

On the stability and accuracy of least squares approximations

Albert Cohen, Mark A. Davenport and Dany Leviatan*

June 19, 2018

Abstract

We consider the problem of reconstructing an unknown function f on a domain X from samples of f at n randomly chosen points with respect to a given measure ρ_X . Given a sequence of linear spaces $(V_m)_{m>0}$ with $\dim(V_m) = m \leq n$, we study the least squares approximations from the spaces V_m . It is well known that such approximations can be inaccurate when m is too close to n , even when the samples are noiseless. Our main result provides a criterion on m that describes the needed amount of regularization to ensure that the least squares method is stable and that its accuracy, measured in $L^2(X, \rho_X)$, is comparable to the best approximation error of f by elements from V_m . We illustrate this criterion for various approximation schemes, such as trigonometric polynomials, with ρ_X being the uniform measure, and algebraic polynomials, with ρ_X being either the uniform or Chebyshev measure. For such examples we also prove similar stability results using deterministic samples that are equispaced with respect to these measures.

1 Introduction and main results

Let X be a domain of \mathbb{R}^d and ρ_X be a probability measure on X . We consider the problem of estimating an unknown function $f : X \rightarrow \mathbb{R}$ from samples $(y_i)_{i=1, \dots, n}$ which are either noiseless or noisy observations of f at the points $(x_i)_{i=1, \dots, n}$, where the x_i are i.i.d. with respect to ρ_X . We measure the error between f and its estimator \tilde{f} in the $L^2(X, \rho_X)$ norm

$$\|v\| := \left(\int_X |v(x)|^2 d\rho_X(x) \right)^{1/2},$$

and we denote by $\langle \cdot, \cdot \rangle$ the associated inner product.

Given a fixed sequence of finite dimensional spaces $(V_m)_{m \geq 1}$ of $L^2(X, \rho_X)$ such that $\dim(V_m) = m$. We would like to compute the best approximation of f in V_m . This is given by the $L^2(X, \rho_X)$ orthogonal projector onto V_m , which we denote by P_m :

$$P_m f := \operatorname{argmin}_{v \in V_m} \|f - v\|.$$

We let

$$e_m(f) = \|f - P_m f\|$$

denote the best approximation error.

*This research has been partially supported by the ANR Defi08 ‘‘ECHANGE’’ and by the US NSF grant DMS-1004718. Portions of this work were completed while M.A.D. and D.L. were visitors at University Pierre et Marie Curie

In general, we may not have access to either ρ_X or any information about f aside from the observations at the points $(x_i)_{i=1,\dots,n}$. In this case we cannot explicitly compute $P_m f$. A natural approach in this setting is to consider the solution of the least squares problem

$$w = \operatorname{argmin}_{v \in V_m} \sum_{i=1}^n |y_i - v(x_i)|^2.$$

Typically, we are interested in the case where $m \leq n$ which is the regime where this problem may admit a unique solution.

In the noiseless case $y_i = f(x_i)$, and hence w may be viewed as the application of the least squares projection operator onto V_m to f , i.e., we can write

$$w = P_m^n f := \operatorname{argmin}_{v \in V_m} \|f - v\|_n$$

where

$$\|v\|_n := \left(\frac{1}{n} \sum_{i=1}^n |v(x_i)|^2 \right)^{1/2}$$

is the L^2 norm with respect to the empirical measure and, analogously, $\langle \cdot, \cdot \rangle_n$ the associated empirical inner product.

It is well known that least squares approximations may be inaccurate even when the measured samples are noiseless. For example, if V_m is the space \mathbb{P}_{m-1} of algebraic polynomials of degree $m - 1$ over the interval $[-1, 1]$ and if we choose $m = n$, this corresponds to Lagrange interpolation, which is known to be highly unstable, failing to converge towards f when given values at uniformly spaced samples, even when f is infinitely smooth (the ‘‘Runge phenomenon’’). Regularization by taking m substantially smaller than n may therefore be needed even in a noise-free context. The goal of this paper is to provide a mathematical analysis on the exact needed amount of such regularization.

Stability of the least squares problem. The solution of the least squares problem can be computed by solving an $m \times m$ system: specifically, if (L_1, \dots, L_m) is an arbitrary basis for V_m , then we can write

$$w = \sum_{j=1}^m u_j L_j,$$

where $\mathbf{u} = (u_j)_{j=1,\dots,m}$ is the solution of the $m \times m$ system

$$\mathbf{G}\mathbf{u} = \mathbf{f}, \tag{1.1}$$

with $\mathbf{G} := (\langle L_j, L_k \rangle_n)_{j,k=1,\dots,m}$ and $\mathbf{f} = (\frac{1}{n} \sum_{i=1}^n y_i L_k(x_i))_{k=1,\dots,m}$. In the noiseless case $y_i = f(x_i)$, so that we can also write $\mathbf{f} := (\langle f, L_k \rangle_n)_{k=1,\dots,m}$. In the event that \mathbf{G} is singular, we simply set $w = 0$.

For the purposes of our analysis, suppose that the basis (L_1, \dots, L_m) is orthonormal in the sense of $L^2(X, \rho_X)$.¹ In this case we have

$$\mathbf{E}(\mathbf{G}) = (\langle L_j, L_k \rangle)_{j,k=1,\dots,m} = \mathbf{I}.$$

¹While such a basis is generally not accessible when ρ_X is unknown, we require it only for the analysis. The actual computation of the estimator can be made using any known basis of V_m , since the solution w is independent of the basis used in computing it.

