is a given RKHS. In particular, if f^* is smooth enough so that $f^* \in \mathcal{H}$, then we have $A_{\mathcal{H}}(\lambda, f^*) \leq \lambda \|f^*\|_{\mathcal{H}}^2$, while if $f^* \notin \mathcal{H}$, *e.g.*, if f^* is only Lipschitz, then $A_{\mathcal{H}}(\lambda, f^*)$ typically grows much faster with λ. We now show a useful result for invariant targets, showing that in this case the approximation error is the same for the kernels K and K_G.

Lemma 3 (Approximation error for invariant functions.)**.** *If f^* is invariant to the group G, so that $f^* = S_G f^*$, then we have*

$$A_{\mathcal{H}_K}(\lambda, f^*) = A_{\mathcal{H}_{K_G}}(\lambda, f^*). \tag{12}$$

Degrees of freedom. The above result suggests that any gains of using K_G instead of K for learning invariant functions should come from estimation rather than approximation error. The estimation error of ridge rigression estimators is typically controlled with the following quantity, often called *degrees of freedom* or *effective dimension* (*e.g.*, [3, 18]):

$$\mathcal{N}_K(\lambda) = \mathrm{Tr}(\Sigma_K (\Sigma_K + \lambda I)^{-1}) = \sum_{m \geq 0} \frac{\lambda_m}{\lambda_m + \lambda}, \tag{13}$$

where $\Sigma_K = \mathbb{E}_x\left[K(x,\cdot) \otimes_{\mathcal{H}_K} K(x,\cdot)\right]$ is the covariance operator and $(\lambda_m)_{m\geq 0}$ its eigenvalues, taking multiplicity into account, which are the same as those of T_K when data is distributed according to $d\tau$ [8]. We then obtain the following simple result relating $\mathcal{N}_{K_G}$ to $\mathcal{N}_K$.

Lemma 4 (Degrees of freedom for K and K_G.)**.** *For any $\ell \geq 0$, we have*

$$\mathcal{N}_{K_G}(\lambda) \leq D(\ell) + \nu_d(\ell)\mathcal{N}_K(\lambda),$$

where $D(\ell) := \sum_{k<\ell} \overline{N}(d,k)$ and $\nu_d(\ell) := \sup_{k\geq\ell} \gamma_d(k)$, with γ_d given in (7)*.*

This suggests that for a fixed ℓ, the effective dimension of K_G is controlled by a factor $\nu_d(\ell)$ times that of K, up to a finite fixed dimension $D(\ell)$. For difficult non-parametric problems which require small λ at large sample sizes, the second term will tend to dominate, so that having a small $\nu_d(\ell)$ may help reduce sample complexity compared to using the vanilla kernel K, an observation which we make rigorous below.

Generalization bound for KRR. Armed with the above lemmas on approximation error and degrees of freedom, we now study generalization of KRR under the following assumptions:

(A1) *capacity condition*: $\mathcal{N}_K(\lambda) \leq C_K \lambda^{-1/\alpha}$ with $\alpha > 1$.

(A2) *source condition*: there exists $r > \frac{\alpha-1}{2\alpha}$ and $g \in L^2(d\tau)$ with $\|g\|_{L^2(d\tau)} \leq C_{f^*}$ such that $f^* = T_K^r g$.

(A3) *invariance*: f^* is G-invariant.

(A4) *problem noise*: ρ is such that $\mathbb{E}_\rho[(y - f^*(x))^2|x] \leq \sigma_\rho^2$.

The first, second, and fourth conditions are commonly used in the kernel methods literature [8]. (A1) characterizes the "size" of the RKHS, and is satisfied when the eigenvalues λ_m of T_K decay as $k^{-\alpha}$. On the sphere, $\alpha = \frac{2s}{d-1}$ corresponds to having s bounded derivatives, *e.g.*, we have $s = d/2 + 1$ for the arc-cosine kernel. The parameter r in (A2) defines the regularity of f^* relative to that of the kernel: $r = 1/2$ corresponds to $f^* \in \mathcal{H}_K$, while larger (resp. smaller) r implies f^* is more (resp. less) smooth. The condition on r is needed for our specific bound, which is based on [3, Proposition 7.2], but may be bypassed using different algorithms or analyses [17, 26]. We now present our bound on the excess risk.

Theorem 5 (Generalization of invariant kernel.)**.** *Assume (A1-4). Let $\nu_d(\ell)$ be as in Lemma 4, or an upper bound thereof, and assume $\nu_0 := \inf_{\ell\geq 0} \nu_d(\ell) > 0$. Let*

$$n \geq \max\left\{ \|f^*\|_\infty^2/\sigma_\rho^2, (C_1/\nu_0)^{\frac{\alpha}{2\alpha r+1-\alpha}} \right\}, \tag{14}$$

and define

$$\ell_n := \sup\{\ell : D(\ell) \leq C_2 \nu_d(\ell)^{\frac{2\alpha r}{2\alpha r+1}} n^{\frac{1}{2\alpha r+1}}\}. \tag{15}$$

We then have, for $\lambda = C_3(\nu_d(\ell_n)/n)^{\alpha/(2\alpha r+1)}$,

$$\mathbb{E}[R(\hat{f}_\lambda) - R(f^*)] \leq C_4 \left(\frac{\nu_d(\ell_n)}{n}\right)^{\frac{2\alpha r}{2\alpha r+1}}. \tag{16}$$

In the same setting, KRR with kernel K and $\lambda = C_3 n^{\frac{-\alpha}{2\alpha r+1}}$ achieves $\mathbb{E}[R(\hat{f}_\lambda) - R(f^*)] \leq C_4 n^{\frac{-2\alpha r}{2\alpha r+1}}$, where C_3, C_4 are the same constants as for the invariant kernel. Here, the constants $C_{1:4}$ only depend on the parameters of assumptions (A1-4).

The theorem shows that the generalization error for the invariant kernel behaves as if it effectively had access to $n/\nu_d(\ell_n)$ samples, so that $\nu_d(\ell_n)$ plays the role of an *effective* inverse sample complexity gain at sample size n. Note that $\nu_d(\ell_n) \leq 1$ and ν_d is decreasing, so that we always have some improvement in sample complexity, and this gets better when ℓ_n is large. In particular, we show in Section 4 that$\gamma_d(k)$, and hence $\nu_d(\ell)$ converge to $1/|G|$, so that asymptotically the gain in sample complexity can be as large as the size of the group, which in some cases may grow *exponentially* in d.

Estimating ℓ_n and the effective gain $\nu_d(\ell_n)$. Since $D(\ell)$ increases with ℓ, Eq. (15) suggests that ℓ_n increases with n. We now provide intuition on how we might expect ℓ_n and $\nu_d(\ell_n)$ to behave in a situation of interest where we know an asymptotic equivalent of ν_d. Namely, assume that we have

$$\nu_d(\ell) \approx \nu_0 + c\ell^{-\beta}.$$

We provide such asymptotic equivalents in Section 4, where $\nu_0 = 1/|G|$, and β depends on spectral properties of the elements of G. For some large groups, β may be small, in which case we may consider other approximations with larger β, at the cost of a larger ν_0. When ℓ_n is large, using the approximation $N(d,k) \approx k^{d-2}$, we have $D(\ell) \approx \sum_{k=0}^{\ell-1} k^{d-2} \approx \ell^{d-1}$. Hiding constants other than ν_0, we may then consider ℓ_n to be solution of

$$\ell^{\frac{(d-1)(2\alpha r+1)}{2\alpha r}} = n^{\frac{1}{2\alpha r}}(\nu_0 + \ell^{-\beta}).$$

Since the l.h.s. increases, while the r.h.s. decreases with ℓ, we must have $\ell_n \geq \max(\ell_{n,1}, \ell_{n,2})$, with

$$\ell_{n,1}^{\frac{(d-1)(2\alpha r+1)}{2\alpha r}} = n^{\frac{1}{2\alpha r}}\nu_0, \quad \text{and} \quad \ell_{n,2}^{\frac{(d-1)(2\alpha r+1)}{2\alpha r}} = n^{\frac{1}{2\alpha r}}\ell_{n,2}^{-\beta}.$$

This yields

$$\nu_d(\ell_n) \leq \nu_0 + C\min\left\{(\nu_0^{2\alpha r}n)^{\frac{-\beta}{(d-1)(2\alpha r+1)}}, n^{\frac{-\beta}{(d-1)(2\alpha r+1)+2\beta\alpha r}}\right\}. \tag{17}$$

Notice that when $\beta \ll d$, both exponents of n display a curse of dimensionality, but this curse goes away as β grows. Note also that the first exponent yields a faster rate, but one that is only achieved for large n due to the factor $\nu_0^{2\alpha r}$, which may be small for large groups.

Curse of dimensionality and optimality. Note that the bound obtained in Theorem 5 is still cursed for an invariant target f^*, in the sense that the exponent in the rate is of order $1/d$ when f^* is only assumed to be Lipschitz. Indeed, a Lipschitz assumption on f^* corresponds to taking r and α such that $2\alpha r \approx 2/(d-1)$, which makes the source condition (A2) similar to a bound on $\|\nabla f^*\|_{L^2(d\tau)}$. This then leads to a cursed rate $n^{-2/(2+d-1)}$, raising the question of whether this can be improved. We note that since $\gamma_d(k) = \Omega(1/|G|)$ (as we show in Section 4), the asymptotic decays (as a function of k) of the coefficients of f^* and of the eigenvalues of T_{K_G} are similar to those for the non-invariant case, which implies that these rates cannot be improved (see, *e.g.*, [8]), and that constant-factor improvements in the sample complexity may be the best we can hope for. Nevertheless, this may still lead to significant gains in practice.

Interpretation as Lipschitz under invariant metric. When assuming that f^* is Lipschitz, our assumptions can be interpreted as a 'geometric' Lipschitz assumption on the target function. Indeed, combining assumptions (A2) and (A3), together with an r such that $2\alpha r \approx 2/(d-1)$, which corresponds to a Lipschitz condition on f^* as argued in the previous paragraph, we may interpret these as saying that f^* is Lipschitz with respect to a weaker G-invariant (semi)-norm:

$$|f^*(x) - f^*(x')| \leq B \inf_{\sigma\in G} \|\sigma \cdot x - x'\| .$$

In this setting, our main result, Theorem 5, establishes that the gains brought by this stronger Lipschitz condition do not generally improve the rate (which stays cursed by dimension; see previous paragraph), but nevertheless may be important in practice thanks to the potentially exponential-in-d improvements in the effective sample size.

4 Counting Invariant Polynomials

In this section, we study the asymptotic behavior of $\gamma_d(k)$, given in (7), when $k \to \infty$ and the dimension d and the group G are fixed. This quantity can be seen as capturing the fraction of orthogonal spherical harmonics of degree k that are invariant to G, and helps us control the possible gains in sample complexity for learning invariant functions, as described in Section 3. Denoting $\gamma_{d,\sigma}(k) := \mathbb{E}_x\left[P_{d,k}(\langle\sigma\cdot x, x\rangle)\right]$, we will show that $\gamma_{d,\sigma}(k)$ vanishes for large k for any σ that is not the identity. This implies that $\gamma_d(k)$ converges

to $1/|G|$, since we trivially have $\gamma_{d,Id}(k) = P_{d,k}(1) = 1$. We further characterize the asymptotic behavior of $\gamma_d(k)$ in terms of properties of the group elements. In the following we consider the case of G being a subgroup of S_d, the groups of permutations on d elements.

Decay of $\gamma_{d,\sigma}(k)$. Our main insight is to leverage the fact that when σ is not the identity, then the random variable $z_\sigma = \langle \sigma \cdot x, x \rangle$ when $x \sim \tau$ admits a density on $[-1, 1]$, which we denote q_σ. This by itself will prove sufficient to show that $\gamma_{d,\sigma}(k)$ decays for large k, thanks to the oscillatory behavior of $P_{d,k}$. We can then further characterize its asymptotic behavior by studying the singularities of q_σ, leveraging the seminal work of Saldanha and Tomei [28]. In particular, these depend on spectral properties of the matrix associated to σ. We summarize this in the next proposition.

