
Learning Sparse Additive Models with Interactions in High Dimensions

Hemant Tyagi
ETH Zürich

Anastasios Kyrillidis
UT Austin, Texas

Bernd Gärtner
ETH Zürich

Andreas Krause
ETH Zürich

Abstract

A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is referred to as a Sparse Additive Model (SPAM), if it is of the form $f(\mathbf{x}) = \sum_{l \in \mathcal{S}} \phi_l(x_l)$, where $\mathcal{S} \subset [d]$, $|\mathcal{S}| \ll d$. Assuming ϕ_l 's and \mathcal{S} to be unknown, the problem of estimating f from its samples has been studied extensively. In this work, we consider a generalized SPAM, allowing for *second order* interaction terms. For some $\mathcal{S}_1 \subset [d]$, $\mathcal{S}_2 \subset \binom{[d]}{2}$, the function f is assumed to be of the form:

$$f(\mathbf{x}) = \sum_{p \in \mathcal{S}_1} \phi_p(x_p) + \sum_{(l,l') \in \mathcal{S}_2} \phi_{(l,l')}(x_l, x_{l'}).$$

Assuming $\phi_p, \phi_{(l,l')}$, \mathcal{S}_1 and \mathcal{S}_2 to be unknown, we provide a randomized algorithm that queries f and *exactly recovers* $\mathcal{S}_1, \mathcal{S}_2$. Consequently, this also enables us to estimate the underlying $\phi_p, \phi_{(l,l')}$. We derive sample complexity bounds for our scheme and also extend our analysis to include the situation where the queries are corrupted with noise – either stochastic, or arbitrary but bounded. Lastly, we provide simulation results on synthetic data, that validate our theoretical findings.

1 Introduction

Many scientific problems involve estimating an unknown function f , defined over a compact subset of \mathbb{R}^d , with d large. Such problems arise for instance, in modeling complex physical processes [1, 2, 3]. Information about f is typically available in the form of point values $(x_i, f(x_i))_{i=1}^n$, which are then used for learning f . It is well known that the problem suffers from the curse of dimensionality, if only smoothness assumptions are placed on f . For example, if f is C^s smooth, then for uniformly approximating f within error $\delta \in (0, 1)$, one needs $n = \Omega(\delta^{-d/s})$ samples [4].

A popular line of work in recent times considers the setting where f possesses an intrinsic low dimensional structure, *i.e.*, depends on only a small subset of d variables. There exist algorithms for estimating such f (tailored to the underlying structural assumption), along with attractive theoretical guarantees that do not suffer from the curse of dimensionality; see [5, 6, 7, 8]. One such assumption leads to the class of sparse additive models (SPAMs), wherein:

$$f(x_1, \dots, x_d) = \sum_{l \in \mathcal{S}} \phi_l(x_l),$$

for some unknown $\mathcal{S} \subset \{1, \dots, d\}$ with $|\mathcal{S}| = k \ll d$. There exist several algorithms for learning these models; we refer to [9, 10, 11, 12, 13] and references therein.

In this paper, we focus on a generalized SPAM model, where f can also contain a small number of *second order interaction terms*, *i.e.*,

$$f(x_1, \dots, x_d) = \sum_{p \in \mathcal{S}_1} \phi_p(x_p) + \sum_{(l,l') \in \mathcal{S}_2} \phi_{(l,l')}(x_l, x_{l'}); \tag{1.1}$$

$\mathcal{S}_1 \subset [d]$, $\mathcal{S}_2 \subset \binom{[d]}{2}$, with $|\mathcal{S}_1| \ll d$, $|\mathcal{S}_2| \ll d^2$. There exist relatively few results for learning models of the form (1.1), with the existing work being in the regression framework [14, 15, 16]. Here, $(x_i, f(x_i))_{i=1}^n$ are typically samples from an unknown probability measure \mathbb{P} .

We consider the setting where we have the freedom to query f at any desired set of points. We propose a strategy for querying f , along with an efficient recovery algorithm, which leads to much stronger guarantees, compared to those known in the regression setting. In particular, we provide the first *finite sample bounds* for exactly recovering sets \mathcal{S}_1 and \mathcal{S}_2 . Subsequently, we *uniformly* estimate the individual components: $\phi_p, \phi_{(l,l')}$ via additional queries of f along the subspaces corresponding to $\mathcal{S}_1, \mathcal{S}_2$.

Contributions. We make the following contributions for learning models of the form (1.1).

- (i) Firstly, we provide a randomized algorithm which provably recovers $\mathcal{S}_1, \mathcal{S}_2$ *exactly*, with $O(k\rho_m(\log d)^3)$ noiseless point queries. Here, ρ_m denotes the maximum number of occurrences of a

variable in \mathcal{S}_2 , and captures the underlying *complexity* of the interactions.

- (ii) An important tool in our analysis is a compressive sensing based sampling scheme, for recovering each row of a sparse Hessian matrix, for functions that also possess sparse gradients. This might be of independent interest.
- (iii) We theoretically analyze the impact of additive noise in the point queries on the performance of our algorithm, for two noise models: arbitrary bounded noise and independent, identically distributed (i.i.d.) noise. In particular, for additive Gaussian noise, we show that with $O(\rho_m^5 k^2 (\log d)^4)$ noisy point queries, our algorithm recovers $\mathcal{S}_1, \mathcal{S}_2$ exactly. We also provide simulation results on synthetic data that validate our theoretical findings.

Notation. For any vector $\mathbf{x} \in \mathbb{R}^d$, we denote its ℓ_p -norm by $\|\mathbf{x}\|_p := \left(\sum_{i=1}^d |x_i|^p\right)^{1/p}$. For a set \mathcal{S} , $(\mathbf{x})_{\mathcal{S}}$ denotes the restriction of \mathbf{x} onto \mathcal{S} , i.e., $((\mathbf{x})_{\mathcal{S}})_l = x_l$ if $l \in \mathcal{S}$ and 0 otherwise. For a function $g : \mathbb{R}^m \rightarrow \mathbb{R}$ of m variables, $\mathbb{E}_p[g]$, $\mathbb{E}_{(l,l')}[g]$, $\mathbb{E}[g]$ denote expectation with respect to uniform distributions over x_p , $(x_l, x_{l'})$ and (x_1, \dots, x_m) , respectively. For any compact $\Omega \subset \mathbb{R}^n$, $\|g\|_{L_\infty(\Omega)}$ denotes the L_∞ norm of g in Ω . The partial derivative operator $\partial/\partial x_i$ is denoted by ∂_i . For instance, $\partial_1^2 \partial_2 g$ denotes $\partial^3 g / \partial x_1^2 \partial x_2$.

2 Problem statement

We are interested in the problem of approximating functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ from point queries. For some unknown sets $\mathcal{S}_1 \subset [d]$, $\mathcal{S}_2 \subset \binom{[d]}{2}$, the function f is assumed to have the following form.

$$f(x_1, \dots, x_d) = \sum_{p \in \mathcal{S}_1} \phi_p(x_p) + \sum_{(l,l') \in \mathcal{S}_2} \phi_{(l,l')}(x_l, x_{l'}). \quad (2.1)$$

Here, $\phi_{(l,l')}$ is considered to be “truly bivariate” meaning that $\partial_l \partial_{l'} \phi_{(l,l')} \neq 0$. The set of all variables that occur in \mathcal{S}_2 , is denoted by $\mathcal{S}_2^{\text{var}}$. For each $l \in \mathcal{S}_2^{\text{var}}$, we refer to $\rho(l)$ as the *degree* of l , i.e., the number of occurrences of l in \mathcal{S}_2 , formally defined as:

$$\rho(l) := |\{l' \in \mathcal{S}_2^{\text{var}} : (l, l') \in \mathcal{S}_2 \text{ or } (l', l) \in \mathcal{S}_2\}|; \quad l \in \mathcal{S}_2^{\text{var}}.$$

The largest such degree is denoted by $\rho_m := \max_{l \in \mathcal{S}_2^{\text{var}}} \rho(l)$.

Our goal is to query f at suitably chosen points in its domain, in order to estimate it within the compact region¹ $[-1, 1]^d$. To this end, note that representation (2.1) is not

¹One could more generally consider the region $[\alpha, \beta]^d$ and transform the variables to $[-1, 1]^d$ via scaling and transformation.

unique². This is avoided by re-writing (2.1) in the following unique ANOVA form [17]:

$$f(x_1, \dots, x_d) = c + \sum_{p \in \mathcal{S}_1} \phi_p(x_p) + \sum_{(l,l') \in \mathcal{S}_2} \phi_{(l,l')}(x_l, x_{l'}) + \sum_{q \in \mathcal{S}_2^{\text{var}}: \rho(q) > 1} \phi_q(x_q), \quad (2.2)$$

where $\mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}} = \emptyset$. Here, $c = \mathbb{E}[f]$ and $\mathbb{E}_p[\phi_p] = \mathbb{E}_{(l,l')}[\phi_{(l,l')}] = 0$; $\forall p \in \mathcal{S}_1, (l, l') \in \mathcal{S}_2$, with expectations being over uniform distributions with respect to variable range $[-1, 1]$. In addition, $\mathbb{E}_l[\phi_{(l,l')}] = 0$ if $\rho(l) = 1$. The univariate ϕ_q corresponding to $q \in \mathcal{S}_2^{\text{var}}$ with $\rho(q) > 1$, represents the net marginal effect of the variable and has $\mathbb{E}_q[\phi_q] = 0$. We note that $\mathcal{S}_1, \mathcal{S}_2^{\text{var}}$ are disjoint in (2.2) as each $p \in \mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}}$ can be merged with their bivariate counterparts, uniquely. The uniqueness of (2.2) is shown formally in the appendix.