Our analysis requires an understanding of how the random matrix \mathbf{G} deviates from its expectation \mathbf{I} in probability. Towards this end, we introduce the quantity

$$K(m) := \sup_{x \in X} \sum_{j=1}^m |L_j(x)|^2.$$

Note that the function $\sum_{j=1}^m |L_j(x)|^2$ is invariant with respect to a rotation applied to (L_1, \dots, L_m) and therefore independent of the choice of the orthonormal basis: it only depends on the space V_m and on the measure ρ_X , and hence $K(m)$ also depends only on V_m and ρ_X . Also note that

$$K(m) \geq \sum_{j=1}^m \|L_j\|^2 = m.$$

We also will use the notation

$$\|\mathbf{M}\| = \max_{\mathbf{v} \neq 0} \frac{|\mathbf{M}\mathbf{v}|}{|\mathbf{v}|},$$

for the spectral norm of a matrix.

Our first result is a probabilistic estimate of the comparability of the norms $\|\cdot\|$ and $\|\cdot\|_n$ uniformly over the space V_m . This is equivalent to the proximity of the matrices \mathbf{G} and \mathbf{I} in spectral norm, since we have that for all $\delta \in [0, 1]$,

$$\|\mathbf{G} - \mathbf{I}\| \leq \delta \Leftrightarrow \left| \|v\|_n^2 - \|v\|^2 \right| \leq \delta \|v\|^2, \quad v \in V_m.$$

Theorem 1 *For $0 < \delta < 1$, one has the estimate*

$$\Pr \{ \|\mathbf{G} - \mathbf{I}\| > \delta \} = \Pr \{ \exists v \in V_m : \left| \|v\|_n^2 - \|v\|^2 \right| > \delta \|v\|^2 \} \leq 2m \exp \left\{ -\frac{c_\delta n}{K(m)} \right\}, \quad (1.2)$$

where $c_\delta := (1 + \delta) \log(1 + \delta) - \delta > 0$.

The proof of Theorem 1 is a simple application of tail bounds for sums of random matrices obtained in [1]. A consequence of this result is that the norms $\|\cdot\|$ and $\|\cdot\|_n$ are comparable with high probability if $K(m)$ is smaller than n by a logarithmic factor: for example taking $\delta = \frac{1}{2}$, we find that for any $r > 0$,

$$\Pr \left\{ \|\mathbf{G} - \mathbf{I}\| > \frac{1}{2} \right\} = \Pr \left\{ \exists v \in V_m : \left| \|v\|_n^2 - \|v\|^2 \right| > \frac{1}{2} \|v\|^2 \right\} \leq 2n^{-r}, \quad (1.3)$$

if m is such that

$$K(m) \leq \kappa \frac{n}{\log n}, \quad \text{with } \kappa := \frac{c_{1/2}}{1+r} = \frac{3 \log(3/2) - 1}{2 + 2r}. \quad (1.4)$$

The above condition thus ensures that \mathbf{G} is well conditioned with high probability. It can also be thought of as ensuring that the least squares problem is *stable* with high probability. Indeed the right side of the least squares system can be written as $\mathbf{f} = \mathbf{M}\mathbf{y}$ with

$$\mathbf{M} = \frac{1}{n} (L_j(x_i))_{j,i \in \{1, \dots, m\} \times \{1, \dots, n\}},$$

an $m \times n$ matrix. Observing that $\langle \mathbf{G}\mathbf{v}, \mathbf{v} \rangle = n |\mathbf{M}^T \mathbf{v}|^2$, we find that

$$\|\mathbf{M}\| = \|\mathbf{M}^T\| = \left(\frac{1}{n} \|\mathbf{G}\| \right)^{1/2}.$$

Therefore, if $\|\mathbf{G} - \mathbf{I}\| \leq \frac{1}{2}$, then we have that for any data vector \mathbf{y} the solution $w = \sum_{j=1}^m u_j L_j$ satisfies

$$\|w\| = |\mathbf{u}| \leq \|\mathbf{G}^{-1}\| \cdot \|\mathbf{M}\| \cdot |\mathbf{y}| \leq \frac{1}{\sqrt{n}} 2\sqrt{\frac{3}{2}} |\mathbf{y}|,$$

which thus gives the stability estimate

$$\|w\| \leq C \left(\frac{1}{n} \sum_{i=1}^n |y_i|^2 \right)^{1/2}, \quad C = \sqrt{6}.$$

In the noiseless case, this can be written as $\|P_m^n f\| \leq C \|f\|_n$, i.e., the least squares projection is stable between the norms $\|\cdot\|_n$ and $\|\cdot\|$. Note that since $K(m)$ not only depends on V_m but also on the measure ρ_X , the range of m such that the condition (1.4) holds is strongly tied to the choice of the measure. This issue is illustrated further in our numerical experiments.

Let us mention that similar probabilistic bounds have been previously obtained, see in particular §5.2 in [2]. These earlier results allow us to obtain the bound (1.3), however relying on the stronger condition

$$K(m) \lesssim \left(\frac{n}{\log(n)} \right)^{1/2}.$$

The numerical results for polynomial least squares that we present in §3 hint that the weaker condition $K(m) \lesssim \frac{n}{\log(n)}$ is sharp. The quantity $K(m)$ was also used in [3] in order to control the $L^\infty(X)$ norm and the $L^2(\rho_X)$ norm.

Accuracy of least squares approximation. As an application, we can derive an estimate for the error of least squares approximation in expectation. Here, we make the assumption that a uniform bound

$$|f(x)| \leq L, \tag{1.5}$$

holds for almost every x with respect to ρ_X . For $m \leq n$, we consider the truncated least squares estimator

$$\tilde{f} = T_L(w),$$

where $T_L(t) = \text{sign}(t) \max\{L, |t|\}$. Our first result deals with the noiseless case.