Proposition 6 (Asymptotic behavior of $\gamma_{d,\sigma}(k)$.)**.** *Let A_σ be the matrix associated to $\sigma \neq \mathrm{Id}$, that is such that $\sigma \cdot x = A_\sigma x$. Denote by Λ_σ the set of (complex) eigenvalues of A_σ, and by m_λ the multiplicity of $\lambda \in \Lambda_\sigma$. When $k \to \infty$, we have the asymptotic equivalent $\gamma_{d,\sigma}(k) = \sum_{\lambda \in \Lambda_\sigma} \gamma_{d,\sigma,\lambda}(k)$, where*

$$\gamma_{d,\sigma,\lambda}(k) \lesssim \begin{cases} k^{-d+m_\lambda} + o(k^{-d+m_\lambda}) \,, & \text{if } \lambda \in \{\pm 1\} \,, \\ k^{-d+m_\lambda+4} + o(k^{-d+m_\lambda+4}) \,, & \text{otherwise}, \end{cases} \tag{18}$$

where $\lesssim$ hides constants that may depend on d, σ and λ.

Every permutation $\sigma \in \mathcal{S}_d$ (where $\mathcal{S}_d$ is the symmetric group of permutations) can be decomposed into cycles on disjoint orbits; the eigenvalues λ (and their multiplicities m_λ) of a matrix A_σ admit an interpretation based on such decomposition. Indeed, since A_σ is unitary, its eigenvalues are of the form $\lambda = e^{2\pi i \theta}$, and one can verify that necessarily $\theta = \frac{p}{q} \in \mathbb{Q}$. Furthermore, assuming w.l.o.g. that q is prime, such eigenvalue appears whenever σ contains a cycle of length a multiple of q. In particular, the multiplicity of the eigenvalue 1, m_1, corresponds to the total number of cycles in such a decomposition, which we will denote by $c(\sigma)$. Then $\gamma_{d,\sigma}(k)$ can be controlled as follows.

Corollary 7 (Decay of $\gamma_{d,\sigma}(k)$)**.** *Let $\sigma \neq \mathrm{Id}$, and let $c(\sigma)$ denote the number of cycles in σ. Then,*

$$\gamma_{d,\sigma}(k) \lesssim \begin{cases} k^{-d+c(\sigma)} \,, & \text{if } c(\sigma) > \frac{d+3}{2} \,, \\ k^{-d/2+6} \,, & \text{otherwise} \,. \end{cases}$$

Decay on specific subgroups. We may now use Corollary 7 to study the asymptotic behaviour of $\gamma_d(k)$ as $k \to \infty$ for various choices of subgroups of $\mathcal{S}_d$, using the following result.

Corollary 8 (Upper bounds with permutation statistics)**.** *Let G be a subset of $\mathcal{S}_d$ and define*

$$\zeta(G, s) := |\{\sigma \in G \ : \ c(\sigma) > s\}|$$

for $s \in [d-1]$. Then, for any s, we have

$$\gamma_d(k) \leq \frac{\zeta(G, s)}{|G|} + O\left(k^{-d+\max\{s,\, d/2+6\}}\right) \,, \tag{19}$$

with equality if s is such that $\zeta(G, s) = 1$.

Note that such an upper bound immediately yields a similar upper bound for $\nu_d(\ell)$ as defined in Section 3, which then controls the effective gain in sample complexity in Theorem 5. Indeed, (19) implies that there is a constant C such that for all $k > 0$, $\gamma_d(k) \leq \zeta(G, s)/|G| + Ck^{-d+\max\{s, d/2+6\}}$. Since this upper bound decreases with k, we obtain

$$\nu_d(\ell) \leq \frac{\zeta(G, s)}{|G|} + C\ell^{-d+\max\{s, d/2+6\}}.$$

In the context of the generalization bound of Theorem 5 and our heuristic derivation thereafter, the *effective* gain in sample complexity is then governed an upper bound on $\nu_d(\ell_n)$ as in (17), with asymptotic gain ν_0 and rate β given by

$$\nu_0 := \frac{\zeta(G, s)}{|G|} \quad \text{and} \quad \beta := d - \max\{s, d/2 + 6\}. \tag{20}$$

Example 9 (Translations). *Let $G = C_d$ be the cyclic group on d elements. Then it holds*

$$\gamma_d(k) = \frac{1}{d} + O\left(k^{-d/2+6}\right) .$$

This follows by noticing that every translation σ (but the identity) satisfies $c(\sigma) \leq d/2$. This leads to an asymptotic gain $\nu_0^{-1} = d$ and $\beta = d/2 - 6$ leads to fast convergence in (17) even when d is large.

Example 10 (Local translations). *Let $d = s \cdot r$ (with $r, s \geq 5$ for simplicity), and consider the group composed of traslations over r blocks of coordinates of size s;* i.e., *the block-cyclic group*

$$G = \{\sigma \ : \ \sigma = \sigma^{(1)} \circ \cdots \circ \sigma^{(r)}\}$$

where each $\sigma^{(i)}$ is a translation over the set $\{(i-1)s+1, \dots, is\}$, for $i \in [r]$. Then it holds

$$\gamma_d(k) = \frac{1}{s^r} + O\left(k^{-s/2+1}\right) . \tag{21}$$

This follows by noticing that every local translation σ (but the identity) satisfies $c(\sigma) \leq (d-s) + s/2$. Here the asymptotic gain is $\nu_0^{-1} = s^r = s^{d/s}$, which can be exponential in d when s is small. We have $\beta = s/2 - 1$, which leads to much slower convergence than the translation case, unless s is large and of order d.

Example 11 (Full permutation group). *For the case of $G = \mathcal{S}_d$, we can split the group based on the value of $\mathrm{Fix}(\sigma)$, the number of elements fixed by a permutation σ. Denote*

$$\xi(G, s) := |\sigma \in G \ : \ \mathrm{Fix}(\sigma) > s| = \sum_{j=s+1}^{d} \binom{d}{j} \, !(d-j) \, .$$

for $s \in [d-1]$, where $!k$ denotes the k-th subfactorial. Then we have

$$\gamma_d(k) \leq \frac{\xi(G, s)}{d!} + O\left(k^{-d/2+\max\{s/2,\, 6\}}\right) ,$$

with equality for $s = d-1$. This follows from the fact that $c(\sigma) \leq \mathrm{Fix}(\sigma) + (d - \mathrm{Fix}(\sigma))/2$. In particular, it follows

$$\gamma_d(k) \leq \frac{2}{(s+1)!} + O\left(k^{-d/2+\max\{s/2,\, 6\}}\right) .$$

When considering the full group, we may get a large asymptotic improvement of order $\nu_0^{-1} = |G| = d!$ in sample complexity, but a slow convergence with $\beta = -1$ as per Corollary 8 (assuming d large enough). Using different values of s may yield different upper bounds with faster convergence rates $\beta = d/2 - \max\{s/2, 6\}$, but smaller asymptotic gains in sample complexity, given by $\nu_0^{-1} = (s+1)!/2$. For instance, with $s = d/2$ and d large enough, we have $\beta = d/4$, leading to a potentially fast convergence rate in 17 towards a sample complexity gain that is still significantly large, of order $(d/2+1)!/2$.

Overall, these examples show that the size of the group determines the best possible improvement in sample complexity, while the spectral properties of its permutations dictate how quickly we may achieve these gains.

5 Beyond Group Invariance: Geometric Stability

In this section, we study gains in sample complexity when the target function f^* is not fully invariant to a group, but may be stable under small geometric changes on the input. We formalize this by considering a similar averaging operator S_G, but we allow G to be a generic set of permutations instead of a group, and allow for a weighted average:

$$S_G f(x) := \sum_{\sigma \in G} h(\sigma) f(\sigma \cdot x), \tag{22}$$

where $h(\sigma) \geq 0$ for all $\sigma \in G$ and $\sum_{\sigma \in G} h(\sigma) = 1$. We assume that G is "symmetric", *i.e.*, $\sigma^{-1} \in G$ when $\sigma \in G$ and $h(\sigma^{-1}) = h(\sigma)$, so that S_G is self-adjoint. In this case, images of S_G are not invariant functions, but may nevertheless exhibit a form of "local" stability to small perturbations of the input data. For instance, if G consists of local translations by at most a few pixels, or if G consists of all translations but h is localized around the identity, then applying S_G yields functions that are stable to local translations. We may also consider a more structured set G of permutations that resemble local deformations, consisting of both a global translation as well as different local translations at different scales, as we describe below.

Spectral properties of S_G. Note that in this setup, we no longer have that S_G is a projection, however we may still view it as a *smoothing* operator, which attenuates certain harmonics that are "less" invariant than others. The next lemma shows related spectral properties to the invariant case.

Lemma 12 (Spectral properties of S_G)**.** *There exists a basis of spherical harmonics $\overline{Y}_{k,j}$, for $k \geq 0$, and $j = 1, \dots, N(d,k)$, in which the operator S_G is diagonal, with eigenvalues $\lambda_{k,j} \geq 0$. In analogy to Lemma 1, we have*

$$\gamma_d(k) := N(d,k)^{-1} \sum_{j=1}^{N(d,k)} \lambda_{k,j} = \sum_{\sigma \in G} h(\sigma) \mathbb{E}_x \left[P_{d,k}(\langle \sigma \cdot x, x \rangle) \right]. \tag{23}$$

We also define $\nu_d(\ell) := \sup_{k \geq \ell} \gamma_d(k)$.

Sample complexity of stable kernel. In analogy to Section 3, we may consider a stable kernel

$$K_G(x, x') = \sum_{\sigma \in G} h(\sigma) \kappa(\langle \sigma \cdot x, x' \rangle). \tag{24}$$

Then, it is easy to check that the integral operator of K_G is given by $T_{K_G} = S_G T_K$. In contrast to Section 3, we no longer have that the approximation errors of K and K_G are the same in general on "geometrically stable" functions, since the notion is not precisely defined. Nevertheless, we may represent favorable targets f^* as those whose coefficients decay similarly at each frequency k to those of S_G, by viewing it as a smoothing of some L^2 function g^*, *i.e.*, $f^* = S_G^r g^*$ for some exponent r. With this in mind, we make the following assumptions, replacing assumptions (A1-3) of Section 3.

(A5) *capacity*: the eigenvalues $(\xi_m)_{m \geq 0}$ of T_K satisfy $\xi_m \leq C(m+1)^{-\alpha}$.

(A6) *source condition*: there exists $r > \frac{\alpha - 1}{2\alpha}$ and $g \in L^2(d\tau)$ with $\|g\|_{L^2(d\tau)} \leq C_{f^*}$ such that $f^* = S_G^r T_K^r g$.

Note that (A6) corresponds to a standard source condition with the kernel K_G (since $T_{K_G}^r = S_G^r T_K^r$), yet it reveals how K_G jointly performs smoothing on the sphere, through T_K, as well as on the permutation group through S_G. While these two forms of smoothing appear "entangled" in this assumption, one may balance them by choosing different levels of smoothing in the kernel function κ, or by averaging multiple times in (24). Assumption (A5) is needed for obtaining a variant of Lemma 4, and implies (A1) with $C_K \propto C^{1/\alpha}$. We then obtain the following generalization bound.