We assume the setting $|\mathcal{S}_1| = k_1 \ll d$, $|\mathcal{S}_2| = k_2 \ll d^2$. The set of *all* active variables i.e., $\mathcal{S}_1 \cup \mathcal{S}_2^{\text{var}}$ is denoted by \mathcal{S} , with $k := |\mathcal{S}| = k_1 + |\mathcal{S}_2^{\text{var}}|$ being the *total sparsity* of the problem.

Due to the special structure of f in (2.2), we note that if $\mathcal{S}_1, \mathcal{S}_2$ were known beforehand, then one can estimate f via standard results from approximation theory or from regression³. Hence, our primary focus in the paper is to recover $\mathcal{S}_1, \mathcal{S}_2$. Our main assumptions for this problem are listed below.

Assumption 1. f can be queried from the slight enlargement: $[-(1+r), (1+r)]^d$, for some small $r > 0$.

Assumption 2. Each $\phi_{(l,l')}, \phi_p$ is three times continuously differentiable, within $[-(1+r), (1+r)]^2$ and $[-(1+r), (1+r)]$ respectively. Since these domains are compact, there exist constants $B_m \geq 0$ ($m = 0, 1, 2, 3$) so that:

$$\|\partial_l^{m_1} \partial_{l'}^{m_2} \phi_{(l,l')}\|_{L_\infty[-(1+r), (1+r)]^2} \leq B_m; \quad m_1 + m_2 = m,$$

where $(l, l') \in \mathcal{S}_2$, and

$$\|\partial_p^m \phi_p\|_{L_\infty[-(1+r), (1+r)]} \leq B_m,$$

where $p \in \mathcal{S}_1$ or, $p \in \mathcal{S}_2^{\text{var}}$ and $\rho(p) > 1$.

Our next assumption is for identifying \mathcal{S}_1 .

Assumption 3. For some constants $D_1, \lambda_1 > 0$, we assume that for each $p \in \mathcal{S}_1$, \exists connected $\mathcal{I}_p \subset [-1, 1]$, of Lebesgue measure at least $\lambda_1 > 0$, such that $|\partial_p \phi_p(x_p)| > D_1, \forall x_p \in \mathcal{I}_p$. This assumption is in a sense necessary. If say $\partial_p \phi_p$ was zero throughout $[-1, 1]$, then it implies that $\phi_p \equiv 0$, since each ϕ_p has zero mean in (2.2).

²Firstly, we could add constants to each $\phi_l, \phi_{(l,l')}$, which sum up to zero. Furthermore, for each $l \in \mathcal{S}_2^{\text{var}} : \rho(l) > 1$, or $l \in \mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}} : \rho(l) = 1$, we could add univariates that sum to zero.

³This is discussed later.

Our last assumption concerns the identification of \mathcal{S}_2 .

Assumption 4. For some constants $D_2, \lambda_2 > 0$, we assume that for each $(l, l') \in \mathcal{S}_2$, \exists connected $\mathcal{I}_l, \mathcal{I}_{l'} \subset [-1, 1]$, each interval of Lebesgue measure at least $\lambda_2 > 0$, such that $|\partial_l \partial_{l'} \phi_{(l, l')}(x_l, x_{l'})| > D_2$, $\forall (x_l, x_{l'}) \in \mathcal{I}_l \times \mathcal{I}_{l'}$.

Given the above, our problem specific parameters are: (i) B_i ; $i = 0, \dots, 3$, (ii) D_j, λ_j ; $j = 1, 2$ and, (iii) k, ρ_m . We do not assume k_1, k_2 to be known, but instead assume that k is known. Furthermore it suffices to use estimates for the problem parameters instead of exact values: In particular, we can use upper bounds for: k, ρ_m, B_i ; $i = 0, \dots, 3$ and lower bounds for: D_j, λ_j ; $j = 1, 2$.

3 Our sampling scheme and algorithm

We start by explaining our sampling scheme, followed by our algorithm for identifying $\mathcal{S}_1, \mathcal{S}_2$. Our algorithm proceeds in two phases – we first estimate \mathcal{S}_2 and then \mathcal{S}_1 . Its theoretical properties for the *noiseless* query setting are described in Section 4. Section 5 then analyzes how the sampling conditions can be adapted to handle the *noisy* query setting.

3.1 Sampling scheme for estimating \mathcal{S}_2

Our main idea for estimating \mathcal{S}_2 is to estimate the off-diagonal entries of the Hessian of f , at appropriately chosen points. The motivation is the observation that for any $(l, l') \in \binom{[d]}{2}$:

$$\partial_l \partial_{l'} f = \begin{cases} \partial_l \partial_{l'} \phi_{(l, l')} & \text{if } (l, l') \in \mathcal{S}_2, \\ 0 & \text{otherwise.} \end{cases}$$

To this end, consider the Taylor expansion of the gradient ∇f , at $\mathbf{x} \in \mathbb{R}^d$, along the direction $\mathbf{v}' \in \mathbb{R}^d$, with step size μ_1 . Since f is C^3 smooth, we have for $\zeta_q = \mathbf{x} + \theta_q \mathbf{v}'$, for some $\theta_q \in (0, \mu_1)$, $q = 1, \dots, d$:

$$\frac{\nabla f(\mathbf{x} + \mu_1 \mathbf{v}') - \nabla f(\mathbf{x})}{\mu_1} = \nabla^2 f(\mathbf{x}) \mathbf{v}' + \frac{\mu_1}{2} \begin{pmatrix} \mathbf{v}'^T \nabla^2 \partial_1 f(\zeta_1) \mathbf{v}' \\ \vdots \\ \mathbf{v}'^T \nabla^2 \partial_d f(\zeta_d) \mathbf{v}' \end{pmatrix}. \quad (3.1)$$

We see from (3.1) that the l^{th} entry of $(\nabla f(\mathbf{x} + \mu_1 \mathbf{v}') - \nabla f(\mathbf{x}))/\mu_1$, corresponds to a “noisy” linear measurement of the l^{th} row of $\nabla^2 f(\mathbf{x})$ with \mathbf{v}' . The noise corresponds to the third order Taylor remainder terms of f .

Denoting the l^{th} row of $\nabla^2 f(\mathbf{x})$ by $\nabla \partial_l f(\mathbf{x}) \in \mathbb{R}^d$, we make the following crucial observation: if $l \in \mathcal{S}_2^{\text{arr}}$ then $\nabla \partial_l f(\mathbf{x})$ has at most ρ_m non-zero off-diagonal entries, implying that it is $(\rho_m + 1)$ sparse. This follows on account of the structure of f (2.2). Furthermore, if $l \in \mathcal{S}_1$ then $\nabla \partial_l f(\mathbf{x})$ has at most one non zero entry (namely the diagonal entry), while if $l \notin \mathcal{S}$, then $\nabla \partial_l f(\mathbf{x}) \equiv 0$.

Compressive sensing based estimation. Assuming for now that we have access to an oracle that provides us with gradient estimates of f , this suggests the following idea. We can obtain random linear measurements, for *each row* of $\nabla^2 f(\mathbf{x})$ via gradient differences, as in (3.1). As each row is sparse, it is known from compressive sensing (CS) [18, 19] that it can be recovered with only a few measurements.