Theorem 2 *In the noiseless case, for any $r > 0$, if m is such that the condition (1.4) holds, then*

$$\mathbb{E}(\|f - \tilde{f}\|^2) \leq (1 + \varepsilon(n)) e_m(f)^2 + 8L^2 n^{-r}, \tag{1.6}$$

where $\varepsilon(n) := \frac{4\kappa}{\log(n)} \rightarrow 0$ as $n \rightarrow +\infty$, with κ as in (1.4)

At this point a few remarks are due regarding the implications of this result in terms of the convergence rate of the estimate.

Consider the following general setting of regression on a random design: we observe independent samples

$$z_i = (x_i, y_i)_{i=1, \dots, n} \tag{1.7}$$

of a variable $z = (x, y)$ of law ρ over $X \times Y$ and marginal law ρ_X over X , and we want to estimate from these samples the regression function defined as the conditional expectation

$$f(x) := \mathbb{E}(y|x). \tag{1.8}$$

We assume that the maximal variance

$$\sigma^2 := \sup_{x \in X} \mathbb{E}(|y - f(x)|^2 | x), \quad (1.9)$$

is bounded. We thus think of the y_i as noisy observations of f at x_i with additive noise of variance at most σ^2 , namely

$$y_i = f(x_i) + \eta_i, \quad (1.10)$$

where the η_i are independent realizations of the variable $\eta := y - f(x)$.

Assuming that f satisfies the uniform bound (1.5), one computes the truncated least squares estimator now with y_i in place of $f(x_i)$. A typical convergence bound for this estimator, see for example Theorem 11.3 in [6], is

$$\mathbb{E}(\|f - \tilde{f}\|^2) \leq C \left(e_m(f)^2 + \max\{L^2, \sigma^2\} \frac{m \log n}{n} \right). \quad (1.11)$$

Convergence rates may be found after balancing the two terms, but they are limited by the optimal learning rate n^{-1} , and this limitation persists even in the noiseless case $\sigma^2 = 0$ due to the presence of L^2 in the right side of (1.11). In contrast, Theorem 2 yields fast convergence rates, provided that the approximation error e_m has fast decay and that the value of m satisfying (1.4) can be chosen large enough.

One motivation for studying the noiseless case is the numerical treatment of parameter dependent PDEs of the general form

$$\mathcal{F}(f, x) = 0,$$

where x is a vector of parameters in some compact set $\mathcal{P} \in \mathbb{R}^d$. We can consider the solution map $x \mapsto f(x)$ either as giving the exact solution to the PDE for the given value of the parameter vector x or as the exact result of a numerical solver for this value of x . In the stochastic PDE context, x is random and obeys a certain law which may be known or unknown. From a random draw $(x_i)_{i=1, \dots, n}$, we obtain solutions $f_i = f(x_i)$ which are noiseless observations of the solution map, and are interested in reconstructing this map. In instances such as elliptic problems with parameters in the diffusion coefficients, the solution map can be well-approximated by polynomials in x (see [4]). In this context, an initial study of the needed amount of regularization was given in [7], however specifically targeted towards polynomial least squares.

For the noisy regression problem described above, our analysis can also be adapted in order to derive the following result.

Theorem 3 *For any $r > 0$, if m is such that the condition (1.4) holds, then*

$$\mathbb{E}(\|f - \tilde{f}\|^2) \leq (1 + 2\varepsilon(n))e_m(f)^2 + 8L^2n^{-r} + 8\sigma^2\frac{m}{n}, \quad (1.12)$$

with $\varepsilon(n)$ as in Theorem 2 and σ is the maximal variance given by (1.9).

In the noiseless case, the bound in Theorem 2 suggests that m should be chosen as large as possible under the constraint that (1.4) holds. In the noisy case, the value of m minimizing the bound in Theorem 3 also depends on the decay of e_m , which is generally unknown. In such a situation, a classical way of choosing the value of m is by a model selection procedure, such as adding a complexity penalty in the least squares or using an independent validation sample. Such procedures can also be of interest in the noiseless case when the measure ρ_X is unknown, since the maximal value of m such that (1.4) holds is then also unknown.

Let us give an example of how the results in Theorems 2 and 3 lead to specific rates of convergence in terms of the number of samples: assume that $X = [-1, 1]$ is equipped with the uniform measure $\rho_X = \frac{dx}{2}$ and that $V_m = \mathbb{P}_{m-1}$ is the space of algebraic polynomials of degree $m - 1$. Then, if f belongs to $C^r(X)$ the

space of r -times differentiable functions, it is well-known that $e_m(f)^2 \lesssim m^{-2r}$. On the one hand the results in §3 show that condition (1.4) can be ensured with $m \sim (n/\log n)^{1/2}$. Therefore, in the noiseless case, we obtain a bound proportional to n^{-r} for the mean squared error, up the logarithmic factor. In the noisy case, after balancing the approximation and variance terms, we obtain a bound proportional to $\sigma^{2r/(r+1)}n^{-r/(r+1)}$. On the other hand, these rates can be improved with r replaced by $2r$ if we use the Chebyshev non-uniform measure that concentrates near the end-points, since in that case the results in §3 show that condition (1.4) can be ensured with $m \sim n/\log n$.

The rest of our paper is organized as follows: we give the proofs of the above results in §2 and we present in §3 examples of applications to classical approximation schemes such as piecewise constants, trigonometric polynomials, or algebraic polynomials. For such examples, we study the range of m such that (1.4) holds and show that this range is in accordance with stability results that can be proved for deterministic sampling. Numerical illustrations are given for algebraic polynomial approximation.