Theorem 13 (Generalization with geometric stability.)**.** *Assume (A4-6), and assume $\nu_0 := \inf_{\ell \geq 0} \nu_d(\ell) > 0$. Let $n \geq \max(\|f^*\|_\infty^2 / \sigma_\rho^2, (C_1/\nu_0)^{1/(2\alpha r + 1 - \alpha)})$, and define*

$$\ell_n := \sup\{\ell : D(\ell) \leq C_2 \nu_d(\ell)^{\frac{2r}{2\alpha r+1}} n^{\frac{1}{2\alpha r+1}}\}. \tag{25}$$

We then have, for $\lambda = C_3 (\nu_d(\ell_n)^{1/\alpha}/n)^{\alpha/(2\alpha r+1)}$,

$$\mathbb{E}[R(\hat{f}_\lambda) - R(f^*)] \leq C_4 \left(\frac{\nu_d(\ell_n)^{1/\alpha}}{n} \right)^{\frac{2\alpha r}{2\alpha r+1}}. \tag{26}$$

In the same setting, KRR with kernel K obtains the same bound with $\nu_d(\ell_n)^{1/\alpha}$ replaced by 1. Here, the constants $C_{1:4}$ only depend on the parameters of assumptions (A4-6).

Note that the theorem is very similar to Theorem 5, albeit with a gain in sample complexity dictated by $\nu_d(\ell_n)^{1/\alpha}$ instead of $\nu_d(\ell_n)$. This is due to the fact that in contrast to the invariant case, where $\gamma_d(k)$ in (7) can help precisely control the number of invariant spherical harmonics, in this case $\gamma_d(k)$ as computed in (23) can only give information about the sum of the eigenvalues $\lambda_{k,j}$ at frequency k, which may be insufficient to precisely estimate the gains in effective dimension. For smooth kernels with fast decays (large α), this discrepancy can weaken the sample complexity gains shown by our bounds, but for slow decays ($\alpha \approx 1$), the gains are similar to the invariant case. We note that in this case we may similarly derive an estimate of $\nu_d(\ell_n)$ as in (17), namely if $\nu_d(\ell) \approx \nu_0 + c\ell^{-\beta}$, then we have

$$\nu_d(\ell_n) \leq \nu_0 + C \min\left\{ (\nu_0^{2r} n)^{\frac{-\beta}{(d-1)(2\alpha r+1)}}, n^{\frac{-\beta}{(d-1)(2\alpha r+1)+2\beta r}} \right\}. \tag{27}$$

Deformation-like stability. For inputs x defined as signals $x \in L^2(\Omega)$ over a continuous domain $\Omega \subseteq \mathbb{R}^s$, $s = 1, 2$, the action of 'small' diffeomorphisms $\varphi : \Omega \to \Omega$ as $(\varphi \cdot x)(u) = x(\varphi^{-1}(u))$ is a powerful diagnostic of performance of trainable CNNs [25], and a key guiding principle for scattering representations [22, 7]. In these works, the basic deformation cost is measured as $\|\varphi\| := \sup_u \|\nabla\varphi(u) - \mathbf{I}\|$. We instantiate an equivalent of small deformations in our finite-dimensional setting as follows.

$$\Phi_\varepsilon := \{\sigma \in \mathcal{S}_d \ : \ |\sigma(u) - \sigma(u') - (u - u')| \leq \varepsilon |u - u'|\}, \tag{28}$$

where the differences are taken modulo d. For $\varepsilon = 0$, we recover the translation group described in Example 9. We can verify that $\varepsilon = 1$ also corresponds to the translation group (due to the constraint that $\sigma(u) \neq \sigma(u')$ whenever $u \neq u'$), thus the first non-trivial model corresponds to $\varepsilon = 2$.

Proposition 14 (Upper bound on $\gamma_d(k)$ for deformations.)**.** *It holds $\Phi_2^{-1} = \Phi_2$. Moreover, $|\Phi_2| \geq \tau^d$ for $\tau \approx 1.714$, and*

$$\gamma_d(k) \leq C \left(\frac{e^{2\eta}}{(2\eta)^{2\eta} \tau^{1-2\eta}} \right)^d + O\left(k^{-\eta d}\right) \tag{29}$$

for $\eta < 1/4$. In particular, $\tilde{\tau}^{-1} := \frac{e^{2\eta}}{(2\eta)^{2\eta}\tau^{1-2\eta}} < 1$ for $\eta < 0.07$, leading to an effective gain in sample complexity exponential in d, $\nu_0^{-1/\alpha} = \Theta(\tilde{\tau}^{d/\alpha})$; and $\beta = \eta d$ resulting in fast convergence in (27) even for large d.

We thus verify that small deformations, already with $\varepsilon = 2$, provide a substantial gain relative to rigid translations, since Φ_2 now grows exponentially with the dimension, rather than linearly. Let us remark that our small deformation model (28) acting on $\{1, d\}$ differs in important ways from diffeomorphisms acting on a continuous domain. In our case they define unitary operators (since they are constructed as subsets of the permutation group), as opposed to diffeorphims, for which $\|\varphi \cdot x\|_{L^2(\Omega)} \neq \|x\|_{L^2(\Omega)}$ generally. In other words, the 'deformations' in Φ_ε are more akin to local shufflings of the pixels rather than local distortions. That said, our model does roughly capture the size of small deformation classes. An interesting question for future work is to extend our framework to non-unitary transformations, which could accommodate appropriate discretisations of continuous diffeomorphisms.

6 Numerical Experiments

In this section, we provide simple numerical experiments on synthetic data which illustrate our theoretical results. In Figure 1, we consider KRR on 5 000 training samples with inputs uniformly distributed on $\mathbb{S}^{d-1}$, and outputs generated according to a target non-smooth function $f^* = S_G g^*$, with $g^*(x) = \mathbb{1}\{w_*^\top x \geq 0.7\}$, where the averaging operator S_G is over different groups in each plot. The regularization parameter λ is optimized on 5 000 test samples. We use the dot-product kernel function $\kappa(u) = (u+1)\kappa_1(u)$, where κ_1 is the arc-cosine kernel of degree 1, which corresponds to an infinite-width shallow ReLU network [9].

When the target is permutation-invariant, we can see in Figure 1(left) that the kernel based on permutation invariance leads to the largest gain in sample complexity compared to those which use cyclic or block-cyclic groups. Since the permutation group has the largest cardinality, this is consistent with our finding that the gains may be of the order of the size of the group. Figures 1(center/right) consider the example of cyclic translations on local blocks of size 2 or 6 (Example 10 with $s = 2$ or 6), and illustrate that the improvement in sample complexity happens later for $s = 2$ than $s = 6$, which is consistent with the slower decays of γ_d obtained in (21) due to the larger number of cycles.

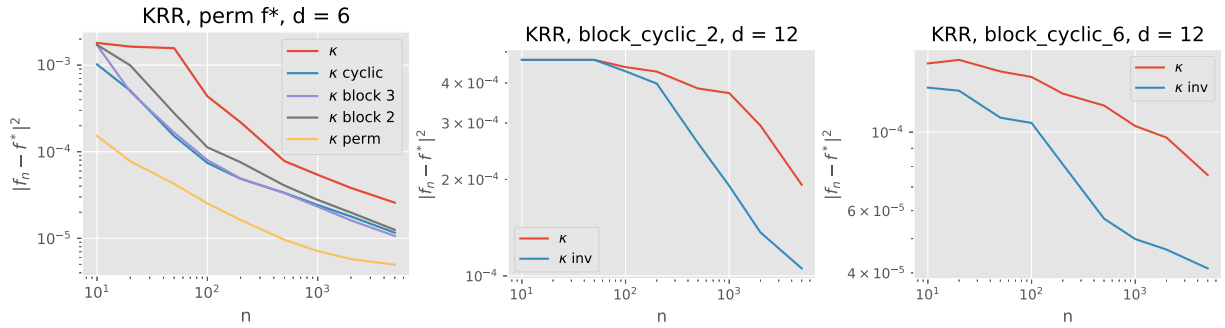

Figure 1: Comparison of KRR with invariant and non-invariant kernels. (left) permutation-invariant target with $d = 6$, comparison between various invariant kernels (cyclic, block-cyclic, and permutation groups). (center/right) invariant vs non-invariant kernels on invariant target functions with $d = 12$, for block-cyclic groups G of two different sizes.

7 Discussion and Conclusion

We have studied how geometric invariance or stability assumptions on target functions enable more efficient learning, with improvements in sample complexity which may be as large as the number of permutations considered in the group or set of elements to which the target is invariant or stable. In particular, this gain can be exponential in the dimension if we consider, *e.g.*, all permutations, local translations on small blocks, or permutations that resemble small deformations. This last example provides a strong justification for seeking models and architectures that are stable to deformations, a natural prior when learning functions on images [25]. In that respect, our results provide a theoretical baseline to assess learning guarantees under geometric priors: by designing appropriate geometrically stable kernels, we simultaneously address approximation and generalisation errors within a framework of convex optimization.

That said, while these gains may be large in practice, the obtained rates are still generically cursed by dimension if the target is non-smooth. In other words, invariance or geoemtric stability allows us to express Lipschitz assumptions with respect to weaker metrics. While these stronger regularity assumptions result in important gains in sample complexity, they do not overcome the inherent difficulty of learning non-smooth structures in high-dimensions. This suggests that further assumptions may be needed to learn efficiently on high-dimensional geometric data, for instance with more structured forms of regularity beyond our invariant/stable setup, which may be exploited perhaps through architectures that involve local connectivity, hierarchy (which would explain the benefits of depth, as opposed to our current results), or feature learning [2, 4, 16, 21, 27]. A natural further question is to study stability to transformations that are not necessarily permutations, and which may then provide more realistic models of continuous deformations. Another interesting question is to study whether it is possible to adapt to general symmetries present in the target, instead of encoding them in the model with an appropriately designed kernel as done here.

Acknowledgements

LV and JB acknowledge partial support from the Alfred P. Sloan Foundation, NSF RI-1816753, NSF CAREER CIF 1845360, NSF CHS-1901091, and Samsung Electronics.

References

[1] K. Atkinson and W. Han. *Spherical harmonics and approximations on the unit sphere: an introduction*, volume 2044. Springer Science & Business Media, 2012.

[2] F. Bach. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research (JMLR)*, 18(1):629–681, 2017.

[3] F. Bach. *Learning Theory from First Principles (draft)*. 2021.

[4] A. Bietti. Approximation and learning with deep convolutional models: a kernel perspective. *arXiv preprint arXiv:2102.10032*, 2021.

[5] A. Bietti and F. Bach. Deep equals shallow for ReLU networks in kernel regimes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

[6] A. Bietti and J. Mairal. Group invariance, stability to deformations, and complexity of deep convolutional representations. *Journal of Machine Learning Research (JMLR)*, 20(25):1–49, 2019.

[7] J. Bruna and S. Mallat. Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 35(8):1872–1886, 2013.

[8] A. Caponnetto and E. De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.

[9] Y. Cho and L. K. Saul. Kernel methods for deep learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2009.

[10] C. Ciliberto, F. Bach, and A. Rudi. Localized structured prediction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[11] N. Cohen and A. Shashua. Inductive bias of deep convolutional networks through pooling geometry. *arXiv preprint arXiv:1605.06743*, 2016.

[12] F. Cucker and S. Smale. On the mathematical foundations of learning. *Bulletin of the American mathematical society*, 39(1):1–49, 2002.

[13] S. S. Du, Y. Wang, X. Zhai, S. Balakrishnan, R. Salakhutdinov, and A. Singh. How many samples are needed to estimate a convolutional neural network? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

[14] C. Efthimiou and C. Frye. *Spherical harmonics in p dimensions*. World Scientific, 2014.

[15] B. Elesedy and S. Zaidi. Provably strict generalisation benefit for equivariant models. *arXiv preprint arXiv:2102.10333*, 2021.

[16] A. Favero, F. Cagnetta, and M. Wyart. Locality defeats the curse of dimensionality in convolutional teacher-student scenarios. *In preparation*, 2021.