Inspired by this observation, consider an oracle that provides us with the estimates: $\widehat{\nabla} f(\mathbf{x}), \{\widehat{\nabla} f(\mathbf{x} + \mu_1 \mathbf{v}'_j)\}_{j=1}^{m_{v'}}$ where \mathbf{v}'_j belong to the set:

$$\mathcal{V}' := \{\mathbf{v}'_j \in \mathbb{R}^d : v'_{j,q} = \pm 1/\sqrt{m_{v'}} \text{ w.p. } 1/2 \text{ each;} \\ j = 1, \dots, m_{v'} \text{ and } q = 1, \dots, d\}.$$

Let $\widehat{\nabla} f(\mathbf{x}) = \nabla f(\mathbf{x}) + \mathbf{w}(\mathbf{x})$, where $\mathbf{w}(\mathbf{x}) \in \mathbb{R}^d$ denotes the gradient estimation noise. Denoting $\mathbf{V}' = [\mathbf{v}'_1 \dots \mathbf{v}'_{m_{v'}}]^T$, we obtain d linear systems, by employing (3.1) at each $\mathbf{v}'_j \in \mathcal{V}'$:

$$\mathbf{y}_q = \mathbf{V}' \nabla \partial_q f(\mathbf{x}) + \eta_{\mathbf{q},1} + \eta_{\mathbf{q},2}; \quad q = 1, \dots, d. \quad (3.2)$$

$\mathbf{y}_q \in \mathbb{R}^{m_{v'}}$ represents the measurement vector for the q^{th} row, with

$$(\mathbf{y}_q)_j = ((\widehat{\nabla} f(\mathbf{x} + \mu_1 \mathbf{v}'_j) - \widehat{\nabla} f(\mathbf{x}))_q) / \mu_1$$

while $\eta_{\mathbf{q},1}, \eta_{\mathbf{q},2} \in \mathbb{R}^{m_{v'}}$ represent noise with $(\eta_{\mathbf{q},1})_j = (\mu_1/2) \mathbf{v}'_j^T \nabla^2 \partial_q f(\zeta_{q,j}) \mathbf{v}'_j$ and $(\eta_{\mathbf{q},2})_j = (w_q(\mathbf{x} + \mu_1 \mathbf{v}'_j) - w_q(\mathbf{x})) / \mu_1$. Given the measurement vector \mathbf{y}_q , we can then obtain the estimate $\widehat{\nabla} \partial_q f(\mathbf{x})$ individually for each $q = 1, \dots, d$, via ℓ_1 minimization [18, 19, 20].

Estimating sufficiently many Hessian’s. Having estimated *each row* of $\nabla^2 f$ at some fixed \mathbf{x} , we have at hand an estimate of the set: $\{\partial_i \partial_j f(\mathbf{x}) : (i, j) \in \binom{[d]}{2}\}$. Our next goal is to repeat the process, at sufficiently many \mathbf{x} ’s within $[-1, 1]^d$.

We will denote the set of such points as χ . This will then enable us to sample each underlying $\partial_l \partial_{l'} \phi_{(l, l')}$ within its respective critical interval, as defined in Assumption 4. Roughly speaking, since $|\partial_l \partial_{l'} \phi_{(l, l')}|$ is “suitably large” in such an interval, we will consequently be able to detect each $(l, l') \in \mathcal{S}_2$, via a thresholding procedure. To this end, we make use of a family of hash functions, defined as follows.

Definition 1. For some $t \in \mathbb{N}$ and $j = 1, 2, \dots, t$ let $h_j : [d] \rightarrow \{1, 2, \dots, t\}$. Then, the set $\mathcal{H}_t^d = \{h_1, h_2, \dots\}$ is a (d, t) -hash family if for any distinct $i_1, \dots, i_t \in [d]$, $\exists h \in \mathcal{H}_t^d$ such that h is an injection when restricted to i_1, i_2, \dots, i_t .

Hash functions are widely used in theoretical computer science, such as in finding juntas [21]. There exists a simple probabilistic method for constructing such a family, so that




Figure 1: (a) $\nabla^2 f(\mathbf{x})$ estimated using: $\widehat{\nabla} f(\mathbf{x})$ (at red disk) and neighborhood gradient estimates (at blue disks) (b) Geometric picture: $d = 3$, $h \in \mathcal{H}_2^3$ with $h(1) = h(3) \neq h(2)$. Red disks are points in $\chi(h)$.

for any constant $C > 1$, $|\mathcal{H}_t^d| \leq (C + 1)te^t \log d$ with high probability (w.h.p.)⁴ [5]. For our purposes, we consider the family \mathcal{H}_2^d so that for any distinct i, j , there exists $h \in \mathcal{H}_2^d$ such that $h(i) \neq h(j)$.

For any $h \in \mathcal{H}_2^d$, let us now denote the vectors $\mathbf{e}_1(h), \mathbf{e}_2(h) \in \mathbb{R}^d$ where

$$(\mathbf{e}_i(h))_q = \begin{cases} 1 & \text{if } h(q) = i, \\ 0 & \text{otherwise,} \end{cases}$$

for $i = 1, 2$ and $q = 1, \dots, d$. Given at hand \mathcal{H}_2^d , we construct our set χ using the procedure⁵ in [5]. Specifically, for some $m_x \in \mathbb{Z}^+$, we construct for each $h \in \mathcal{H}_2^d$ the set:

$$\chi(h) := \left\{ \mathbf{x}(h) \in [-1, 1]^d : \mathbf{x}(h) = \sum_{i=1}^2 c_i \mathbf{e}_i(h); \right. \\ \left. c_1, c_2 \in \left\{ -1, -\frac{m_x - 1}{m_x}, \dots, \frac{m_x - 1}{m_x}, 1 \right\} \right\}.$$

Then, we obtain $\chi = \cup_{h \in \mathcal{H}_2^d} \chi(h)$ as the set of points at which we will recover $\nabla^2 f$. Observe that χ has the property of discretizing *any* 2-dimensional canonical subspace, within $[-1, 1]^d$ with $|\chi| \leq (2m_x + 1)^2 |\mathcal{H}_2^d| = O(\log d)$.

Estimating sparse gradients. Note that ∇f is at most k sparse, due to the structure of f . We now describe the oracle that we use, for estimating sparse gradients. As f is \mathcal{C}^3 smooth, therefore the third order Taylor's expansion of f at \mathbf{x} , along $\mathbf{v}, -\mathbf{v} \in \mathbb{R}^d$, with step size $\mu > 0$, and $\zeta = \mathbf{x} + \theta \mathbf{v}$, $\zeta' = \mathbf{x} - \theta' \mathbf{v}$; $\theta, \theta' \in (0, \mu)$ leads to

$$\frac{f(\mathbf{x} + \mu \mathbf{v}) - f(\mathbf{x} - \mu \mathbf{v})}{2\mu} \\ = \langle \mathbf{v}, \nabla f(\mathbf{x}) \rangle + (R_3(\zeta) - R_3(\zeta')) / (2\mu). \quad (3.3)$$

(3.3) corresponds to a noisy-linear measurement of $\nabla f(\mathbf{x})$, with \mathbf{v} . The “noise” here arises on account of the third order terms $R_3(\zeta), R_3(\zeta') = O(\mu^3)$, in the Taylor expansion.

⁴With probability $1 - O(d^{-c})$ for some constant $c > 0$.

⁵Such sets were used in [5] for a more general problem involving functions that are intrinsically k variate.

Let \mathcal{V} denote the set of measurement vectors:

$$\mathcal{V} := \{v_j \in \mathbb{R}^d : v_{j,q} = \pm 1/\sqrt{m_v} \text{ w.p. } 1/2 \text{ each;} \\ j = 1, \dots, m_v \text{ and } q = 1, \dots, d\}.$$

Employing (3.3) at each $\mathbf{v}_j \in \mathcal{V}$, we obtain:

$$\mathbf{y} = \mathbf{V} \nabla f(\mathbf{x}) + \mathbf{n}. \quad (3.4)$$

Here, $\mathbf{y} \in \mathbb{R}^{m_v}$ denotes the measurement vector with $(\mathbf{y})_j = (f(\mathbf{x} + \mu \mathbf{v}_j) - f(\mathbf{x} - \mu \mathbf{v}_j)) / (2\mu)$. Also, $\mathbf{V} = [\mathbf{v}_1 \dots \mathbf{v}_{m_v}]^T \in \mathbb{R}^{m_v \times d}$ denotes the measurement matrix and $\mathbf{n} \in \mathbb{R}^{m_v}$ denotes the noise terms. We then estimate $\nabla f(\mathbf{x})$ via standard ℓ_1 minimization⁶ [18, 19, 20]. Estimating sparse gradients via CS, has been considered previously in [8, 13], albeit using *second order* Taylor expansions, for different function models.

3.2 Sampling scheme for estimating \mathcal{S}_1

Having obtained an estimate $\widehat{\mathcal{S}}_2$ of \mathcal{S}_2 we now proceed to estimate \mathcal{S}_1 . Let $\widehat{\mathcal{S}}_2^{\text{var}}$ denote the set of variables in $\widehat{\mathcal{S}}_2$ and $\mathcal{P} := [d] \setminus \widehat{\mathcal{S}}_2^{\text{var}}$. Assuming $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$, we are now left with a SPAM on the *reduced* variable set \mathcal{P} . Consequently, we employ the sampling scheme of [13], wherein the gradient of f is estimated at equispaced points, along a diagonal of $[-1, 1]^d$. For $m'_x \in \mathbb{Z}^+$, this set is defined as:

$$\chi_{\text{diag}} := \left\{ \mathbf{x} = (x \ x \ \dots \ x) \in \mathbb{R}^d : \right. \\ \left. x \in \left\{ -1, -\frac{m'_x - 1}{m'_x}, \dots, \frac{m'_x - 1}{m'_x}, 1 \right\} \right\}.$$

Note that $|\chi_{\text{diag}}| = 2m'_x + 1$. The motivation for estimating ∇f at $\mathbf{x} \in \chi_{\text{diag}}$ is that we obtain estimates of $\partial_p \phi_p$ at equispaced points within $[-1, 1]$, for $p \in \mathcal{S}_1$. With a sufficiently fine discretization, we would “hit” the critical regions associated with each $\partial_p \phi_p$, as defined in Assumption 3. By applying a thresholding operation, we would then be able to identify each $p \in \mathcal{S}_1$.