2 Proofs

Proof of Theorem 1: The matrix \mathbf{G} can be written as

$$\mathbf{G} = \mathbf{X}_1 + \cdots + \mathbf{X}_n,$$

where the \mathbf{X}_i are i.i.d. copies of the random matrix

$$\mathbf{X} = \frac{1}{n}(L_j(x)L_k(x))_{j,k=1,\dots,m},$$

where x is distributed according to ρ_X . We use the following Chernoff bound from [8], originally obtained by [1]: if $\mathbf{X}_1, \dots, \mathbf{X}_n$ are independent $m \times m$ random self-adjoint and positive matrices satisfying

$$\lambda_{\max}(\mathbf{X}_i) = \|\mathbf{X}_i\| \leq R,$$

almost surely, then with

$$\mu_{\min} := \lambda_{\min}\left(\sum_{i=1}^n \mathbb{E}(\mathbf{X}_i)\right) \quad \text{and} \quad \mu_{\max} := \lambda_{\max}\left(\sum_{i=1}^n \mathbb{E}(\mathbf{X}_i)\right),$$

one has

$$\Pr \left\{ \lambda_{\min}\left(\sum_{i=1}^n \mathbf{X}_i\right) \leq (1 - \delta)\mu_{\min} \right\} \leq m \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{\mu_{\min}/R}, \quad 0 \leq \delta < 1,$$

and

$$\Pr \left\{ \lambda_{\max}\left(\sum_{i=1}^n \mathbf{X}_i\right) \geq (1 + \delta)\mu_{\max} \right\} \leq m \left(\frac{e^{\delta}}{(1 + \delta)^{1+\delta}} \right)^{\mu_{\max}/R}, \quad \delta \geq 0$$

In our present case, we have $\sum_{i=1}^n \mathbb{E}(\mathbf{X}_i) = n\mathbb{E}(\mathbf{X}) = \mathbf{I}$ so that $\mu_{\min} = \mu_{\max} = 1$. It is easily checked that $\frac{e^{\delta}}{(1+\delta)^{1+\delta}} \geq \frac{e^{-\delta}}{(1-\delta)^{1-\delta}}$ for $0 < \delta < 1$, and therefore

$$\Pr \{ \|\mathbf{G} - \mathbf{I}\| > \delta \} \leq 2m \left(\frac{e^{\delta}}{(1 + \delta)^{1+\delta}} \right)^{1/R} = 2m \exp\left(-\frac{c\delta}{R}\right).$$

We next use the fact that a rank 1 symmetric matrix $ab^T = (b_j a_k)_{j,k=1,\dots,m}$ has its spectral norm equal to the product of the Euclidean norms of the vectors a and b , and therefore

$$\|\mathbf{X}\| \leq \frac{1}{n} \sum_{j=1}^m |L_j(x)|^2 = \frac{K(m)}{n},$$

almost surely. We may therefore take $R = \frac{K(m)}{n}$ which concludes the proof. \square

Proof of Theorem 2: We denote by $d\rho_X^n := \otimes^n d\rho_X$ the probability measure of the draw. We also denote by Ω the set of all possible draws, that we divide into the set Ω_+ of all draw such that

$$\|\mathbf{G} - \mathbf{I}\| \leq \frac{1}{2},$$

and the complement set $\Omega_- := \Omega \setminus \Omega_+$. According to (1.3), we have

$$\Pr\{\Omega_-\} = \int_{\Omega_-} d\rho_X^n \leq 2n^{-r}, \quad (2.1)$$

under the condition (1.4). This leads to

$$\mathbb{E}(\|f - \tilde{f}\|^2) = \int_{\Omega} \|f - \tilde{f}\|^2 d\rho_X^n \leq \int_{\Omega_+} \|f - P_m^n f\|^2 d\rho_X^n + 8L^2 n^{-r},$$

where we have used $\|f - \tilde{f}\|^2 \leq 2L^2$, as well as the fact that T_L is a contraction that preserves f .

It remains to prove that the first term in the above right side is bounded by $(1 + \varepsilon(n))e_m(f)^2$. With $g := f - P_m f$, we observe that

$$f - P_m^n f = f - P_m f + P_m^n P_m f - P_m^n f = g - P_m^n g.$$

Since g is orthogonal to V_m , we thus have

$$\|f - P_m^n f\|^2 = \|g\|^2 + \|P_m^n g\|^2 = \|g\|^2 + \sum_{j=1}^m |a_j|^2,$$

where $\mathbf{a} = (a_j)_{j=1, \dots, m}$ is solution of the system

$$\mathbf{G}\mathbf{a} = \mathbf{b},$$

with $\mathbf{b} := (\langle g, L_k \rangle_n)_{k=1, \dots, m}$. When the draw belongs to Ω_+ , we have $\|\mathbf{G}^{-1}\|_2 \leq 2$ and therefore

$$\sum_{j=1}^m |a_j|^2 \leq 4 \sum_{k=1}^m |\langle g, L_k \rangle_n|^2.$$

It follows that

$$\int_{\Omega_+} \|f - P_m^n f\|^2 d\rho_X^n \leq \int_{\Omega_+} \left(\|g\|^2 + 4 \sum_{k=1}^m |\langle g, L_k \rangle_n|^2 \right) d\rho_X^n \leq \|g\|^2 + 4 \sum_{k=1}^m \mathbb{E}(|\langle g, L_k \rangle_n|^2).$$