[17] S. Fischer and I. Steinwart. Sobolev norm learning rates for regularized least-squares algorithms. *Journal of Machine Learning Research (JMLR)*, 21(205):1–38, 2020.

[18] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.

[19] Z. Li, Y. Zhang, and S. Arora. Why are convolutional nets more sample-efficient than fully-connected nets? In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

[20] J. Mairal, P. Koniusz, Z. Harchaoui, and C. Schmid. Convolutional kernel networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.

[21] E. Malach and S. Shalev-Shwartz. Computational separation between convolutional and fully-connected networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

[22] S. Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.

[23] S. Mei, T. Misiakiewicz, and A. Montanari. Learning with invariances in random features and kernel models. In *Conference on Learning Theory (COLT)*, 2021.

[24] H. N. Mhaskar and T. Poggio. Deep vs. shallow networks: An approximation theory perspective. *Analysis and Applications*, 14(06):829–848, 2016.

[25] L. Petrini, A. Favero, M. Geiger, and M. Wyart. Relative stability toward diffeomorphisms in deep nets indicates performance. *arXiv preprint arXiv:2105.02468*, 2021.

[26] L. Pillaud-Vivien, A. Rudi, and F. Bach. Statistical optimality of stochastic gradient descent on hard learning problems through multiple passes. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2018.

[27] T. Poggio, H. Mhaskar, L. Rosasco, B. Miranda, and Q. Liao. Why and when can deep-but not shallow-networks avoid the curse of dimensionality: a review. *International Journal of Automation and Computing*, 14(5):503–519, 2017.

[28] N. C. Saldanha and C. Tomei. The accumulated distribution of quadratic forms on the sphere. *Linear algebra and its applications*, 245:335–351, 1996.

[29] A. J. Smola, Z. L. Ovari, and R. C. Williamson. Regularization with dot-product kernels. In *Advances in Neural Information Processing Systems* (*NIPS*), 2001.

[30] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

This appendix contains additional background on spherical harmonics (Appendix A), and proofs of the results from Section 3, 4 and 5 (in Appendix B, C and D, respectively).

A Background on Spherical Harmonics and Legendre Polynomials

In this section, we provide background on spherical harmonic and Legendre/Gegenbauer polynomials, which are used extensively in our analysis. See, *e.g.*, [1, 14] for references. We consider inputs on the $d-1$ sphere $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d, \|x\| = 1\}$.

We recall some properties of the spherical harmonics $Y_{k,j}$ introduced in Section 2. For $j = 1, \dots, N(d,k)$, where $N(d,k) = \frac{2k+d-2}{k}\binom{k+d-3}{d-2}$, the spherical harmonics $Y_{k,j}$ are homogeneous harmonic polynomials of degree k that are orthonormal with respect to the uniform distribution τ on the d–1 sphere. The degree k plays the role of an integer frequency, as in Fourier series, and the collection $\{Y_{k,j}, k \geq 0, j = 1, \dots, N(d,k)\}$ forms an orthonormal basis of $L^2(\mathbb{S}^{d-1}, d\tau)$. As with Fourier series, there are tight connections between decay of coefficients in this basis w.r.t. k, and regularity/differentiability of functions, in this case differentiability on the sphere. In particular, this is a key property that we exploit for obtaining decays related to the number of invariant polynomials in Proposition 6. This follows from the fact that spherical harmonics are eigenfunctions of the Laplace-Beltrami operator on the sphere $\Delta_{\mathbb{S}^{d-1}}$ (see [14, Proposition 4.5]):

$$\Delta_{\mathbb{S}^{d-1}} Y_{k,j} = -k(k+d-2) Y_{k,j}. \tag{30}$$

For a given frequency k, we have the following addition formula:

$$\sum_{j=1}^{N(d,k)} Y_{k,j}(x) Y_{k,j}(y) = N(d,k) P_{d,k}(x^\top y), \tag{31}$$

where $P_{d,k}$ is the k-th Legendre polynomial in dimension d (also known as Gegenbauer polynomial when using a different scaling), given by the Rodrigues formula:

$$P_{d,k}(t) = (-1/2)^k \frac{\Gamma(\frac{d-1}{2})}{\Gamma(k+\frac{d-1}{2})} (1-t^2)^{(3-d)/2} \left(\frac{d}{dt}\right)^k (1-t^2)^{k+(d-3)/2}. \tag{32}$$

The polynomials $P_{d,k}$ are orthogonal in $L^2([-1,1], dw)$ where the measure dw is given by the weight function $dw(t) = (1-t^2)^{(d-3)/2} dt$, and we have

$$\int_{-1}^{1} P_{d,k}^2(t)(1-t^2)^{(d-3)/2} dt = \frac{\omega_{d-1}}{\omega_{d-2}} \frac{1}{N(d,k)}, \tag{33}$$

where $\omega_{p-1} = \frac{2\pi^{p/2}}{\Gamma(p/2)}$ denotes the surface of the sphere $\mathbb{S}^{p-1}$ in p dimensions. Using the addition formula (31) and orthogonality of spherical harmonics, we can show

$$\int P_{d,j}(w^\top x) P_{d,k}(w^\top y) d\tau(w) = \frac{\delta_{jk}}{N(d,k)} P_{d,k}(x^\top y) \tag{34}$$

We will use the following recurrence relation of Legendre polynomials [14, Eq. 4.36]

$$t P_{d,k}(t) = \frac{k}{2k+d-2} P_{d,k-1}(t) + \frac{k+d-2}{2k+d-2} P_{d,k+1}(t), \tag{35}$$

for $k \geq 1$, and for $k = 0$ we simply have $t P_{d,0}(t) = P_{d,1}(t)$. We will also use the following pointwise upper bound on $P_{d,k}(t)$ (see [1, Eq. 2.117]):

$$|P_{d,k}(t)| \leq \frac{\Gamma\left(\frac{d-1}{2}\right)}{\sqrt{\pi}} \left(\frac{4}{k(1-t^2)}\right)^{(d-2)/2} \tag{36}$$

The Funk-Hecke formula is helpful for computing Fourier coefficients in the basis of spherical harmonics in terms of Legendre polynomials: for any $j = 1, \dots, N(d,k)$, we have

$$\int f(x^\top y) Y_{k,j}(y) d\tau(y) = \frac{\omega_{d-2}}{\omega_{d-1}} Y_{k,j}(x) \int_{-1}^{1} f(t) P_{d,k}(t) (1-t^2)^{(d-3)/2} dt. \tag{37}$$

For example, we may use this to obtain decompositions of dot-product kernels by computing Fourier coefficients of functions $\kappa(\langle x, \cdot \rangle)$. Indeed, denoting

$$\mu_k = \frac{\omega_{d-2}}{\omega_{d-1}} \int_{-1}^{1} \kappa(t) P_{d,k}(t) (1-t^2)^{(d-3)/2} dt,$$

writing the decomposition of $\kappa(\langle x, \cdot \rangle)$ using (37) leads to the following Mercer decomposition of the kernel:

$$\kappa(x^\top y) = \sum_{k=0}^{\infty} \mu_k \sum_{j=1}^{N(d,k)} Y_{k,j}(x) Y_{k,j}(y) = \sum_{k=0}^{\infty} \mu_k N(d,k) P_{d,k}(x^\top y). \tag{38}$$

B Proofs for Section 3 (Sample complexity of Invariant Kernels)

B.1 Proof of Lemma 2 (spectral properties of K and K_G)

Proof. That μ_k are eigenvalues of T_K with multiplicity $N(d,k)$ is standard and follows from the Funk-Hecke formula (see, *e.g.*, [29]). In particular, for any spherical harmonic $Y_k \in V_{d,k}$, we have $T_K Y_k = \mu_k Y_k$, and there are $N(d,k)$ orthogonal spherical harmonics in $V_{d,k}$.

For K_G, note that we have $T_{K_G} = S_G T_K$, so that for any G-invariant spherical harmonic $\overline{Y}_k \in \overline{V}_{d,k}$ we have $T_{K_G} \overline{Y}_k = \mu_k S_G \overline{Y}_k = \mu_k \overline{Y}_k$, while for any $Y_k \in V_{d,k} \cap \overline{V}_{d,k}^{\perp}$, we have $T_{K_G} Y_k = 0$ since $S_G Y_k = 0$. □

B.2 Proof of Lemma 3 (approximation error)

Proof. Let $\overline{Y}_{k,j}$, for $j = 1, \dots, N(d,k)$ be an orthonormal basis of $V_{d,k}$ such that $(\overline{Y}_{k,j})_{j \le \overline{N}(d,k)}$ form an orthonormal basis of $\overline{V}_{d,k}$. Then, the collection of $\overline{Y}_{k,j}$ for $k \ge 0$ and $j = 1, \dots, N(d,k)$ forms an orthonormal basis of $L^2(d\tau)$.

For a function $f \in L^2(d\tau)$ with decomposition

$$f(x) = \sum_{k \ge 0} \sum_{j=1}^{N(d,k)} a_{k,j} \overline{Y}_{k,j}(x),$$

we have the following expressions of the RKHS norms for K and K_G by Mercer's theorem (*e.g.*, [12]):

$$\|f\|_{\mathcal{H}_K}^2 = \begin{cases} \sum_{k:\mu_k > 0} \sum_{j=1}^{N(d,k)} \frac{a_{k,j}^2}{\mu_k}, & \text{if } a_{k,j} = 0 \text{ whenever } \mu_k = 0 \\ \infty, & \text{otherwise.} \end{cases}$$

$$\|f\|_{\mathcal{H}_{K_G}}^2 = \begin{cases} \sum_{k:\mu_k > 0} \sum_{j=1}^{\overline{N}(d,k)} \frac{a_{k,j}^2}{\mu_k}, & \text{if } a_{k,j} = 0 \text{ whenever } \mu_k = 0 \text{ or } j > \overline{N}(d,k) \\ \infty, & \text{otherwise.} \end{cases}$$

Assume now that f^* is invariant, so that its coefficients $a^*_{k,j}$ satisfy $a^*_{k,j} = 0$ for $j > \overline{N}(d,k)$. We have

$$\begin{aligned}
A_{\mathcal{H}_K}(\lambda, f^*) &= \inf_{f \in \mathcal{H}_K} \|f - f^*\|^2_{L^2(d\tau)} + \lambda \|f\|^2_{\mathcal{H}_K} \\
&= \inf_{a_{k,j}} \sum_{k:\mu_k>0} \sum_{j=0}^{N(d,k)} \left((a_{k,j} - a^*_{k,j})^2 + \lambda \frac{a^2_{k,j}}{\mu_k} \right) \\
&= \inf_{a_{k,j}} \sum_{k:\mu_k>0} \sum_{j=0}^{\overline{N}(d,k)} \left((a_{k,j} - a^*_{k,j})^2 + \lambda \frac{a^2_{k,j}}{\mu_k} \right) + \sum_{k:\mu_k>0} \sum_{j=\overline{N}(d,k)+1}^{N(d,k)} (1 + \lambda/\mu_k) a^2_{k,j} \\
&= \inf_{a_{k,j}} \sum_{k:\mu_k>0} \sum_{j=0}^{\overline{N}(d,k)} \left((a_{k,j} - a^*_{k,j})^2 + \lambda \frac{a^2_{k,j}}{\mu_k} \right) \\
&= A_{\mathcal{H}_{K_G}}(\lambda, f^*),
\end{aligned}$$

which proves the lemma. □

B.3 Proof of Lemma 4 (degrees of freedom)

Proof. The result immediately follows from the following expressions of degrees of freedom for K and K_G:

$$\mathcal{N}_K(\lambda) = \sum_{k \geq 0} N(d,k) \frac{\mu_k}{\mu_k + \lambda}, \qquad \mathcal{N}_{K_G}(\lambda) = \sum_{k \geq 0} \overline{N}(d,k) \frac{\mu_k}{\mu_k + \lambda}. \tag{39}$$