To this end, consider the set of sampling directions:

$$\mathcal{V}'' := \{\mathbf{v}''_j \in \mathbb{R}^d : v''_{j,q} = \pm 1/\sqrt{m_{v''}} \text{ w.p. } 1/2 \text{ each;} \\ j = 1, \dots, m_{v''} \text{ and } q = 1, \dots, d\},$$

and let $\mu' > 0$ denote the step size. For each $\mathbf{x} \in \chi_{\text{diag}}$, we will query f at points: $(\mathbf{x} + \mu' \mathbf{v}''_j)_{\mathcal{P}}, (\mathbf{x} - \mu' \mathbf{v}''_j)_{\mathcal{P}}; \mathbf{v}''_j \in \mathcal{V}''$, *restricted* to \mathcal{P} . Then, as described earlier, we can form a linear system consisting of $m_{v''}$ equations, and solve it via ℓ_1 minimization to obtain the gradient estimate. The complete procedure for estimating $\mathcal{S}_1, \mathcal{S}_2$, is described formally in Algorithm 1.

⁶Can be solved efficiently using interior point methods [22]

Algorithm 1 Algorithm for estimating $\mathcal{S}_1, \mathcal{S}_2$

1: **Input:** $m_v, m_{v'}, m_x, m'_x \in \mathbb{Z}^+$; $\mu, \mu_1, \mu' > 0$; $\tau' > 0, \tau'' > 0$.

2: **Initialization:** $\widehat{\mathcal{S}}_1, \widehat{\mathcal{S}}_2 = \emptyset$.

3: **Output:** Estimates $\widehat{\mathcal{S}}_2, \widehat{\mathcal{S}}_1$.

4:

5: Construct $(d, 2)$ -hash family \mathcal{H}_2^d and sets $\mathcal{V}, \mathcal{V}'$.

6: **for** $h \in \mathcal{H}_2^d$ **do**

7: Construct the set $\chi(h)$.

8: **for** $i = 1, \dots, (2m_x + 1)^2$ and $\mathbf{x}_i \in \chi(h)$ **do**

9: $(\mathbf{y}_i)_j = \frac{f(\mathbf{x}_i + \mu \mathbf{v}_j) - f(\mathbf{x}_i - \mu \mathbf{v}_j)}{2\mu}$; $j = 1, \dots, m_v$; $\mathbf{v}_j \in \mathcal{V}$.

10: $\widehat{\nabla} f(\mathbf{x}_i) := \operatorname{argmin}_{\mathbf{y}_i = \mathbf{V}\mathbf{z}} \|\mathbf{z}\|_1$.

11: **for** $p = 1, \dots, m_{v'}$ **do**

12: $(\mathbf{y}_{i,p})_j = \frac{f(\mathbf{x}_i + \mu_1 \mathbf{v}'_p + \mu \mathbf{v}_j) - f(\mathbf{x}_i + \mu_1 \mathbf{v}'_p - \mu \mathbf{v}_j)}{2\mu}$; $j = 1, \dots, m_v$; $\mathbf{v}'_p \in \mathcal{V}'$. ESTIMATION OF \mathcal{S}_2

13: $\widehat{\nabla} f(\mathbf{x}_i + \mu_1 \mathbf{v}'_p) := \operatorname{argmin}_{\mathbf{y}_{i,p} = \mathbf{V}\mathbf{z}} \|\mathbf{z}\|_1$.

14: **end for**

15: **for** $q = 1, \dots, d$ **do**

16: $(\mathbf{y}_q)_j = \frac{(\widehat{\nabla} f(\mathbf{x}_i + \mu_1 \mathbf{v}'_j) - \widehat{\nabla} f(\mathbf{x}_i))_q}{\mu_1}$; $j = 1, \dots, m_{v'}$.

17: $\widehat{\nabla} \partial_q f(\mathbf{x}_i) := \operatorname{argmin}_{\mathbf{y}_q = \mathbf{V}'\mathbf{z}} \|\mathbf{z}\|_1$.

18: $\widehat{\mathcal{S}}_2 = \widehat{\mathcal{S}}_2 \cup \left\{ (q, q') : q' \in \{q+1, \dots, d\} \ \& \ |(\widehat{\nabla} \partial_q f(\mathbf{x}_i))_{q'}| > \tau' \right\}$.

19: **end for**

20: **end for**

21: **end for**

22:

23: Construct the sets $\chi_{\text{diag}}, \mathcal{V}''$ and initialize $\mathcal{P} := [d] \setminus \widehat{\mathcal{S}}_2^{\text{var}}$.

24: **for** $i = 1, \dots, (2m'_x + 1)$ and $\mathbf{x}_i \in \chi_{\text{diag}}$ **do**

25: $(\mathbf{y}_i)_j = \frac{f((\mathbf{x}_i + \mu' \mathbf{v}'_j)_{\mathcal{P}}) - f((\mathbf{x}_i - \mu' \mathbf{v}'_j)_{\mathcal{P}})}{2\mu'}$; $j = 1, \dots, m_{v''}$; $\mathbf{v}_j \in \mathcal{V}''$.

26: $(\widehat{\nabla} f((\mathbf{x}_i)_{\mathcal{P}}))_{\mathcal{P}} := \operatorname{argmin}_{\mathbf{y}_i = (\mathbf{V}'')_{\mathcal{P}}(\mathbf{z})_{\mathcal{P}}} \|\mathbf{z}\|_1$. ESTIMATION OF \mathcal{S}_1

27: $\widehat{\mathcal{S}}_1 = \widehat{\mathcal{S}}_1 \cup \left\{ q \in \mathcal{P} : |(\widehat{\nabla} f((\mathbf{x}_i)_{\mathcal{P}}))_q| > \tau'' \right\}$.

28: **end for**

4 Theoretical guarantees for noiseless case

Next, we provide sufficient conditions on our sampling parameters that guarantee exact recovery of $\mathcal{S}_1, \mathcal{S}_2$, in the noiseless query setting. This is stated in the following Theorem. All proofs are deferred to the appendix.

Theorem 1. \exists positive constants $\{c'_i\}_{i=1}^3, \{C_i\}_{i=1}^3$ so that if: $m_x \geq \lambda_2^{-1}$, $m_v > c'_1 k \log(d/k)$, and $m_{v'} > c'_2 \rho_m \log(d/\rho_m)$, then the following holds. Denoting $a = \frac{(4\rho_m + 1)B_3}{2\sqrt{m_{v'}}$, $b = \frac{C_1 \sqrt{m_{v'}}((4\rho_m + 1)k)B_3}{3m_v}$, $a' = \frac{D_2}{4aC_2}$, let μ, μ_1 satisfy: $\mu^2 < (a'^2 a)/b$ and

$$\mu_1 \in (a' - \sqrt{a'^2 - (b\mu^2/a)}, a' + \sqrt{a'^2 - (b\mu^2/a)}).$$

We then have for $\tau' = C_2(a\mu_1 + \frac{b\mu^2}{\mu_1})$, that $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$ w.h.p. Provided $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$, if $m'_x \geq \lambda_1^{-1}$, $m_{v''} > c'_3(k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \log(\frac{|\mathcal{P}|}{k - |\widehat{\mathcal{S}}_2^{\text{var}}|})$ and $\mu'^2 < \frac{3m_{v''}D_1}{C_3(k - |\widehat{\mathcal{S}}_2^{\text{var}}|)B_3}$, then

$$\tau'' = \frac{C_3(k - |\widehat{\mathcal{S}}_2^{\text{var}}|)\mu'^2 B_3}{6m_{v''}}, \text{ implies } \widehat{\mathcal{S}}_1 = \mathcal{S}_1 \text{ w.h.p.}$$

Remark 1. We note that the condition on μ' is less strict than in [13] for identifying \mathcal{S}_1 . This is because in [13], the gradient is estimated via a forward difference procedure, while we perform a central difference procedure in (3.3).

Query complexity. Estimating $\nabla f(\mathbf{x})$ at some fixed \mathbf{x} requires $2m_v = O(k \log d)$ queries. Estimating $\nabla^2 f(\mathbf{x})$ involves computing an additional $m_{v'} = O(\rho_m \log d)$ gradient vectors in a neighborhood of \mathbf{x} – implying $O(m_v m_{v'}) = O(k \rho_m (\log d)^2)$ point queries. This consequently implies a total query complexity of $O(k \rho_m (\log d)^2 |\chi|) = O(\lambda_2^{-2} k \rho_m (\log d)^3)$, for estimating \mathcal{S}_2 . We make an additional $O(\lambda_1^{-1}(k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \log(d - |\widehat{\mathcal{S}}_2^{\text{var}}|))$ queries of f , in order to estimate \mathcal{S}_1 . Therefore, the overall query complexity for estimating $\mathcal{S}_1, \mathcal{S}_2$ is $O(\lambda_2^{-2} k \rho_m (\log d)^3)$.