We estimate each of the $\mathbb{E}(|\langle g, L_k \rangle_n|^2)$ as follows:

$$\begin{aligned} \mathbb{E}(|\langle g, L_k \rangle_n|^2) &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}(g(x_i)g(x_j)L_k(x_i)L_k(x_j)) \\ &= \frac{1}{n^2} \left(n(n-1) |\mathbb{E}(g(x)L_k(x))|^2 + n \mathbb{E}(|g(x)L_k(x)|^2) \right) \\ &= \left(1 - \frac{1}{n} \right) |\langle g, L_k \rangle|^2 + \frac{1}{n} \int_X |g(x)|^2 |L_k(x)|^2 d\rho_X \\ &= \frac{1}{n} \int_X |g(x)|^2 |L_k(x)|^2 d\rho_X, \end{aligned}$$

where we have used the fact that g is orthogonal to V_m and thus to L_k . Summing over k , we obtain

$$\sum_{k=1}^m \mathbb{E}(|\langle g, L_k \rangle_n|^2) \leq \frac{K(m)}{n} \|g\|^2 \leq \frac{\kappa}{\log(n)} \|g\|^2,$$

where we have used (1.4). We have thus proven that

$$\int_{\Omega_+} \|f - P_m^n f\|^2 d\rho_X^n \leq \left(1 + \frac{4\kappa}{\log(n)}\right) \|g\|^2 = (1 + \varepsilon(n)) e_m(f)^2,$$

which concludes the proof. \square

Proof of Theorem 3: We define the additive noise in the sample by writing

$$y_i = f(x_i) + \eta_i,$$

and thus the η_i are i.i.d. copies of the variable

$$\eta = y - f(x).$$

Note that η and x are not assumed to be independent. However we have

$$\mathbb{E}(\eta|x) = 0,$$

which implies the decorrelation property

$$\mathbb{E}(\eta h(x)) = 0,$$

for any function h . As in the proof of Theorem 2 we split Ω into Ω_+ and Ω_- and find that

$$\mathbb{E}(\|f - \tilde{f}\|^2) \leq \int_{\Omega_+} \|f - w\|^2 d\rho_X^n + 8M^2 n^{-r},$$

where w now stands for the solution to the least squares problem with noisy data (y_1, \dots, y_n) . With the same definition of $g = f - P_m f$, we can write

$$f - w = g - P_m^n g - \tilde{w},$$

where \tilde{w} stands for the solution to the least squares problem for the noise data (η_1, \dots, η_n) . Therefore

$$\|f - P_m^n f\|^2 = \|g\|^2 + \|P_m^n g + \tilde{w}\|^2 \leq \|g\|^2 + 2\|P_m^n g\|^2 + 2\|\tilde{w}\|^2 = \|g\|^2 + 2 \sum_{j=1}^m |a_j|^2 + 2 \sum_{j=1}^m |d_j|^2,$$

where $\mathbf{a} = (a_j)_{j=1, \dots, m}$ is as in the proof of Theorem 2 and $\mathbf{d} = (d_j)_{j=1, \dots, m}$ is solution of the system

$$\mathbf{G}\mathbf{d} = \mathbf{n},$$

with $\mathbf{n} := (\frac{1}{n} \sum_{i=1}^n \eta_i L_k(x_i))_{k=1, \dots, m} = (n_k)_{k=1, \dots, m}$. By the same arguments as in the proof of Theorem 2, we thus obtain

$$\mathbb{E}(\|f - \tilde{f}\|^2) \leq (1 + 2\varepsilon(n)) e_m(f)^2 + 8L^2 n^{-r} + 8 \sum_{k=1}^m \mathbb{E}(|n_k|^2).$$

We are left to show that $\sum_{k=1}^m \mathbb{E}(|n_k|^2) \leq \frac{\sigma^2 m}{n}$. For this we simply write that

$$\mathbb{E}(|n_k|^2) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}(\eta_i L_k(x_i) \eta_j L_k(x_j)).$$

For $i \neq j$, we have

$$\mathbb{E}(\eta_i L_k(x_i) \eta_j L_k(x_j)) = (\mathbb{E}(\eta L_k(x)))^2 = 0.$$

For $i = j$, we have

$$\begin{aligned} \mathbb{E}(|\eta_i L_k(x_i)|^2) &= \mathbb{E}(|\eta L_k(x)|^2) \\ &= \int_X \mathbb{E}(|\eta L_k(x)|^2 | x) d\rho_X \\ &= \int_X \mathbb{E}(|\eta|^2 | x) |L_k(x)|^2 d\rho_X \\ &\leq \sigma^2 \int_X |L_k(x)|^2 d\rho_X = \sigma^2. \end{aligned}$$

It follows that $\mathbb{E}(|n_k|^2) \leq \frac{\sigma^2}{n}$, which concludes the proof. \square

3 Examples and numerical illustrations

We now give several examples of approximation schemes for which one can compute the quantity $K(m)$ and therefore estimate the range of m such that the condition (1.4) holds. For each of these examples, we also exhibit a deterministic sampling (x_1, \dots, x_n) for which the stability property

$$\|\mathbf{G} - \mathbf{I}\| \leq \frac{1}{2},$$

or equivalently

$$|\|v\|_n^2 - \|v\|^2| \leq \frac{1}{2} \|v\|^2, \quad v \in V_m,$$

is ensured for the same range of m (actually slightly better by a logarithmic factor). For the sake of simplicity, we work in the one dimensional setting, with X a bounded interval.