□

B.4 Proof of Theorem 5 (generalization bound)

Proof. We start from the following bound from [3, Proposition 7.2], which holds for any $\lambda \leq 1$, assuming $K_G(x,x) \leq 1$ almost surely (this is satisfied when $\kappa(u) \leq 1$), and for $n \geq \frac{5}{\lambda}(1 + \log(1/\lambda))$:

$$\mathbb{E}[R(\hat{f}_\lambda)] - R(f^*) \leq 16 \frac{\sigma_q^2}{n} \mathcal{N}_{K_G}(\lambda) + 16 A_{\mathcal{H}_{K_G}}(\lambda, f^*) + \frac{24}{n^2} \|f^*\|^2_\infty. \tag{40}$$

Under assumption (A2), we have (see, *e.g.*, [12, Theorem 7], using that $\|f\|_{\mathcal{H}_K} = \|T_K^{-1/2} f\|_{L^2(d\tau)}$)

$$A_{\mathcal{H}_K}(\lambda, f^*) \leq C_{f^*} \lambda^{2r}, \tag{41}$$

with $C_{f^*} := \|T_K^{-r} f^*\|^2_{L^2(d\tau)}$ By Lemma 3, we also have

$$A_{\mathcal{H}_{K_G}}(\lambda, f^*) \leq C_{f^*} \lambda^{2r}. \tag{42}$$

Using Lemma 4 for some integer $\ell \geq 0$ and (A1), the bound (40) becomes

$$\mathbb{E}[R(\hat{f}_\lambda)] - R(f^*) \leq 16 C_{f^*} \lambda^{2r} + 16 \frac{\sigma_q^2 D(\ell)}{n} + 16 \frac{C_K \sigma_q^2 \nu_d(\ell)}{n} \lambda^{-1/\alpha} + \frac{24}{n^2} \|f^*\|^2_\infty. \tag{43}$$

Jointly optimizing the first and third terms for λ yields

$$\lambda_n = \left(\frac{C_K \sigma_q^2 \nu_d(\ell)}{2r\alpha C_{f^*} n} \right)^{\frac{\alpha}{2\alpha r + 1}} \tag{44}$$

The bound then becomes

$$\mathbb{E}[R(\hat{f}_{\lambda_n})] - R(f^*) \lesssim C_{f^*}^{-\frac{\alpha}{2\alpha r+1}} \left(\frac{C_K \sigma_q^2 \nu_d(\ell)}{n} \right)^{\frac{2\alpha r}{2\alpha r+1}} + \frac{\sigma_q^2 D(\ell)}{n} + \frac{1}{n^2} \|f^*\|^2_\infty. \tag{45}$$

Here, $\lesssim$ hides only absolute constants that depend on α and r. Now we choose $\ell = \ell_n$, given by (15), with constant corresponding to:

$$\ell_n := \sup\{\ell : \sigma_q^2 D(\ell) \leq C_{f^*}^{-\frac{\alpha}{2\alpha r+1}} \left(C_K \sigma_q^2 \nu_d(\ell)\right)^{\frac{2\alpha r}{2\alpha r+1}} n^{\frac{1}{2\alpha r+1}}\}, \tag{46}$$

so that the second term is smaller than the first term. The last term is of the same order when

$$n \gtrsim \frac{\|f^*\|_\infty^2}{\sigma_q^2 D(\ell_n)}, \tag{47}$$

which is verified under the condition (14) since $D(\ell_n) \geq 1$. Note that for the specific bound (40) to hold we also need $\lambda_n \gtrsim 1/n$ (up to logarithmic terms). This imposes the qualification condition $r > (\alpha - 1)/2\alpha$, and leads to the additional requirement

$$n \gtrsim \left(\frac{C_{f^*}}{\sigma_q^2 C_K \nu_d(\ell)}\right)^{\frac{\alpha}{2\alpha r+1-\alpha}}. \tag{48}$$

This is verified under condition (14) with $C_1 = C_{f^*}/\sigma_q^2 C_K$.

For the KRR estimator with kernel K, the same bound (43) holds, but without the factor $\nu_d(\ell)$ and without the term involving $D(\ell)$. The resulting bound follows from a similar analysis. □

C Proofs for Section 4 (Decays of $\gamma_d(k)$)

C.1 Proof of Proposition 6 (decay of $\gamma_{d,\sigma}(k)$)

The proof of Proposition 6 is technical and relies on identifying and analyzing the singularities in the density q_σ of the random variable $Z_\sigma = \langle \sigma \cdot x, x \rangle$, with $x \sim \tau$, using results in [28]. Lemma 15 provides a general integration by parts result which is useful throughout the proof to obtain asymptotic decays from regularity properties. Lemma 17 and Lemma 18 provide asymptotic decays for singularities $\phi(t)$ localized around some λ in $(-1, 1)$ and $\{\pm 1\}$, respectively, either through integration by parts or using closed form expressions of certain integrals. Proposition 6 is then proved by appropriately "cancelling" the singularities in q_σ using such localized functions ϕ, as explained in Lemma 19, and applying the integration by parts lemma on the resulting function, which is of higher smoothness and thus leads to faster-decaying terms.

Lemma 15 (Integration by parts)**.** *Let $g : [-1, 1] \to \mathbb{R}$ be $2s$-times differentiable, with all derivatives bounded on $[-1, 1]$. We then have*

$$\int_{-1}^{1} g(t) P_{d,k}(t)(1-t^2)^{\frac{d-3}{2}} dt = \frac{1}{(k(k+d-2))^s} \int_{-1}^{1} \tilde{g}_{d,s}(t) P_{d,k}(t)(1-t^2)^{\frac{d-3}{2}} dt, \tag{49}$$

where $\tilde{g}_{d,s}$ is a bounded function on $[-1, 1]$.

Proof. We use the following relation, derived in [5, Lemma 4] for a function f_0:

$$\int_{-1}^{1} f_0(t) P_{d,k}(t)(1-t^2)^{\frac{d-3}{2}} dt = \frac{1}{k(k+d-2)} \Big(- f_0(t)(1-t^2)^{1+\frac{d-3}{2}} P'_{d,k}(t) \Big|_{-1}^{1} + f_0'(t)(1-t^2)^{1+\frac{d-3}{2}} P_{d,k}(t) \Big|_{-1}^{1} + \int_{-1}^{1} f_1(t) P_{d,k}(t)(1-t^2)^{(d-3)/2} dt \Big),$$

where $f_1(t) = -f_0''(t)(1-t^2) + (d-1)t f_0'(t)$. Note that the terms in brackets vanish when f_0 and f_0' are bounded, and that f_1 is $2s-2$ times differentiable with bounded derivatives if f_0 is $2s$ times differentiable. We may thus apply this recursively s times to $f_0 = g$, with

$$f_k(t) = -f_{k-1}''(t)(1-t^2) + (d-1)t f_{k-1}'(t),$$

and we obtain the desired result, with $\tilde{g} = f_s$. □

Lemma 16. *Let $g \in L^\infty([-1,1])$. It holds that*

$$\left| \int_{-1}^{1} g(t) P_{d,k}(t) (1-t^2)^{(d-3)/2} \, dt \right| \leq 2\pi d^{-1/2} \|g\|_\infty \left(\frac{d}{k} \right)^{(d-2)/2} .$$

Proof. By [1, equation (2.117)], we get that

$$\begin{aligned} |P_{d,k}(t)| &\leq \frac{1}{\sqrt{\pi}} \Gamma\left(\frac{d-1}{2} \right) \left(\frac{4}{k(1-t^2)} \right)^{(d-2)/2} \\ &\leq \frac{1}{\sqrt{\pi}} \left(\frac{d-1}{4} \right)^{(d-3)/2} \left(\frac{4}{k(1-t^2)} \right)^{(d-2)/2} \\ &\leq \frac{1}{\sqrt{\pi}} \left(\frac{d}{4} \right)^{-1/2} \left(\frac{d}{k(1-t^2)} \right)^{(d-2)/2} \leq 2d^{-1/2} \left(\frac{d}{k(1-t^2)} \right)^{(d-2)/2} . \end{aligned}$$

Therefore it follows that

$$\left| \int_{-1}^{1} g(t) P_{d,k}(t) (1-t^2)^{(d-3)/2} \, dt \right| \leq 2d^{-1/2} \|g\|_\infty \left(\frac{d}{k} \right)^{(d-2)/2} \int_{-1}^{1} (1-t^2)^{-1/2} \, dt \, ,$$

which concludes the proof. □

Lemma 17 (Decay for $\lambda \in (-1,1)$)**.** *Let $\phi_{\lambda+,\alpha}(t) := (t-\lambda)_+^\alpha \varphi_{\lambda+,\alpha}(t)$ and $\phi_{\lambda-,\alpha}(t) := (t-\lambda)_-^\alpha \varphi_{\lambda-,\alpha}(t)$, where $\varphi_{\lambda\pm,\alpha} \in C^\infty([-1,1])$ have support $(-1+\epsilon, 1-\epsilon)$ for some $\epsilon > 0$ and take the value 1 at $t = \lambda$. Then we have that*

$$\left| \int_{-1}^{1} \phi_{\lambda\pm,\alpha}(t) P_{d,k}(t) dt \right| \leq C(d,\alpha) k^{-d/2-\alpha+3} . \tag{50}$$

Also let, for α integer, $\phi^*_{\lambda\pm,\alpha}(t) := (t-\lambda)_\pm^\alpha \log|t-\lambda| \varphi^*_{\lambda\pm,\alpha}(t)$, where $\varphi^*_{\lambda\pm,\alpha} \in C^\infty([-1,1])$ have support $(-1+\epsilon, 1-\epsilon)$ for some $\epsilon > 0$ and take the value 1 at $t = \lambda$. Then we have that

$$\left| \int_{-1}^{1} \phi^*_{\lambda\pm,\alpha}(t) P_{d,k}(t) dt \right| \leq C(d,\alpha) k^{-d/2-\alpha+3} . \tag{51}$$

Proof. Let $\psi_{\lambda\pm,\alpha}(t) := \phi_{\lambda\pm,\alpha}(t)(1-t^2)^{-(d-3)/2}$. Notice that $\psi_{\lambda+,\alpha}$ satisfies the assumption of Lemma 15 with $2s = 2\lfloor \frac{\alpha}{2} \rfloor \geq \alpha - 2$. Therefore we obtain

$$\begin{aligned} \left| \int_{-1}^{1} \phi_{\lambda+,\alpha}(t) P_{d,k}(t) dt \right| &= \left| \int_{-1}^{1} \psi_{\lambda+,\alpha}(t) P_{d,k}(t) (1-t^2)^{(d-3)/2} dt \right| \\ &\leq k^{-\alpha+2} \left| \int_{-1}^{1} \tilde{\psi}_{\lambda+,\alpha,d}(t) P_{d,k}(t) \left(1-t^2\right)^{(d-3)/2} dt \right| \\ &\leq C(d,\alpha) k^{-\alpha+2-d/2+1} , \end{aligned}$$

where $\tilde{\psi}_{\lambda+,\alpha,d}$ is a bounded function given by Lemma 15, and where we used Lemma 16 to obtain the last inequality. The second inequality follows in the same way, by noticing that the function $\psi^*_{\lambda\pm,\alpha}(t) := \phi^*_{\lambda\pm,\alpha}(t)(1-t^2)^{-(d-3)/2}$ satisfies the assumption of Lemma 15 with $2s = 2\lfloor \frac{\alpha-1}{2} \rfloor \geq \alpha - 2$ (since we assume that α is integer in this case). □

Lemma 18 (Decay for $\lambda = \pm 1$)**.** *Let $\phi_{1,\alpha,s}(t) := (\frac{1+t}{2})^{\alpha+s-\lfloor\alpha\rfloor}(1-t)^\alpha$, with α non-integer and s integer. Then, $\phi_{1,\alpha,s}$ is s times differentiable at -1 and obeys the decay*

$$\left| \int_{-1}^{1} \phi_{1,\alpha,s}(t) P_{d,k}(t) dt \right| \leq C(d,\alpha,s) k^{-2(\alpha+1)}, \tag{52}$$

where the constant $C(d,\alpha,s)$ may be different depending on the parity of k.