\mathcal{H}_2^d can be constructed in $\text{poly}(d)$ time. For each $\mathbf{x} \in \chi$, we first solve $m_{v'} + 1$ linear programs (Steps 10, 13), each solvable in $\text{poly}(m_v, d)$ time. We then solve d linear programs (Step 17), with each taking $\text{poly}(m_{v'}, d)$ time. This is done at $|\chi| = O(\lambda_2^{-2} \log d)$ points, hence the overall *computation cost* for estimation of \mathcal{S}_2 (and later \mathcal{S}_1) is polynomial in: the number of queries, and d . Lastly, we note that [23] also estimates sparse Hessians via CS, albeit for the function optimization problem. Their scheme entails a sample complexity⁷ of $O(k\rho_m(\log(k\rho_m))^2(\log d)^2)$ for estimating $\nabla^2 f(\mathbf{x})$; this is worse by a $O((\log(k\rho_m))^2)$ term compared to our method.

Recovering the components of the model. Having estimated $\mathcal{S}_1, \mathcal{S}_2$, we can now estimate each underlying component in (2.2) by sampling f along the *subspace* corresponding to the component. Using these samples, one can then construct via standard techniques, a spline based quasi interpolant [24] that *uniformly* approximates the component. This is shown formally in the appendix.

5 Impact of noise

We now consider the case where the point queries are corrupted with external noise. This means that at query \mathbf{x} , we observe $f(\mathbf{x}) + z'$, where $z' \in \mathbb{R}$ denotes external noise.

In order to estimate $\nabla f(\mathbf{x})$, we obtain the samples: $f(\mathbf{x} + \mu\mathbf{v}_j) + z'_{j,1}$ and $f(\mathbf{x} - \mu\mathbf{v}_j) + z'_{j,2}$; $j = 1, \dots, m_v$. This changes (3.4) to the linear system $\mathbf{y} = \mathbf{V}\nabla f(\mathbf{x}) + \mathbf{n} + \mathbf{z}$, where $z_j = (z'_{j,1} - z'_{j,2})/(2\mu)$. Hence, the step-size μ needs to be chosen carefully now – a small value would blow up the external noise component, while a large value would increase perturbation due to the higher order Taylor’s terms.

Arbitrary bounded noise. In this scenario, we assume the external noise to be arbitrary and bounded, meaning that $|z'| < \varepsilon$, for some finite $\varepsilon \geq 0$. If ε is too large, then we would expect recovery of $\mathcal{S}_1, \mathcal{S}_2$ to be impossible as the structure of f would be destroyed.

We show in Theorem 2 that if $\varepsilon < \varepsilon_1 = O(D_2^3/(B_3^2\rho_m^2\sqrt{k}))$, then Algorithm 1 recovers \mathcal{S}_2 with appropriate choice of sampling parameters. Furthermore, assuming \mathcal{S}_2 is recovered exactly, and provided ε additionally satisfies $\varepsilon < \varepsilon_2 = O(D_1^{3/2}/\sqrt{(k-|\mathcal{S}_2^{\text{var}}|)B_3})$, then the algorithm also recovers \mathcal{S}_1 exactly. In contrast to Theorem 1, the step size μ cannot be chosen arbitrarily small now, due to external noise.

Theorem 2. *Let $m_x, m'_x, m_v, m_{v'}, m_{v''}$ be as defined in Theorem 1. Say $\varepsilon < \varepsilon_1 = O\left(\frac{D_2^3}{B_3^2\rho_m^2\sqrt{k}}\right)$. Denoting $b' = 2C_1\sqrt{m_v m_{v'}}$, $\exists 0 < A_1 < A_2$ and $0 < A_3 < A_4$ so that for $\mu \in (A_1, A_2)$, $\mu_1 \in (A_3, A_4)$ and $\tau' = C_2(a\mu_1 + \frac{b\mu^2}{\mu_1} + \frac{b'\varepsilon}{\mu\mu_1})$, we have $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$ w.h.p. Given $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$,*

⁷See [23, Corollary 4.1]

denote $a_1 = (k - |\widehat{\mathcal{S}}_2^{\text{var}}|)B_3/(6m_{v''})$, $b_1 = \sqrt{m_{v''}}$ and say $\varepsilon < \varepsilon_2 = O\left(\frac{D_1^{3/2}}{\sqrt{(k-|\mathcal{S}_2^{\text{var}}|)B_3}}\right)$. $\exists 0 < A_5 < A_6$ so that $\mu' \in (A_5, A_6)$, $\tau'' = C_3(a_1\mu'^2 + \frac{b_1\varepsilon}{\mu'})$ implies $\widehat{\mathcal{S}}_1 = \mathcal{S}_1$ w.h.p.

Stochastic noise. We now assume the point queries to be corrupted with i.i.d. Gaussian noise, so that $z' \sim \mathcal{N}(0, \sigma^2)$ with variance σ^2 . We consider resampling each point query a sufficient number of times, and averaging the values. During the \mathcal{S}_2 estimation phase, we resample each query N_1 times so that $z' \sim \mathcal{N}(0, \sigma^2/N_1)$. For any $0 < \varepsilon < \varepsilon_1$, if N_1 is suitably large, then we can uniformly bound $|z'| < \varepsilon$ – via standard tail bounds for Gaussians – over all noise samples, with high probability. Consequently, we can use the result of Theorem 2 for estimating \mathcal{S}_2 . The same reasoning applies to Step 25, i.e., the \mathcal{S}_1 estimation phase, where we resample each query N_2 times.

Theorem 3. *Let $m_x, m'_x, m_v, m_{v'}, m_{v''}$ be as defined in Theorem 1. For any $\varepsilon < \varepsilon_1$, $0 < p_1 < 1$, say we resample each query in Steps 9, 12, $N_1 > \frac{\sigma^2}{\varepsilon^2} \log\left(\frac{\sqrt{2}\sigma}{\varepsilon p_1} m_v(m_{v'} + 1)(2m_x + 1)^2|\mathcal{H}_2^d|\right)$ times, and take the average. For μ, μ_1, τ' as in Theorem 2, we have $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$ with probability at least $1 - p_1 - o(1)$. Given $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$, with $\varepsilon' < \varepsilon_2$, $0 < p_2 < 1$, say we resample each query in Step 25, $N_2 > \frac{\sigma^2}{\varepsilon'^2} \log\left(\frac{\sqrt{2}\sigma(2m'_x+1)m_{v''}}{\varepsilon' p_2}\right)$ times, and take the average. Then for μ', τ'' as in Theorem 2 (with ε replaced by ε'), we have $\widehat{\mathcal{S}}_1 = \mathcal{S}_1$ with probability at least $1 - p_2 - o(1)$.*

Query complexity. In the case of arbitrary, but bounded noise, the query complexity remains the same as for the noiseless case. In case of i.i.d. Gaussian noise, for estimating \mathcal{S}_2 , we have $\varepsilon = O(\rho_m^{-2}k^{-1/2})$. Choosing $p_1 = d^{-\delta}$ for any constant $\delta > 0$ gives us $N_1 = O(\rho_m^4 k \log d)$. This means that with $O(N_1 k \rho_m (\log d)^3 |\chi|) = O(\rho_m^5 k^2 (\log d)^4 \lambda_2^{-2})$ queries, $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$ holds w.h.p. Next, for estimating \mathcal{S}_1 , we have $\varepsilon' = O((k - |\mathcal{S}_2^{\text{var}}|)^{-1/2})$. Choosing $p_2 = ((d - |\mathcal{S}_2^{\text{var}}|)^{-\delta})$ for any constant $\delta > 0$, we get $N_2 = O((k - |\mathcal{S}_2^{\text{var}}|) \log(d - |\mathcal{S}_2^{\text{var}}|))$. This means the query complexity for estimating \mathcal{S}_1 is $O(N_2 \lambda_1^{-1} (k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \log(d - |\widehat{\mathcal{S}}_2^{\text{var}}|)) = O(\lambda_1^{-1} (k - |\widehat{\mathcal{S}}_2^{\text{var}}|)^2 (\log(d - |\widehat{\mathcal{S}}_2^{\text{var}}|))^2)$. Therefore, the overall query complexity of Algorithm 1 for estimating $\mathcal{S}_1, \mathcal{S}_2$ is $O(\rho_m^5 k^2 (\log d)^4 \lambda_2^{-2})$.

Remark 2. *We saw above that $O(k^2(\log d)^2)$ samples are sufficient for estimating \mathcal{S}_1 in presence of i.i.d Gaussian noise. This improves the corresponding bound in [13] by a $O(k)$ factor, and is due to the less strict condition on μ' .*

Recovering the components of the model. Having identified $\mathcal{S}_1, \mathcal{S}_2$, we can estimate the underlying components in (2.2), via standard nonparametric regression for ANOVA type models [25]. Alternately, for each component, we could also sample f along the subspace corresponding to

the component and then perform regression, to obtain its estimate with *uniform* error bounds. This is shown formally in the appendix.