Piecewise constant functions. Here $X = [a, b]$ and V_m is the space of piecewise constant functions over a partition of X into intervals I_1, \dots, I_m . In such a case, an orthonormal basis with respect to $L^2(X, \rho_X)$ is given by the characteristic functions $L_k := (\rho_X(I_k))^{-1/2} \chi_{I_k}$, and therefore

$$K(m) = \max_{k=1, \dots, m} (\rho_X(I_k))^{-1}.$$

Given a measure ρ_X , the partition that minimizes $K(m)$, and therefore allows us to fulfill (1.4) for the largest range of m , is one that evenly distributes the measure ρ_X . With such partitions, $K(m)$ reaches its minimal value

$$K(m) = m,$$

and (1.4) can be achieved with $m \sim \frac{n}{\log n}$.

If we now choose $n = m$ deterministic points x_1, \dots, x_m with $x_k \in I_k$, we clearly have

$$\|v\|_n^2 = \|v\|^2, \quad v \in V_m.$$

Therefore the stability of the least squares problem can be ensured with m up to the value n using a deterministic sample.

Trigonometric polynomials and uniform measure. Without loss of generality, we take $X = [-\pi, \pi]$, and we consider for odd $m = 2p + 1$ the space V_m of trigonometric polynomials of degree p , which is spanned by the functions $L_k(x) = e^{ikx}$ for $k = -p, \dots, p$. Assuming that ρ_X is the uniform measure, this is an orthonormal basis with respect to $L^2(X, \rho_X)$. In this example, we again obtain the minimal value

$$K(m) = m.$$

Therefore (1.4) can be achieved with $m \sim \frac{n}{\log n}$.

We now consider the deterministic uniform sampling $x_i := -\pi + \frac{2\pi i}{n}$ for $i = 1, \dots, n$. With such a sampling, one has the identity

$$\int_{-\pi}^{\pi} v(x) d\rho_X = \frac{1}{2\pi} \int_{-\pi}^{\pi} v(x) dx = \frac{1}{n} \sum_{i=1}^n v(x_i),$$

for all trigonometric polynomials v of degree $n - 1$ (this is easily seen by checking the identity on every basis element). When $v \in V_m$ with $m = 2p + 1$, we know that $|v|^2$ is a trigonometric polynomial of degree $2p$. We thus find that

$$\|v\|_n^2 = \|v\|^2, \quad v \in V_m,$$

provided that $2p \leq n - 1$, or equivalently $m \leq n$. Therefore the stability of the least squares problem can be ensured with m up to the value n using a deterministic sample.

Algebraic polynomials and uniform measure. Without loss of generality, we take $X = [-1, 1]$, and we consider $V_m = \mathbb{P}_{m-1}$ the space of algebraic polynomials of degree $m - 1$. When ρ_X is the uniform measure, an orthonormal basis is given by defining L_k as the Legendre polynomial of degree $k - 1$ with normalization

$$\|L_k\|_{L^\infty([-1,1])} = |L_k(1)| = \sqrt{2k - 1},$$

and thus

$$K(m) = \sum_{k=1}^m (2k - 1) = m^2.$$

Therefore (1.4) can be achieved with $m \sim \sqrt{\frac{n}{\log n}}$ which is a lower range compared to the previous examples. We now consider the deterministic sampling obtained by partitioning X into n intervals (I_1, \dots, I_n) of equal

length $\frac{2}{n}$, and picking one point x_i in each I_i . For any $v \in V_m$, we may write

$$\begin{aligned}
\left| \int_{I_i} |v(x)|^2 d\rho_X - \frac{1}{n} |v(x_i)|^2 \right| &= \left| \frac{1}{2} \int_{I_i} |v(x)|^2 dx - \frac{1}{n} |v(x_i)|^2 \right| \\
&= \left| \frac{1}{2} \int_{I_i} (|v(x)|^2 - |v(x_i)|^2) dx \right| \\
&\leq \frac{1}{2} \int_{I_i} \left| |v(x)|^2 - |v(x_i)|^2 \right| dx \\
&\leq \frac{1}{2} |I_i| \int_{I_i} |(v^2)'(x)| dx \\
&= \frac{2}{n} \int_{I_i} |v'(x)v(x)| d\rho_X.
\end{aligned}$$

Summing over i , it follows that

$$\left| \|v\|_n^2 - \|v\|^2 \right| \leq \frac{2}{n} \int_X |v'(x)v(x)| d\rho_X \leq \frac{2}{n} \|v'\| \|v\| \leq \frac{2(m-1)^2}{n} \|v\|^2,$$

where we have used the Cauchy-Schwarz and Markov inequalities. Therefore the stability of the least squares problem can be ensured with m up to the value $\frac{\sqrt{n}}{2} + 1$ using a deterministic sample.

Algebraic polynomials and Chebyshev measure. Consider again algebraic polynomials of degree $m-1$ on $X = [-1, 1]$, now equipped with the measure

$$d\rho_X = \frac{dx}{\pi\sqrt{1-x^2}}.$$

Then an orthonormal basis is given by defining L_k as the Chebyshev polynomial of degree $k-1$, with $L_1 = 1$ and

$$L_k(x) = \sqrt{2} \cos((k-1) \arccos x),$$

for $k > 1$, and thus

$$K(m) = 2m - 1.$$

Therefore (1.4) can be achieved with $m \sim \frac{n}{\log n}$, which expresses the fact that least squares approximations are stable for higher polynomial degrees when working with the Chebyshev measure rather than with the uniform measure.