Similarly, let $\phi_{-1,\alpha,s}(t) := (\frac{1-t}{2})^{\alpha+s-\lfloor\alpha\rfloor}(t+1)^\alpha$, then $\phi_{-1,\alpha,s}$ is s times differentiable at 1, and obeys the same decay.

Proof. We begin by evaluating the decay of $\psi_\alpha(t) := (1-t^2)^\alpha$. Following analogous calculations to [5, Lemma 6], we have[1]

$$\int \psi_\alpha(t) P_{d,k}(t) dt \sim C(d,\alpha) k^{-2(\alpha+1)}, \tag{53}$$

for k even, and the integral is equal to zero for k odd. We have

$$C(d,\alpha) = 2^{4\alpha+3} \frac{\omega_{d-2}}{\omega_{d-1}} \frac{\Gamma(\alpha+1)^2}{\Gamma(2\alpha+2)} \frac{\Gamma(\alpha+\frac{3}{2})\Gamma(\frac{d-1}{2})\Gamma(\alpha+\frac{5-d}{2})}{\Gamma(-\frac{1}{2})\Gamma(\alpha+\frac{7-d}{2})\Gamma(-\alpha+\frac{d-5}{2})},$$

where ω_{p-1} is the surface of $\mathbb{S}^{p-1}$.

Now, let $r := s - \lfloor \alpha \rfloor$, so that we have

$$\phi_{1,\alpha,s}(t) = 2^{-\alpha-r}(1+t)^r \psi_\alpha(t).$$

Let $c_0, \dots, c_r$ denote the coefficients of the degree-r polynomial $p(t) = 2^{-\alpha-r}(1+t)^r$, so that

$$p(t) = 2^{-\alpha-r}(1+t)^r = c_0 + c_1 t + \cdots + c_r t^r.$$

Using the relation (see, *e.g.*, [14, Proposition 4.21])

$$t P_{d,k}(t) = \frac{k}{2k+d-2} P_{d,k-1}(t) + \frac{k+d-2}{2k+d-2} P_{d,k+1}(t),$$

we may then write

$$p(t) P_{d,k}(t) = \sum_{j=-r}^{r} b_j(k) P_{d,k+j}(t),$$

for some coefficients $b_j(k)$ satisfying $b_j(k) = O(1)$ as $k \to \infty$. Then, we have

$$\begin{aligned}\int \phi_{1,\alpha,s} P_{d,k}(t) dt &= \int \psi_\alpha(t) p(t) P_{d,k}(t) dt \\ &= \sum_{j=-r}^{r} b_j(k) \int \psi_\alpha(t) P_{d,k+j}(t) dt.\end{aligned}$$

When $k \to \infty$, this is a sum of at most $2r+1 \le 2s+1$ terms, each of which decays with k as $k^{-2(\alpha+1)}$ by (53). This yields the result.

The decay for $\phi_{-1,\alpha,s}$ is proved analogously. □

Lemma 19 (Cancelling singularities of the density)**.** *Let q_σ denote the density of the random variable $Z_\sigma = \langle \sigma \cdot x, x \rangle$, with $\sigma \neq \mathrm{Id}$, and let $\bar{\Lambda}_\sigma$ be the set of eigenvalues of $\bar{A}_\sigma := (A_\sigma + A_\sigma^\top)/2$, where A_σ is the permutation matrix of σ, and denote by $\bar{m}_\lambda$ the multiplicity of $\lambda \in \bar{\Lambda}_\sigma$. Define*

$$\alpha_\lambda = \frac{d - \bar{m}_\lambda}{2} - 1. \tag{54}$$

There exists constants $\{c_{\lambda,i}\}$ such that the function defined by

$$\begin{aligned}\tilde{q}_\sigma = &\sum_{\lambda \in \Lambda_\sigma \setminus \{1, \lambda_{\min}\}} \sum_{i=0}^{\lceil d+1-\alpha_\lambda \rceil} (c_{\lambda_+,i} \phi_{\lambda_+,\alpha_\lambda+i} + c_{\lambda_-,i} \phi_{\lambda_-,\alpha_\lambda+i} + c^*_{\lambda_+,i} \phi^*_{\lambda_+,\alpha_\lambda+i} + c^*_{\lambda_-,i} \phi^*_{\lambda_-,\alpha_\lambda+i}) \\ &+ \mathbf{1}\{\lambda_{\min} > -1\} \sum_{i=0}^{\lceil d+1-\alpha_{\lambda_{\min}} \rceil} (c_{\lambda_{\min}+,i} \phi_{\lambda_{\min}+,\alpha_{\lambda_{\min}}+i} + c^*_{\lambda_{\min}+,i} \phi^*_{\lambda_{\min}+,\alpha_{\lambda_{\min}}+i}) \\ &+ \sum_{\lambda \in \Lambda_\sigma \cap \{\pm 1\}} \sum_{i=0}^{\lceil d+1+\frac{d-3}{2}-\alpha_\lambda \rceil} c_{\lambda,i} \phi_{\lambda,\alpha_\lambda+i,\lceil d+1+\frac{d-3}{2} \rceil}\end{aligned}$$

satisfies that $t \mapsto (q_\sigma(t) - \tilde{q}_\sigma(t))(1-t^2)^{-\frac{d-3}{2}}$ admits $d+1$ bounded derivatives on $[-1,1]$.

[1] Note that while Lemma 6 in [5] is stated for $\alpha = \nu + \frac{d-3}{2}$ for $\nu > 0$, the derivation still holds for any $\alpha > -1$.

Proof. Note that we have

$$\langle \sigma \cdot x, x\rangle = \frac{1}{2}(\langle A_\sigma x, x\rangle + \langle x, A_\sigma, x\rangle) = \langle \bar{A}_\sigma x, x\rangle,$$

where $\bar{A}_\sigma = \frac{1}{2}(A_\sigma + A_\sigma^\top)$ is symmetric and thus has real eigenvalues. When A_σ is a permutation matrix, these eigenvalues are in $[-1, 1]$, as the real part of complex roots of unity.

We then identify the singularities of the density q_σ, which are the same as those of the cumulative distribution function, up to one fewer degree of smoothness. Such singularities are shown in the following lemma, proved in [28].

Lemma 20 ([28])**.** *Consider $\sigma \in G$ as above, and let $\bar{\Lambda}_\sigma$ be the set of eigenvalues of $\bar{A}_\sigma$. For each $\lambda \in \bar{\Lambda}_\sigma$, we denote by $\bar{m}_\lambda$ its multiplicity. Then the cumulative distribution function Q_σ of $Z_\sigma = \langle \sigma \cdot x, x\rangle$ takes the form*

$$Q_\sigma(t) = \varphi(t) + \sum_{\lambda \in \bar{\Lambda}_\sigma} g_\lambda(t)$$

where φ is analytic and

- $g_\lambda(t) = |t-\lambda|(t-\lambda)^{\frac{d-\bar{m}_\lambda}{2}-1}\varphi_\lambda^1(t) + (t-\lambda)^{\frac{d-\bar{m}_\lambda}{2}}\log(|t-\lambda|)\,\varphi_\lambda^2(t)$ *if $d-\bar{m}_\lambda$ is even,*
- $g_\lambda(t) = (t-\lambda)_+^{\frac{d-\bar{m}_\lambda}{2}}\varphi_\lambda^1(t) + (\lambda-t)_+^{\frac{d-\bar{m}_\lambda}{2}}\,\varphi_\lambda^2(t)$ *if $d-\bar{m}_\lambda$ is odd,*

for some $\varphi_\lambda^1, \varphi_\lambda^2$ analytic. Further, the term involving $\log(|t-\lambda|)$ only appears for $\lambda \in (-1, 1)$.

In particular, it follows from this lemma by differentiation that we may write

$$q_\sigma(t) = \tilde{\varphi}(t) + \sum_{\lambda \in \bar{\Lambda}_\sigma} \tilde{g}_\lambda(t),$$

with $\tilde{\varphi}$ analytic and

- $\tilde{g}_\lambda(t) = (t-\lambda)_+^{\frac{d-\bar{m}_\lambda}{2}-1}\tilde{\varphi}_{\lambda,1}(t) + (t-\lambda)_-^{\frac{d-\bar{m}_\lambda}{2}-1}\tilde{\varphi}_{\lambda,2}(t) + (t-\lambda)_+^{\frac{d-\bar{m}_\lambda}{2}-1}\log(|t-\lambda|)\tilde{\varphi}_{\lambda,3}(t) + (t-\lambda)_-^{\frac{d-\bar{m}_\lambda}{2}-1}\log(|t-\lambda|)\tilde{\varphi}_{\lambda,4}(t)$, if $d-\bar{m}_\lambda$ is even
- $\tilde{g}_\lambda(t) = (t-\lambda)_+^{\frac{d-\bar{m}_\lambda}{2}-1}\tilde{\varphi}_{\lambda,1}(t) + (t-\lambda)_-^{\frac{d-\bar{m}_\lambda}{2}-1}\tilde{\varphi}_{\lambda,2}(t)$, if $d-\bar{m}_\lambda$ is odd,

where $\tilde{\varphi}_{\lambda,i}$ are analytic for $i \in [4]$.

The result then follows by appropriately "cancelling" those singularities up to order d using the simple functions ϕ introduced in the previous lemmas, and noting that for singularities at ± 1, we require an additional $(d-3)/2$ degrees of smoothness, so that we may divide by the weight function $(1-t^2)^{(d-3)/2}$.

For instance, for $\lambda \in (-1, 1)$, an appropriate exponent α and an analytic φ, we may write

$$(t-\lambda)_+^\alpha \varphi(t) = (t-\lambda)_+^{\frac{d-\bar{m}_\lambda}{2}-1}(c_0 + (t-\lambda)\psi(t)), \tag{55}$$

with ψ analytic. Then for $\phi_{\lambda+,\alpha}$ as in Lemma 17, we have

$$(t-\lambda)_+^\alpha \varphi(t) - c_0\phi_{\lambda+,\alpha}(t) = (t-\lambda)_+^{\alpha+1}\tilde{\varphi}(t),$$

with $\tilde{\varphi}$ analytic. We may then repeat this process with functions $\phi_{\lambda+,\alpha+1}$, $\phi_{\lambda+,\alpha+2}$, etc., to finally obtain that

$$(t-\lambda)_+^\alpha \varphi(t) - \sum_{i=0}^{\lceil d+1-\alpha\rceil} c_i\phi_{\lambda+,\alpha+i}(t)$$

is d times differentiable, as desired. A similar reasoning can be applied for other types of singularities. For the terms involving $\log(|t-\lambda|)$, which only appear for $\lambda \in (-1, 1)$, note that the corresponding exponent α_λ is integer since $d-\bar{m}_\lambda$ is even, so that Lemma 17 applies.

□

We are now ready to state the proof of Proposition 6.