6 Related work

Learning SPAMs. We begin with an overview of results for learning SPAMs, in the regression setting. [14] proposed the COSSO algorithm, that extends the Lasso to the reproducing kernel Hilbert space (RKHS) setting. [26] generalizes the non negative garrote to the nonparametric setting. [27, 9, 10] consider least squares methods, regularized by sparsity inducing penalty terms, for learning such models. [12, 28] propose a convex program for estimating f (in the RKHS setting) that achieves the minimax optimal error rates. [11] proposes a method based on the adaptive group Lasso. These methods are designed for learning SPAMs and cannot handle models of the form (1.1).

Learning generalized SPAMs. There exist fewer results for generalized SPAMs of the form (1.1), in the regression setting. The COSSO algorithm [14] can handle (1.1), however its convergence rates are shown only for the case of no interactions. [15] proposes the VANISH algorithm – a least squares method with sparsity constraints. It is shown to be sparsistent, *i.e.*, it asymptotically recovers $\mathcal{S}_1, \mathcal{S}_2$ for $n \rightarrow \infty$. They also show a consistency result for estimating f , similar to [9]. [16] proposes the ACOSSO method, an adaptive version of the COSSO algorithm, which can also handle (1.1). They derive convergence rates and sparsistency results for their method, albeit for the case of no interactions. [29] studies a generalization of (1.1) that allows for the presence of a sparse number of m -wise interaction terms for some additional sparsity parameter m . While they derive⁸ non-asymptotic L_2 error rates for estimating f , they do not guarantee unique identification of the interaction terms for any value of m . A special case of (1.1) – where ϕ_p 's are linear and each $\phi_{(l,\nu)}$ is of the form $x_l x_\nu$ – has been studied considerably. Within this setting, there exist algorithms that recover $\mathcal{S}_1, \mathcal{S}_2$, along with convergence rates for estimating f , but only in the limit of large n [30, 15, 31]. [32] generalized this to the setting of sparse multilinear systems – albeit in the noiseless setting – and derived non-asymptotic sampling bounds for identifying the interaction terms. However finite sample bounds for the non-linear model (1.1) are not known in general.

Learning generic low-dimensional function models. There exists related work in approximation theory – which is also the setting considered in this paper – wherein one assumes freedom to query f at any desired set of points within its domain. [5] considers functions depending on an unknown subset \mathcal{S} ($|\mathcal{S}| = k$) of the variables – a more

general model than (1.1). They provide a choice of query points of size $O(c^k k \log d)$ for some constant $c > 1$, and algorithms that recover \mathcal{S} w.h.p. [33] derives a simpler algorithm with sample complexity $O((C_1^4/\alpha^4)k(\log d)^2)$ for recovering \mathcal{S} w.h.p., where C_1, α depend⁹ on smoothness of f . For general k -variate f : $\alpha = c^{-k}$ for some constant $c > 1$, while for our model (1.1): $C_1 = O(\rho_m)$. This model was also studied in [34, 35] in the regression setting – they proposed an estimator that recovers \mathcal{S} w.h.p, with sample complexity $O(c^k k \log d)$. [8, 7] generalize this model to functions f of the form $f(\mathbf{x}) = g(\mathbf{A}\mathbf{x})$, for unknown $\mathbf{A} \in \mathbb{R}^{k \times d}$. They derive algorithms that approximately recover the row-span of \mathbf{A} w.h.p, with sample complexities typically polynomial in d .

While the above methods could possibly recover \mathcal{S} , they are not designed for identifying *interactions* among the variables. Specifically, their sample complexities exhibit a worse dependence on k, ρ_m and/or d . [13] provides a sampling scheme that specifically learns SPAMs, with sample complexities $O(k \log d), O(k^3(\log d)^2)$, in the absence/presence of Gaussian noise, respectively.

7 Simulation results

Dependence on d . We first consider the following experimental setup: $\mathcal{S}_1 = \{1, 2\}$ and $\mathcal{S}_2 = \{(3, 4), (4, 5)\}$, which implies $k_1 = 2, k_2 = 2, \rho_m = 2$ and $k = 5$. We consider two models:

- (i) $f_1(\mathbf{x}) = 2x_1 - 3x_2^2 + 4x_3x_4 - 5x_4x_5$,
- (ii) $f_2(\mathbf{x}) = 10 \sin(\pi \cdot x_1) + 5e^{-2x_2} + 10 \sin(\pi \cdot x_3x_4) + 5e^{-2x_4x_5}$.

We begin with the relatively simple model f_1 , for which the problem parameters are set to: $\lambda_1 = 0.3, \lambda_2 = 1, D_1 = 2, D_2 = 3, B_3 = 6$. We obtain $m_x = 1, m'_x = 4$. We use the same constant \tilde{C} when we set $m_v := \tilde{C}k \log(d/k), m_{v'} := \tilde{C}\rho_m \log(d/\rho_m)$, and $m_{v''} := \tilde{C}(k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \log(\frac{|P|}{k - |\widehat{\mathcal{S}}_2^{\text{var}}|})$. For the construction of the hash functions, we set the size to $|\mathcal{H}_2^d| = C' \log d$ with $C' = 1.7$, leading to $|\mathcal{H}_2^d| \in [8, 12]$ for $10^2 \leq d \leq 10^3$. We choose step sizes: μ, μ_1, μ' and thresholds: τ', τ'' as in Theorem 2. As CS solver, we use the ALPS algorithm [36], an efficient first-order method.

For the noisy setting, we consider the function values to be corrupted with i.i.d. Gaussian noise. The noise variance values considered are: $\sigma^2 \in \{10^{-4}, 10^{-3}, 10^{-2}\}$ for which we choose resampling factors: $(N_1, N_2) \in \{(50, 20), (85, 36), (90, 40)\}$. We see in Fig. 2, that for $\tilde{C} \approx 5.6$ the probability of successful identification (noiseless case) undergoes a phase transition and becomes close to 1, for different values of d . This validates Theorem 1. Fixing $\tilde{C} = 5.6$, we then see that with the total number of queries growing slowly with d , we have successful identification. For the noisy case, the total number of queries is

⁸In the Gaussian white noise model, which is known to be asymptotically equivalent to the regression model as $n \rightarrow \infty$.

⁹ $C_1 = \max_{i \in \mathcal{S}} \|\partial_i f\|_\infty$ and $\alpha = \min_{i \in \mathcal{S}} \|\partial_i f\|_1$

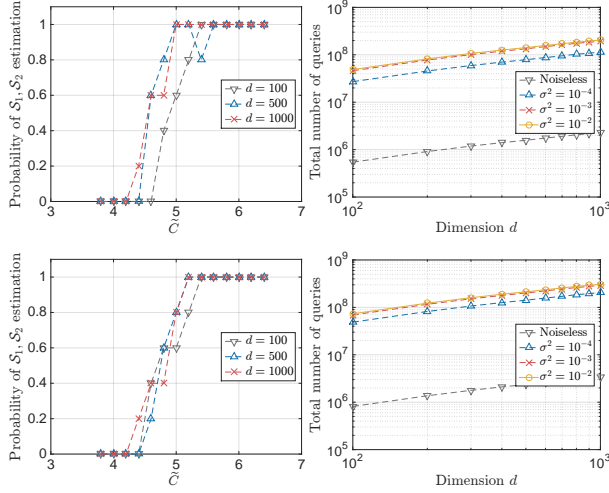


Figure 2: First (resp. second) row is for f_1 (resp. f_2). Left panel depicts the success probability of identifying exactly $\mathcal{S}_1, \mathcal{S}_2$, in the noiseless case. x -axis represent the constant \tilde{C} . The right panel depicts total queries vs. d for exact recovery, with $\tilde{C} = 5.6$ and various noise settings. All results are over 5 independent Monte Carlo trials.

roughly 10^2 times that in the noiseless setting, however the scaling with d is similar to the noiseless case.

We next consider the relatively harder model: f_2 , where the problem parameters are set to: $\lambda_1 = \lambda_2 = 0.3$, $D_1 = 8$, $D_2 = 4$, $B_3 = 35$ and, $m_x = m'_x = 4$. We see in Fig. 2, a phase transition (noiseless case) at $\tilde{C} = 5.6$ thus validating Theorem 1. For noisy cases, we consider σ^2 as before, and $(N_1, N_2) \in \{(60, 30), (90, 40), (95, 43)\}$. The number of queries is seen to be slightly larger than that for f_1 .