We now consider the deterministic sampling obtained by partitioning X into n intervals (I_1, \dots, I_n) of equal

Chebyshev measure $\rho_X(I_i) = \frac{1}{n}$, and picking one point x_i in each I_i . For any $v \in V_m$, we may write

$$\begin{aligned} \left| \int_{I_i} |v(x)|^2 d\rho_X - \frac{1}{n} |v(x_i)|^2 \right| &= \left| \int_{I_i} (|v(x)|^2 - |v(x_i)|^2) d\rho_X \right| \\ &\leq \int_{I_i} ||v(x)|^2 - |v(x_i)|^2| d\rho_X \\ &\leq \rho_X(I_i) \int_{I_i} |(v^2)'(x)| dx \\ &= \frac{1}{n} \int_{I_i} |v'(x)v(x)| dx. \end{aligned}$$

Summing over i , it follows that

$$|\|v\|_n^2 - \|v\|^2| \leq \frac{1}{n} \int_X |v'(x)v(x)| dx \leq \frac{1}{n} \|v\| \left(\pi \int_X |v'(x)|^2 \sqrt{1-x^2} dx \right)^{1/2}.$$

Using the change of variable $x = \cos t$, it is easily seen that the inverse estimate

$$\int_X |v'(x)|^2 \sqrt{1-x^2} dx \leq (m-1)^2 \int_X |v(x)|^2 \frac{1}{\sqrt{1-x^2}} dx,$$

holds for any $v \in V_m$. Therefore

$$|\|v\|_n^2 - \|v\|^2| \leq \frac{1}{n} \int_X |v'(x)v(x)| dx \leq \frac{\pi(m-1)}{n} \|v\|^2$$

which shows that the stability of the least squares problem can be ensured with m up to the value $\frac{n}{2\pi} + 1$ using a deterministic sample.

Let us observe that in several practical scenarios, the measure ρ_X of the observations may be unknown to us, therefore raising the question of the behavior of $K(m)$ for an arbitrary measure.

It is not too difficult to check that when the space V_m is not the trivial space of constant functions (which is the case as soon as $m \geq 2$) the quantity $K(m)$ may become arbitrarily large for certain measures ρ_X . We leave the proof of this general fact as an exercise for the reader, and rather provide a simple illustration: consider the space V_2 of polynomials of degree 1 on $[-1, 1]$ and the measure $\rho_X = \frac{1}{2\varepsilon} \chi_{[-\varepsilon, \varepsilon]}(x) dx$ where $\varepsilon > 0$ is small. Then an orthonormal basis is provided by the functions $L_0(x) = 1$ and $L_1(x) = \frac{\sqrt{3}}{\varepsilon} x$, so that $K(m) \sim \varepsilon^{-2}$. An interesting problem is to understand if for certain families of space (V_m) , the quantity $K(m)$ can be controlled under fairly general assumptions on the measure ρ_X . One typical such assumption is the the strong density assumption, which states that

$$\rho_X(E) \sim |E|, \quad E \text{ measurable}, \quad (3.1)$$

where $|\cdot|$ is the Lebesgue measure. In the case of piecewise constant functions on uniform partitions, or for more general spline functions on uniform grids, it is not difficult to check that this assumption implies the behavior $K(m) \sim m$.

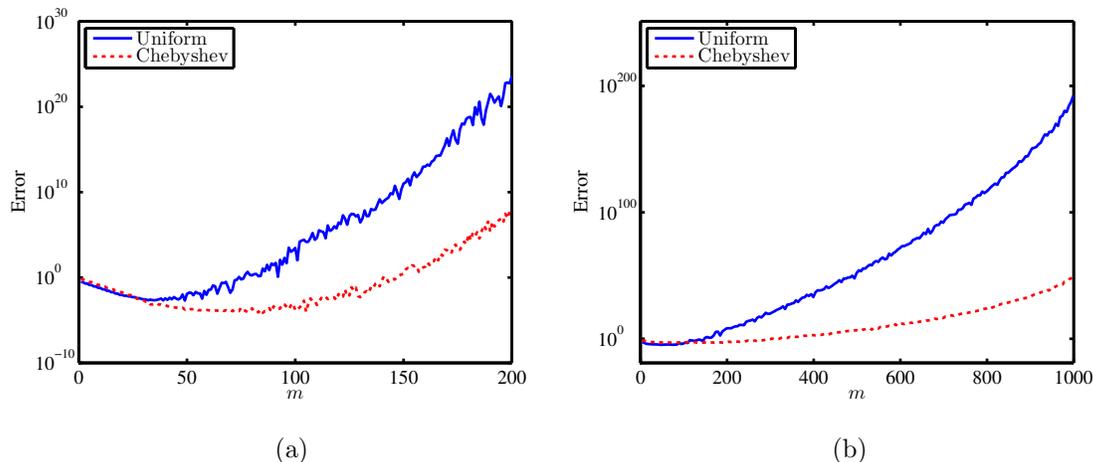


Figure 3.1: The $L^2(X, \rho_X)$ error as m varies (a) for f_1 and (b) for f_2 .

Numerical illustration. We conclude with a brief numerical illustration of our theoretical results for the setting of algebraic polynomials. Specifically, we consider the smooth function $f_1(x) = 1/(1 + 25x^2)$ originally considered by Runge to illustrate the instability of polynomial interpolation at equispaced points, and the non-smooth function $f_2(x) = |x|$, both restricted to the interval $[-1, 1]$.