Proof of Proposition 6. Let q_σ be the density of $\langle \sigma \cdot x, x\rangle$, and let $\tilde{q}_\sigma$ be as in Lemma 19. By Lemma 19 and Lemma 15, we have

$$\int_{-1}^{1} (q_\sigma(t) - \tilde{q}_\sigma(t)) P_{d,k}(t) dt = \int_{-1}^{1} \frac{q_\sigma(t) - \tilde{q}_\sigma(t)}{(1-t^2)^{\frac{d-3}{2}}} P_{d,k}(t)(1-t^2)^{\frac{d-3}{2}} dt \leq Ck^{-d},$$

where we have bounded the integral on the r.h.s. of (49) by a constant. Renaming the terms in $\tilde{q}$ as $\tilde{q}_\sigma = \sum_i c_i q_i$, where each q_i is as in Lemma 17 or Lemma 18, we have

$$\gamma_{d,\sigma}(k) \leq \sum_i c_i \int q_i P_{d,k} + O(k^{-d}).$$

Now, note that the eigenvalues of $\bar{A}_\sigma = \frac{1}{2}(A_\sigma + A_\sigma^\top)$ are the real parts of the complex eigenvalues of the permutation matrix A_σ. Since A_σ is real, its complex eigenvalues come in conjugate pairs, so that for any $\lambda \in \Lambda_\sigma$ with $\lambda \notin \mathbb{R}$, we have $m_\lambda = \bar{m}_{Re(\lambda)}/2$, where $\bar{m}$ are the multiplicities of Lemma 19. Note that in the case of permutations, the eigenvalues of A_σ are roots of unity, so that we have

$$m_\lambda = \begin{cases} \bar{m}_\lambda, & \text{if } \lambda \in \{\pm 1\} \\ \bar{m}_{Re(\lambda)}/2, & \text{otherwise.} \end{cases}$$

The result then follows by applying the decays given by Lemma 17 and Lemma 18 to each component q_i with the appropriate α_λ, and focusing on the leading-order terms. Namely, for $\lambda \in \{\pm 1\}$, by Lemma 18, the leading order term in $\gamma_{d,\sigma,\lambda}(k)$ has decay

$$k^{-2(\alpha_\lambda+1)} = k^{-2(\frac{d-\bar{m}_\lambda}{2}-1+1)} = k^{-d+m_\lambda},$$

while for $\lambda \notin \{\pm 1\}$, we have $Re(\lambda) \in (-1, 1)$, hence using (54) and $\bar{m}_{Re(\lambda)} = 2m_\lambda$, we have $\alpha_{Re(\lambda)} = \frac{d}{2} - m_\lambda - 1$, so that the decay, by Lemma 17, is upper bounded by

$$k^{-\frac{d}{2}-\alpha_{Re(\lambda)}+3} \leq k^{-d+m_\lambda+4}.$$

This concludes the proof. □

C.2 Proof of Corollary 7 (leading order of $\gamma_{d,\sigma}$)

Proof. By Proposition 6, we get that

$$\gamma_{d,k}(\sigma) \lesssim k^{-d+s}$$

where $s = \max_{\lambda \in \Lambda_\sigma} \{m_\lambda + 4 \cdot \mathbf{1}(|\lambda| < 1)\}$. Notice that, for any permutation σ, it holds $m_1 \geq m_\lambda$. Therefore, we have

$$s \leq m_1 + 4 \cdot \mathbf{1}\{\exists |\lambda| < 1 \ : \ m_1 < m_\lambda + 4\} .$$

Now, if $m_1 < m_\lambda + 4$ for some $|lambda| < 1$, since $m_1 + m_\lambda \leq d$, it must hold $2m_1 \leq d + 3$, or, equivalently, $m_1 \leq d/2 + 3/2$. It follows that

$$s \leq \begin{cases} m_1 & \text{if } m_1 > (d+3)/2 \, , \\ d/2 + 5.5 & \text{otherwise.} \end{cases}$$

This concludes the proof, since $m_1 = c(\sigma)$. □

C.3 Proof of Corollary 8 (different upper bounds using permutation statistics)

Proof. For all $\sigma \in G \setminus \zeta(G,s)$, it holds $c(\sigma) \leq s$ and

$$\gamma_{d,\sigma}(k) \lesssim k^{-d+\eta_\sigma} ,$$

with

$$\eta_\sigma = c(\sigma) \cdot \mathbf{1}\left(c(\sigma) > (d+3)/2\right) + \left(\frac{d}{2}+6\right)\mathbf{1}\left(c(\sigma) \leq (d+3)/2\right) .$$

In particular, we have

$$\gamma_d(k) = \frac{\zeta(G,s)}{|G|} + O\left(k^{-d+\eta}\right)$$

where

$$\eta = \max_{c(\sigma)\leq s} \eta_\sigma .$$

Denote $s^*(s) = \max_{\sigma\in G\setminus\zeta(G,s)} c(\sigma)$. If $s^*(s) \geq d/2+6$, then it holds that $s^*(s) = \eta$. Otherwise, $c(\sigma) \leq d/2+7$ for any σ such that $c(\sigma) \leq s$, which implies that $s \leq d/2+6$. It follows that

$$\eta = \max\{s^*(s),\ d/2+6\} \leq \max\{s,\ d/2+6\}.$$

□

C.4 Details on Example 11 (full permutation group)

The number of permutations in $G = S_d$ which fix exactly n elements is given by $\binom{n}{k}\,!(n-k)$, where $!m$ denotes the m-th subfactorial:

$$!m := \left\lfloor \frac{m!+1}{e} \right\rfloor \leq \frac{2m!}{e} .$$

It follows that

$$\begin{aligned}
\frac{\xi(G,s)}{|G|} &= \frac{1}{d!}\sum_{k=s+1}^{d} \frac{d!}{(d-k)!k!}\,!(d-k) \leq \frac{2}{e}\sum_{k=s+1}^{d}\frac{1}{k!} \\
&\leq \frac{2}{e(s+1)!}\sum_{k=s+1}^{d}\frac{1}{(s+2)^{k-(s+1)}} = \frac{2}{e(s+1)!}\frac{1-\frac{1}{(s+2)^{d-s}}}{1-\frac{1}{s+2}} \\
&\leq \frac{2(s+2)}{e(s+1)}\frac{1}{(s+1)!} \leq \frac{2}{(s+1)!} .
\end{aligned}$$

D Proofs for Section 5

D.1 Proof of Lemma 12 (spectral properties of smoothing operator S_G)

Proof. As in the invariant case, we note that for any degree k, the space $V_{d,k}$ of spherical harmonics of degree k is stable by S_G, *i.e.*, $S_G V_{d,k} \subset V_{d,k}$. Since S_G is self-adjoint, we may then find an orthonormal basis of such spherical harmonics, which we denote $\overline{Y}_{k,j}$, for $j = 1,\dots,N(d,k)$, such that the restriction of S_G to $V_{d,k}$ is diagonal, and we have $S_G\overline{Y}_{k,j} = \lambda_{k,j}\overline{Y}_{k,j}$, with $\lambda_{k,j} \geq 0$.

It remains to show (23). Define the operator $\Pi_k f = \mathbb{E}_y\left[P_{d,k}(\langle\cdot,y\rangle)f(y)\right]$. $S_G\Pi_k$ is then an integral operator with kernel

$$H(x,y) = \sum_{\sigma\in G} h(\sigma) P_{d,k}(\langle\sigma\cdot x, y\rangle). \tag{56}$$

Since $\overline{Y}_{k,j}$, $j = 1,\dots,N(d,k)$ forms an orthonormal basis of $V_{d,k}$, by the addition formula of spherical harmonics, we have

$$\Pi_k = \frac{1}{N(d,k)}\sum_{j=1}^{N(d,k)} \overline{Y}_{k,j}\overline{Y}^*_{k,j}.$$

It follows that

$$S_G \Pi_k = \frac{1}{N(d,k)} \sum_{j=1}^{N(d,k)} S_G \overline{Y}_{k,j} \overline{Y}^*_{k,j} = \frac{1}{N(d,k)} \sum_{j=1}^{N(d,k)} \lambda_{k,j} \overline{Y}_{k,j} \overline{Y}^*_{k,j}.$$

This implies that the kernel H of the operator $S_G \Pi_k$ can also be expressed as

$$H(x,y) = \frac{1}{N(d,k)} \sum_{j=1}^{N(d,k)} \lambda_{k,j} \overline{Y}_{k,j}(x) \overline{Y}_{k,j}(y). \tag{57}$$

Fixing $y = x$ and taking expectations over $x \sim \tau$ in both (56) and (57) proves the equality.

□

D.2 Proof of Theorem 13 (generalization with geometric stability)

The proof of the theorem is analogous to that of Theorem 5, replacing the control on $\mathcal{N}_{K_G}(\lambda)$ with that of Lemma 21 below, which provides an extension of Lemma 4 to generic smoothing operators, at the cost of a weaker constant $\nu_d(\ell)^{1/\alpha}$ instead of $\nu_d(\ell)$.

Lemma 21 (Degrees of freedom for K_G with stability.)**.** *Assume (A5). We have*

$$\mathcal{N}_K(\lambda) \leq C_K \lambda^{-1/\alpha},$$

and for any $\ell \geq 0$, we have

$$\mathcal{N}_{K_G}(\lambda) \leq D(\ell) + \nu_d(\ell)^{1/\alpha} C_K \lambda^{-1/\alpha}, \tag{58}$$

with the same constant C_K.

Proof. The first statement is a standard consequence of Assumption (A5). Namely, if ξ_m denote the eigenvalues of T_K (namely, the same as μ_k counted with their multiplicities) and $\xi_m \leq C(m+1)^{-\alpha}$, we have

$$\begin{aligned}
\mathcal{N}_K(\lambda) &= \sum_{m \geq 0} \frac{\xi_m}{\xi_m + \lambda} \\
&\leq \sum_{m \geq 0} \frac{1}{1 + \lambda C^{-1} (m+1)^\alpha} \\
&\leq \int_0^\infty \frac{dt}{1 + \lambda C^{-1} t^\alpha} \\
&\leq \frac{C^{1/\alpha} \lambda^{-1/\alpha}}{\alpha} \int_0^\infty \frac{u^{1/\alpha - 1} du}{1+u} = C_K \lambda^{-1/\alpha},
\end{aligned}$$

with $C_K := \frac{C^{1/\alpha}}{\alpha} \int_0^\infty \frac{u^{1/\alpha - 1} du}{1+u}$.

We now write

$$\begin{aligned}
\mathcal{N}_{K_G}(\lambda) &= \sum_{k \geq 0} \sum_{j=0}^{N(d,k)} \frac{\lambda_{k,j} \mu_k}{\lambda_{k,j} \mu_k + \lambda} \\
&= \sum_{k \geq 0} \sum_{j} \frac{\lambda_{k,j}}{\lambda_{k,j} + \lambda \mu_k^{-1}} \\
&\leq \sum_k N(d,k) \frac{\bar{\lambda}_k}{\bar{\lambda}_k + \lambda \mu_k^{-1}} \qquad \text{(by Jensen's inequality, with } \bar{\lambda}_k = N(d,k)^{-1} \sum_j \lambda_{k,j}) \\
&\leq \sum_k N(d,k) \frac{\bar{\lambda}_k \mu_k}{\bar{\lambda}_k \mu_k + \lambda},
\end{aligned}$$

We may then write, for some $\ell \geq 1$,

$$\mathcal{N}_{K_G}(\lambda) \leq D(\ell) + \sum_{k \geq 0} \frac{\bar{\mu}_k}{\bar{\mu}_k + \lambda},$$

where

$$\bar{\mu}_k = \begin{cases} \bar{\lambda}_k \mu_k, & \text{if } k \geq \ell \\ 0, & \text{o/w.} \end{cases}$$

Note that for $k \geq \ell$, we have $\bar{\lambda}_k = \frac{\sum_j \lambda_{k,j}}{N(d,k)} = \gamma_d(k) \leq \nu_d(\ell)$, and the same holds trivially for $k < \ell$. Then, writing $\bar{\xi}_m$ the collections of $\bar{\mu}_k$ counted with multiplicities, we may write $\bar{\xi}_m \leq \nu_d(\ell) C (m+1)^{-\alpha}$, with the same constant C as in (A5). Repeating the argument above for bounding $\mathcal{N}_{K_G}(\lambda)$ in terms of $\lambda^{-1/\alpha}$ then yields the result. □

D.3 Proof of Proposition 14 (upper bound on $\gamma_d(k)$ for deformations)

Proof. We first show that Φ_2 is stable under inversion, and later proceed to study lower bounds on its number of elements, and cycle statistics.