Dependence on k . We now demonstrate the scaling of the total number of queries versus the sparsity k for identification of $\mathcal{S}_1, \mathcal{S}_2$. Consider the model $f_3(\mathbf{x}) = \sum_{i=1}^T (\alpha_1 \mathbf{x}_{(i-1)5+1} - \alpha_2 \mathbf{x}_{(i-1)5+2} + \alpha_3 \mathbf{x}_{(i-1)5+3} \mathbf{x}_{(i-1)5+4} - \alpha_4 \mathbf{x}_{(i-1)5+4} \mathbf{x}_{(i-1)5+5})$ where $\mathbf{x} \in \mathbb{R}^d$ for $d = 500$. Here, $\alpha_i \in [2, 5], \forall i$; i.e., we randomly selected α_i 's within range and kept the values fixed for all 5 Monte Carlo iterations. Note that $\rho_m = 2$ and the sparsity $k = 5T$; we consider $T \in \{1, 2, \dots, 10\}$. We set $\lambda_1 = 0.3, \lambda_2 = 1, D_1 = 2, D_2 = 3, B_3 = 6$ and $\tilde{C} = 5.6$. For the noisy cases, we consider σ^2 as before, and choose the same values for (N_1, N_2) as for f_1 . In Figure 3 we see that the number of queries scales as $\sim k \log(d/k)$, and is roughly 10^2 more in the noisy case as compared to the noiseless setting.

Dependence on ρ_m . We now demonstrate the scaling of the total queries versus the maximum degree ρ_m for identification of $\mathcal{S}_1, \mathcal{S}_2$. Consider the model $f_4(\mathbf{x}) = \alpha_1 \mathbf{x}_1 - \alpha_2 \mathbf{x}_2^2 + \sum_{i=1}^T (\alpha_{3,i} \mathbf{x}_3 \mathbf{x}_{i+3}) + \sum_{i=1}^5 (\alpha_{4,i} \mathbf{x}_{2+2i} \mathbf{x}_{3+2i})$. We choose $d = 500, \tilde{C} = 6, \alpha_i \in [2, \dots, 5], \forall i$ (as ear-

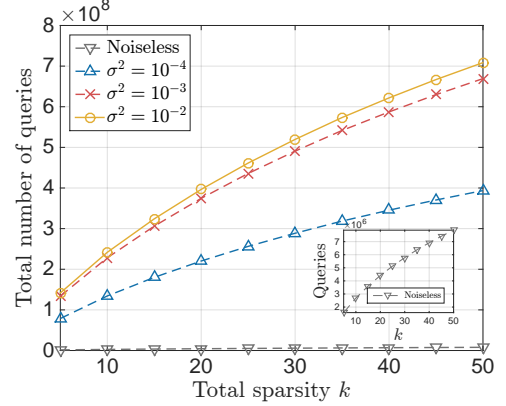


Figure 3: Total number of queries versus k for f_3 . This is shown for both noiseless and noisy cases (i.i.d Gaussian).

lier) and set $\lambda_1 = 0.3, \lambda_2 = 1, D_1 = 2, D_2 = 3, B_3 = 6$. For $T \geq 2$, we have $\rho_m = T$; we choose $T \in \{2, 3, \dots, 10\}$. Also note that $k = 13$ throughout. For the noisy cases, we consider σ^2 as before, and choose $(N_1, N_2) \in \{(70, 40), (90, 50), (100, 70)\}$. In Figure 4, we see that the number of queries scales as $\sim \rho_m \log(d/\rho_m)$, and is roughly 10^2 more in the noisy case as compared to the noiseless setting.

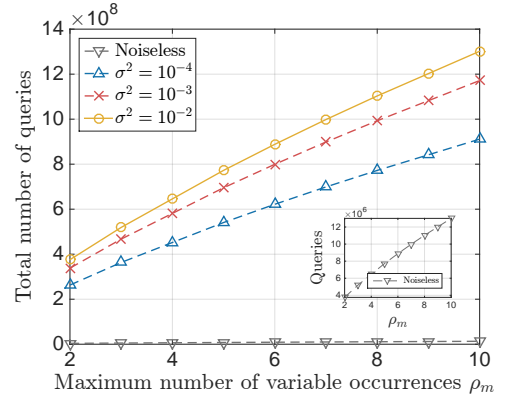


Figure 4: Total number of queries versus ρ_m for f_4 . This is shown for both noiseless and noisy cases (i.i.d Gaussian).

8 Concluding remarks

We proposed a sampling scheme for learning a generalized SPAM and provided finite sample bounds for recovering the underlying structure of such models. We also considered the setting where the point queries are corrupted with noise and analyzed sampling conditions for the same. It would be interesting to improve the sampling bounds that we obtained, and under similar assumptions. We leave this for future work.

Acknowledgements. This research was supported in part by SNSF grant CRSII2.147633.

References

- [1] Th. Muller-Gronbach and K. Ritter. Minimal errors for strong and weak approximation of stochastic differential equations. *Monte Carlo and Quasi-Monte Carlo Methods*, pages 53–82, 2008.
- [2] M.H. Maathuis, M. Kalisch, and P. Bühlmann. Estimating high-dimensional intervention effects from observational data. *Ann. Statist.*, 37(6A):3133–3164, 2009.
- [3] M.J. Wainwright. Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting. *IEEE Trans. Inform. Theory*, 55(12):5728–5741, 2009.
- [4] J.F. Traub, G.W. Wasilkowski, and H. Wozniakowski. *Information-Based Complexity*. Academic Press, New York, 1988.
- [5] R. DeVore, G. Petrova, and P. Wojtaszczyk. Approximation of functions of few variables in high dimensions. *Constr. Approx.*, 33:125–143, 2011.
- [6] A. Cohen, I. Daubechies, R.A. DeVore, G. Kerkyacharian, and D. Picard. Capturing ridge functions in high dimensions from point queries. *Constr. Approx.*, pages 1–19, 2011.
- [7] H. Tyagi and V. Cevher. Active learning of multi-index function models. In *Advances in Neural Information Processing Systems 25*, pages 1466–1474, 2012.
- [8] M. Fornasier, K. Schnass, and J. Vybíral. Learning functions of few arbitrary linear parameters in high dimensions. *Foundations of Computational Mathematics*, 12(2):229–262, 2012.
- [9] P. Ravikumar, J. Lafferty, H. Liu, and L. Wasserman. Sparse additive models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(5):1009–1030, 2009.
- [10] L. Meier, S. Van De Geer, and P. Bühlmann. High-dimensional additive modeling. *Ann. Statist.*, 37(6B):3779–3821, 2009.
- [11] J. Huang, J.L. Horowitz, and F. Wei. Variable selection in nonparametric additive models. *Ann. Statist.*, 38(4):2282–2313, 2010.
- [12] G. Raskutti, M.J. Wainwright, and B. Yu. Minimax-optimal rates for sparse additive models over kernel classes via convex programming. *J. Mach. Learn. Res.*, 13(1):389–427, 2012.
- [13] H. Tyagi, A. Krause, and B. Gärtner. Efficient sampling for learning sparse additive models in high dimensions. In *Advances in Neural Information Processing Systems 27*, pages 514–522, 2014.
- [14] Y. Lin and H.H. Zhang. Component selection and smoothing in multivariate nonparametric regression. *Ann. Statist.*, 34(5):2272–2297, 2006.
- [15] P. Radchenko and G. M. James. Variable selection using adaptive nonlinear interaction structures in high dimensions. *J. Amer. Statist. Assoc.*, 105:1541–1553, 2010.
- [16] C. B. Storlie, H. D. Bondell, B. J. Reich, and H. H. Zhang. Surface estimation, variable selection, and the nonparametric oracle property. *Statistica Sinica*, 21(2):679–705, 2011.
- [17] C. Gu. *Smoothing Spline ANOVA Models*. Springer (New York), 2002.
- [18] E.J. Candès, J.K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [19] D.L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.
- [20] P. Wojtaszczyk. ℓ_1 minimization with noisy data. *SIAM J. Numer. Anal.*, 50(2):458–467, 2012.
- [21] E. Mossel, R. O’Donnell, and R.P. Servedio. Learning juntas. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, pages 206–212, 2003.
- [22] Y. Nesterov and A. Nemirovskii. *Interior-Point Polynomial Algorithms in Convex Programming*. Society for Industrial and Applied Mathematics, Philadelphia, 1994.
- [23] A.S. Bandeira, K. Scheinberg, and L.N. Vicente. Computation of sparse low degree interpolating polynomials and their application to derivative-free optimization. *Mathematical Programming*, 134(1):223–257, 2012.
- [24] C. de Boor. *A practical guide to splines*. Springer Verlag (New York), 1978.
- [25] C.J. Stone. The use of polynomial splines and their tensor products in multivariate function estimation. *Ann. Statist.*, 22(1):118–171, 1994.
- [26] M. Yuan. Nonnegative garrote component selection in functional anova models. In *AISTATS*, volume 2, pages 660–666, 2007.