For both functions, we take n i.i.d. samples x_1, \dots, x_n with respect to a measure ρ_X on $X = [-1, 1]$ and compute the noise-free observations $y_i = f(x_i)$. We consider either the uniform measure $\rho_X := \frac{dx}{2}$ or the Chebyshev measure $\rho_X := \frac{dx}{\pi\sqrt{1-x^2}}$. In both cases, we compute the least squares approximating polynomial of degree m using these points for a range of different values of $m \leq n$. We then numerically compute the error in the $L^2(X, \rho_X)$ norm, with ρ_X the corresponding measure in which the sample have been drawn, using the adaptive Simpson's quadrature rule [5] implemented in Matlab.

Figure 3.1 shows the results of this simulation using $n_1 = 200$ samples for estimating f_1 and $n_2 = 1000$ samples for estimating f_2 . We observe that, in all cases, as m approaches n the solutions become highly inaccurate due to the inherent instability of the problem. However, we can set m to be much larger before instability starts to develop when the points are drawn with respect to the Chebyshev measure, as is expected.

Next we consider the effect of n on the best choice of m . Specifically, for any given sample of points we can compute the value $m(n)$ that corresponds to the polynomial degree for which we obtain the best approximation to f_1 or f_2 and examine how this behaves as a function of n . This is shown in Figure 3.2, that displays as a function of n the average value of $m(n)$ over 50 realizations of the sample, for both measures and both functions f_1 and f_2 (the averaging has the effect of reducing oscillation in the curve $n \mapsto m(n)$ making it more readable). We vary the sample size from $n = 1$ to 1000 for f_2 , but only from $n = 1$ to 200 for the smooth function f_1 , since in that case the $L^2(X, \rho_X)$ error drops below machine precision for larger values of n with m in the regime where the least squares problem is stable and therefore the minimal value $m(n)$ cannot be precisely located.

We observe that, in accordance with our theoretical results, $m(n)$ behaves like \sqrt{n} when the points are drawn with respect to the uniform measure, while it behaves almost linear in n when the points are drawn with respect to the Chebyshev measure.

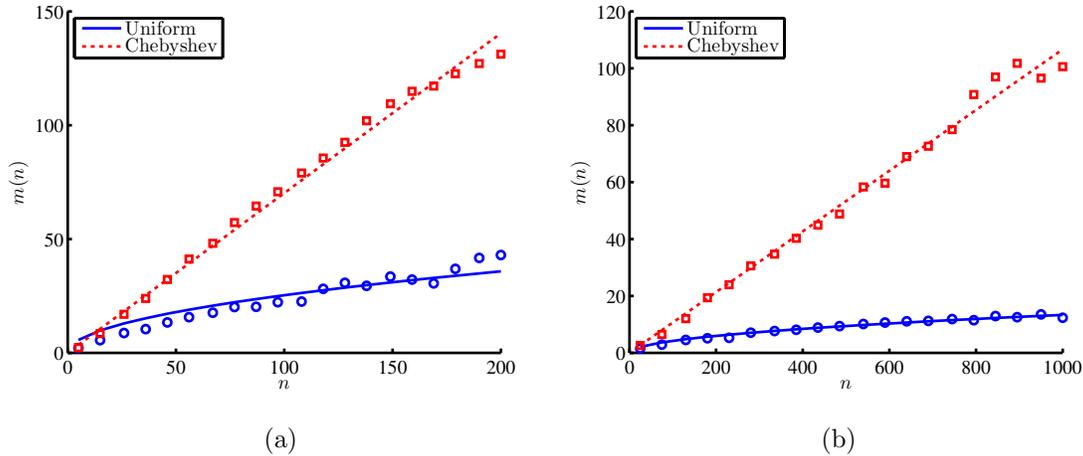


Figure 3.2: Optimal values $m(n)$ as n varies (a) for f_1 (comparison with $0.7n$ and $2.5\sqrt{n}$) and (b) for f_2 (comparison with $0.1n$ and $0.4\sqrt{n}$).

References

- [1] Ahlswede, R. and A. Winter, *Strong converse for identification via quantum channels*, IEEE Trans. Information Theory **48**, 569–579, 2002.
- [2] Baraud, Y., *Model selection for regression on a random design*, ESAIM Prob. Stat. **6**, 127–146, 2002.
- [3] Birgé L. and P. Massart, *Minimum contrast estimators on sieves: Exponential bounds and rates of convergence*, Bernoulli **4**, 329–375, 1998.
- [4] Cohen, A., R. DeVore and C. Schwab, *Analytic regularity and polynomial approximation of parametric elliptic PDE’s*, Analysis and Applications **9**, 11–47, 2011.
- [5] Gander, W. and W. Gautschi, *Adaptive Quadrature – Revisited*, BIT **40**, 84–101, 2000.
- [6] Györfi, L., M. Kohler, A. Krzyzak, A. and H. Walk, *A distribution-free theory of nonparametric regression*, Springer, Berlin, 2002.
- [7] Migliorati, G., F. Nobile, E. von Schweriny and R. Tempone, *Analysis of the point collocation method*, preprint MOX, Politecnico di Milano, 2011.
- [8] Tropp, J. *User friendly tail bounds for sums of random matrices*, to appear in J. FoCM, 2011.

The authors would like to thank Lukas Meier for bringing a small error in the original proof of Theorem 1 to our attention.

Albert Cohen
 Laboratoire Jacques-Louis Lions
 Université Pierre et Marie Curie
 4, Place Jussieu, 75005 Paris, France
 cohen@ann.jussieu.fr

Mark Davenport
School of Electrical and Computer Engineering
Georgia Institute of Technology
777 Atlantic Drive NW
Atlanta, GA 30332, USA
mdav@gatech.edu

Dany Leviatan
Raymond and Beverly Sackler School of Mathematics
Tel Aviv University
69978, Tel Aviv, Israel
leviatan@post.tau.ac.il