Step 1: $\Phi_2^{-1} = \Phi_2$. Let us first establish that Φ_2 is closed under inversion.

First observe that

$$\Phi_2 = \{\sigma; |\sigma(u+1) - \sigma(u) - 1| \leq 2 \; \forall \; u\} \; . \tag{59}$$

The inclusion LHS $\subseteq$ RHS is immediate by definition. The reverse inclusion is obtained by the triangle inequality, by observing that if $u < \tilde{u} < u'$, then

$$\begin{aligned} |\sigma(u) - \sigma(u') - (u - u')| &= |\sigma(u) - \sigma(\tilde{u}) - (u - \tilde{u}) + \sigma(\tilde{u}) - \sigma(u') - (\tilde{u} - u')| \\ &\leq |\sigma(u) - \sigma(\tilde{u}) - (u - \tilde{u})| + |\sigma(\tilde{u}) - \sigma(u') - (\tilde{u} - u')| \; , \end{aligned}$$

so by induction if the condition holds for small pairs $(u, \tilde{u})$, $(\tilde{u}, u')$ it extends to all pairs (u, u').

We directly verify from (59) that $\sigma \in \Phi_2$ iff it holds

$$\forall \; u \; , \; \sigma(u+1) = \sigma(u) + \{3, 2, 1, -1\} \; , \tag{60}$$

since we need to have $\sigma(u) \neq \sigma(u')$ whenever $u \neq u'$.

Let now $\tilde{u} = \sigma(u)$, so $\sigma^{-1}(\tilde{u}) = u$. We will show that σ^{-1} also verifies (60). We want to enumerate all possible u' so that $\sigma(u') = \tilde{u} + 1$. Clearly $\sigma^{-1}(\tilde{u}+1) \neq \sigma^{-1}(\tilde{u})$ so $u' \neq u$.

Suppose by contradiction that $u' < u - 1$. Note that we must have $\sigma(u'+1) \geq \tilde{u} + 2$ since $\tilde{u}$ and $\tilde{u}+1$ already have pre-images (namely u and u'), and smaller values would violate $\sigma(u'+1) - \sigma(u') \in \{3, 2, 1, -1\}$. Similarly $\sigma(u'+s) \geq \tilde{u} + 2$ for all $s = 2, \ldots, u - u' - 1$, since otherwise we would need a step $\sigma(u'+s+1) - \sigma(u'+s) \leq -3$, which is ruled out by (60). Then it must be that $\sigma(u) - \sigma(u-1) = \tilde{u} - \sigma(u-1) \leq -2$, which is a contradiction. We have thus shown $u' \geq u - 1$.

Similarly, let us show $u' \leq u + 3$. Assume, by contradiction that $u' > u + 3$. Note first that the only way to have $\sigma(u+s) < \tilde{u}$ for some $s \in [0, u' - u]$ is to only have $\sigma(u+1) = \tilde{u} - 1$, and $\sigma(u+s) \geq \tilde{u} + 2$ otherwise. Indeed, values smaller than $\tilde{u}$ must happen just following u in order to allow decreasing by 1, and having additional negative steps after $u+1$ (*e.g.*, $\sigma(u+2) = \tilde{u} - 2$) would require a step $\sigma(u+s+1) - \sigma(u+s) > 3$ for some $s \in [2, u' - u]$ (since the values $\tilde{u} - 1, \tilde{u}, \tilde{u} + 1$ already have pre-images, given by $u+1, u, u'$, respectively), which is a contradiction. Then, if $\sigma(u+1) = \tilde{u} - 1$, we must have $\sigma(u+2) = \tilde{u} + 2$ (since we cannot have longer steps), which implies $\sigma(u'-1) \geq \tilde{u} + 3$ and thus $\sigma(u') - \sigma(u'-1) \leq -2$, which is a contradiction. Alternatively, we must have $\sigma(u+s) \geq \tilde{u} + 2$ for all $s \in [1, u' - u - 1]$. This implies $\sigma(u'-1) = \tilde{u} + 2$, $\sigma(u'-2) = \tilde{u} + 3$, and more generally $\sigma(u'-t) = \tilde{u} + t + 1$, since these are the only allowed steps to obtain $\sigma(u') = \tilde{u} + 1$. Then, we have $\sigma(u+1) = \sigma(u' - (u' - u - 1)) = \tilde{u} + (u' - u) > \tilde{u} + 3$, which is in contradiction with $\sigma(u+1) - \sigma(u) \leq 3$. We have thus proved $u' \leq u + 3$. We thus have that σ^{-1} satisfies (60), which shows $\Phi_2^{-1} = \Phi_2$.

Step 2: Lower bound on $|\Phi_2|$. Denote as before $\sigma(u) = \tilde{u}$. By denoting $\Delta_v = \sigma(v) - v$ for arbitrary v, observe that $\Delta_{u'} - \Delta_{u'-1} \in \{2, 1, 0, -1\}$. Similarly we define $\Gamma_{\tilde{u}} := \sigma^{-1}(\tilde{u}) - \tilde{u}$. By the previous argument, we also have $\Gamma_{\tilde{u}+1} = \Gamma_{\tilde{u}} + \{2, 1, 0, -1\}$. Fix an arbitrary u_0, say $u_0 = 1$ and consider the subset of Φ_2 given by

$$\Phi_2^b = \{\sigma \in \Phi_2; \sigma(u_0) = u_0, \sigma(u_0 - 1) = u_0 - 1\} .$$

Φ_2^b thus contains permutations with 'fixed' boundary conditions. For $\sigma \in \Phi_2^b$, the boundary condition prevents $\Delta_{u_0+1} < \Delta_{u_0}$, so we identify the following possible cases:

1-block: $\Delta_{u_0+1} = \Delta_{u_0}$. In this case, $\sigma(u_0 + 1) = \sigma(u_0) + 1$.

2-block: $\Delta_{u_0+1} = \Delta_{u_0} + 1$. This implies $\Gamma_{\tilde{u_0}+2} = \Gamma_{\tilde{u_0}} - 1$, which in turn implies $\Gamma_{\tilde{u_0}+1} = \Gamma_{\tilde{u}_0} + 1$, and finally $\Delta_{u_0+2} = \Delta_{u_0+1} - 2$. In summary, $\sigma(u_0 + 1) = \sigma(u_0) + 2$ and $\sigma(u_0 + 2) = \sigma(u_0) + 1$.

3-block: $\Delta_{u_0+1} = \Delta_{u_0} + 2$. This implies $\Gamma_{\tilde{u_0}+3} = \Gamma_{\tilde{u_0}} - 2$, which necessarily implies $\Gamma_{\tilde{u_0}+1} = \Gamma_{\tilde{u_0}} + 2$, $\Gamma_{\tilde{u_0}+2} = \Gamma_{\tilde{u_0}+1} - 2$ and $\Gamma_{\tilde{u_0}+3} = \Gamma_{\tilde{u_0}+2} - 2$. This corresponds to $\sigma(u_0 + 1) = \sigma(u_0) + 3$, $\sigma(u_0 + 2) = \sigma(u_0) + 2$ and $\sigma(u_0 + 3) = \sigma(u_0) + 1$.

So an element of Φ_2^b can be constructed sequentially by assembling three possible 'blocks' B_i of size $i = \{1, 2, 3\}$. Moreover, we verify immediately that the following transitions are admissible:

$$B_1 \to B_{\{1,2,3\}} \ , \ B_2 \to B_{\{1,2\}} \ , \ B_3 \to B_1 \ .$$

Thus, by denoting $\mathcal{B}(m; B_i)$ the number of permutations in Φ_2^b restricted to their first m elements, and which that start (after u_0) with a block of type B_i, we have the following recursion:

$$\begin{aligned} \mathcal{B}(m; B_1) &= \mathcal{B}(m-1; B_1) + \mathcal{B}(m-1; B_2) + \mathcal{B}(m-1; B_3) \\ \mathcal{B}(m; B_2) &= \mathcal{B}(m-2; B_2) + \mathcal{B}(m-2; B_1) \\ \mathcal{B}(m; B_3) &= \mathcal{B}(m-3; B_1) \ , \end{aligned} \tag{61}$$

with $\mathcal{B}(i; B_i) = 1$. Let $F_i(z) := \sum_{m \geq 0} \mathcal{B}(m, B_i) z^m$ be the generating function associated to each of the above sequences. We have

$$F_1(z) = z^{-1}(F_1(z) + F_2(z) + F_3(z)) \ , \ F_2(z) = z^{-2}(F_1(z) + F_2(z)) \ , \ F_3(z) = z^{-3} F_1(z) \ .$$

By substituitng F_2, F_3 into the first equation, we obtain

$$F_1(z)(1 - z^{-1} - z^{-3}(1 - z^{-2})^{-1} - z^{-4}) = 0 \ ,$$

so F_1 has a pole at $\tau \approx 1.714$, the solution of the associated characterstic equation $z = 1 + \frac{1}{z^2-1} + \frac{1}{z^3}$. Moreover, this pole is also present in F_2 and F_3. This shows that $\mathcal{B}(m, B_i) \asymp C_i \tau^m$, and hence

$$|\Phi_2| \geq |\Phi_2^b| = \sum_{i=1}^{3} \mathcal{B}(d; B_i) = \Theta(\tau^d) \ .$$

Step 3: cycle statistics. Let us now compute a bound for $\gamma_d(k)$ using Corollary 8. Let

$$\xi(\Phi_2, n; d) = \{\sigma \in \Phi_2 \ : \ \mathrm{Fix}(\sigma) \geq n\}$$

denote the set of elements of Φ_2 that fix at least n positions, set $n = (1 - \eta)d$, and assume $\eta < 1/2$. Observe that this necessarily implies that two consecutive indices, say u_0 and $u_0 - 1$, are fixed, by the pigeonhole principle. Thus

$$\xi(\Phi_2, n) \subset \Phi_2^b \ ,$$

and we can use the characterisation of elements in Φ_2^b. We have

$$\begin{aligned} |\xi(\Phi_2, (1-\eta)d; d)| &\leq \sum_{n'=(1-\eta)d}^{d} \binom{d}{d-n'} |\mathcal{B}(n'; B_1) + \mathcal{B}(n'; B_2) + \mathcal{B}(n'; B_3)| \\ &\leq C\tau^{\eta d} \sum_{n'=(1-\eta)d}^{d} \binom{d}{d-n'} \leq C \left(e\eta^{-1}\right)^{\eta d} \tau^{\eta d} \ , \end{aligned}$$

where C is an abolute constant. Finally, from Example 11, we have

$$n < c(\sigma) \Rightarrow 2n - d < \mathrm{Fix}(\sigma) \ ,$$

thus $\zeta(\Phi_2, n) \leq \xi(\Phi_2, 2n-d)$. By picking $n = (1-\eta)d$ with $\eta < 1/4$, we have $2n - d = (1-2\eta)d > d/2$ and

$$\begin{aligned}
\gamma_d(k) &\leq \frac{\xi(\Phi_2, (1-2\eta)d)}{|\Phi_2|} + O\left(k^{-d+\max(n,d/2+7)}\right) && (62)\\
&\leq C(e(2\eta)^{-1})^{2\eta d}\tau^{(2\eta-1)d} + O\left(k^{-d+\max(n,d/2+7)}\right) \ . && (63)\\
&= C\left(\frac{e^{2\eta}}{(2\eta)^{2\eta}\tau^{1-2\eta}}\right)^d + O\left(k^{-\eta d}\right) \ . && (64)
\end{aligned}$$

When $2\eta < 0.15$, we verify that $\frac{e^{2\eta}}{(2\eta)^{2\eta}\tau^{1-2\eta}} < 1$.

□