- [27] V. Koltchinskii and M. Yuan. Sparse recovery in large ensembles of kernel machines. In *COLT*, pages 229–238, 2008.
- [28] V. Koltchinskii and M. Yuan. Sparsity in multiple kernel learning. *Ann. Statist.*, 38(6):3660–3695, 2010.
- [29] A. Dalalyan, Y. Ingster, and A.B. Tsybakov. Statistical inference in compound functional models. *Probability Theory and Related Fields*, 158(3-4):513–532, 2014.
- [30] N.H. Choi, W. Li, and J. Zhu. Variable selection with the strong heredity constraint and its oracle property. *J. Amer. Statist. Assoc.*, 105(489):354–364, 2010.
- [31] J. Bien, J. Taylor, and R. Tibshirani. A lasso for hierarchical interactions. *Ann. Statist.*, 41(3):1111–1141, 2013.
- [32] B. Nazer and R.D. Nowak. Sparse interactions: Identifying high-dimensional multilinear systems via compressed sensing. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pages 1589–1596, 2010.
- [33] K. Schnass and J. Vybiral. Compressed learning of high-dimensional sparse functions. In *ICASSP*, 2011.
- [34] L. Comminges and A.S. Dalalyan. Tight conditions for consistent variable selection in high dimensional nonparametric regression. *J. Mach. Learn. Res.*, 19:187–206, 2012.
- [35] L. Comminges and A.S. Dalalyan. Tight conditions for consistency of variable selection in the context of high dimensionality. *Ann. Statist.*, 40(5):2667–2696, 2012.
- [36] A. Kyrillidis and V. Cevher. Recipes on hard thresholding methods. In *CAMSAP*, 2011.
- [37] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [38] L. Györfi, M. Kohler, A. Krzyzak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer, New York, 2002.
- [39] A.B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2008.
- [40] J. Fan and I. Gijbels. *Local polynomial modeling and its applications*. Chapman & Hall, London, New York, 1996.
- [41] A. Nemirovski. *Topics in non-parametric statistics*. In *Ecole d’Eté de Probabilités de Saint-Flour XVIII*, 1998, 85-277, Springer, New York, 2000.

Supplementary Material : Learning Sparse Additive Models with Interactions in High Dimensions.

In this supplementary material, we prove the results stated in the paper. In Section A, we show that the model representation (2.2) is a unique representation for f of the form (2.1). In Section B, we prove the main results of this paper namely: Theorem 1, Theorem 2 and Theorem 3. In Section C we discuss how the individual components of the model (2.2) can be estimated once $\mathcal{S}_1, \mathcal{S}_2$ are known. This is shown for both the noiseless as well as the noisy setting.

A Model uniqueness

We show here that the model representation (2.2) is a unique representation for f of the form (2.1). We first note that any measurable $f : \mathbb{R}^d \rightarrow \mathbb{R}$, admits a unique ANOVA decomposition (cf., [17]) of the form:

$$f(x_1, \dots, x_d) = c + \sum_{\alpha} f_{\alpha}(x_{\alpha}) + \sum_{\alpha < \beta} f_{\alpha\beta} + \sum_{\alpha < \beta < \gamma} f_{\alpha\beta\gamma} + \dots \quad (\text{A.1})$$

Indeed, for any probability measure μ_{α} on \mathbb{R} ; $\alpha = 1, \dots, d$, let \mathcal{E}_{α} denote the averaging operator, defined as

$$\mathcal{E}_{\alpha}(f)(\mathbf{x}) := \int_{\mathbb{R}} f(x_1, \dots, x_d) d\mu_{\alpha}. \quad (\text{A.2})$$

Then the components of the model can be written as : $c = (\prod_{\alpha} \mathcal{E}_{\alpha})f$, $f_{\alpha} = ((I - \mathcal{E}_{\alpha}) \prod_{\beta \neq \alpha} \mathcal{E}_{\beta})f$, $f_{\alpha\beta} = ((I - \mathcal{E}_{\alpha})(I - \mathcal{E}_{\beta}) \prod_{\gamma \neq \alpha, \beta} \mathcal{E}_{\gamma})f$, and so on. For our purpose, μ_{α} is considered to be the uniform probability measure on $[-1, 1]$. This is because we are interested in estimating f within $[-1, 1]^d$. Given this, we now find the ANOVA decomposition of f defined in (2.1).

As a sanity check, let us verify that $f_{\alpha\beta\gamma} \equiv 0$ for all $\alpha < \beta < \gamma$. Indeed if $p \in \mathcal{S}_1$, then at least two of $\alpha < \beta < \gamma$ will not be equal to p . Similarly for any $(l, l') \in \mathcal{S}_2$, at least one of α, β, γ will not be equal to l and l' . This implies $f_{\alpha\beta\gamma} \equiv 0$. The same reasoning easily applies for high order components of the ANOVA decomposition.

That $c = \mathbb{E}[f] = \sum_{p \in \mathcal{S}_1} \mathbb{E}_p[\phi_p] + \sum_{(l, l') \in \mathcal{S}_2} \mathbb{E}_{(l, l')}[\phi_{(l, l')}]$ is readily seen. Next, we have that

$$(I - \mathcal{E}_{\alpha}) \prod_{\beta \neq \alpha} \mathcal{E}_{\beta} \phi_p = \begin{cases} 0 & ; \alpha \neq p, \\ \phi_p - \mathbb{E}_p[\phi_p] & ; \alpha = p \end{cases}; \quad p \in \mathcal{S}_1. \quad (\text{A.3})$$

$$(I - \mathcal{E}_{\alpha}) \prod_{\beta \neq \alpha} \mathcal{E}_{\beta} \phi_{(l, l')} = \begin{cases} \mathbb{E}_{l'}[\phi_{(l, l')}] - \mathbb{E}_{(l, l')}[\phi_{(l, l')}] & ; \alpha = l, \\ \mathbb{E}_l[\phi_{(l, l')}] - \mathbb{E}_{(l, l')}[\phi_{(l, l')}] & ; \alpha = l', \\ 0 & ; \alpha \neq l, l', \end{cases}; \quad (l, l') \in \mathcal{S}_2. \quad (\text{A.4})$$

(A.3), (A.4) give us the first order components of $\phi_p, \phi_{(l, l')}$ respectively. One can next verify, using the same arguments as earlier, that for any $\alpha < \beta$:

$$(I - \mathcal{E}_{\alpha})(I - \mathcal{E}_{\beta}) \prod_{\gamma \neq \alpha, \beta} \mathcal{E}_{\gamma} \phi_p = 0; \quad \forall p \in \mathcal{S}_1. \quad (\text{A.5})$$

Lastly, we have for any $\alpha < \beta$ that the corresponding second order component of $\phi_{(l, l')}$ is given by:

$$(I - \mathcal{E}_{\alpha})(I - \mathcal{E}_{\beta}) \prod_{\gamma \neq \alpha, \beta} \mathcal{E}_{\gamma} \phi_{(l, l')} = \begin{cases} \phi_{(l, l')} - \mathbb{E}_l[\phi_{(l, l')}] & \\ -\mathbb{E}_{l'}[\phi_{(l, l')}] + \mathbb{E}_{(l, l')}[\phi_{(l, l')}] & ; \alpha = l, \beta = l', \\ 0 & ; \text{otherwise} \end{cases}; \quad (l, l') \in \mathcal{S}_2. \quad (\text{A.6})$$

We now make the following observations regarding the variables in $\mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}}$.

1. For each $l \in \mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}}$ such that: $\rho(l) = 1$, and $(l, l') \in \mathcal{S}_2$, we can simply merge ϕ_l with $\phi_{(l, l')}$. Thus l is no longer in \mathcal{S}_1 .
2. For each $l \in \mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}}$ such that: $\rho(l) > 1$, we can add the first order component for ϕ_l with the total first order component corresponding to all $\phi_{(l, l')}$'s and $\phi_{(l', l)}$'s. Hence again, l will no longer be in \mathcal{S}_1 .

Therefore all $q \in \mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}}$ can essentially be merged with \mathcal{S}_2 . Keeping this re-arrangement in mind, we can to begin with, assume in (2.1) that $\mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}} = \emptyset$. Then with the help of (A.3), (A.4), (A.5), (A.6), we have that any f of the form (2.1) (with $\mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}} = \emptyset$), can be uniquely written as:

$$f(x_1, \dots, x_d) = c + \sum_{p \in \mathcal{S}_1} \tilde{\phi}_p(x_p) + \sum_{(l, l') \in \mathcal{S}_2} \tilde{\phi}_{(l, l')}(x_l, x_{l'}) + \sum_{q \in \mathcal{S}_2^{\text{var}}: \rho(q) > 1} \tilde{\phi}_q(x_q); \quad \mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}} = \emptyset, \quad (\text{A.7})$$

where

$$c = \sum_{p \in \mathcal{S}_1} \mathbb{E}_p[\phi_p] + \sum_{(l, l') \in \mathcal{S}_2} \mathbb{E}_{(l, l')}[\phi_{(l, l')}], \quad (\text{A.8})$$

$$\tilde{\phi}_p = \phi_p - \mathbb{E}_p[\phi_p]; \quad \forall p \in \mathcal{S}_1, \quad (\text{A.9})$$

$$\tilde{\phi}_{(l, l')} = \begin{cases} \phi_{(l, l')} - \mathbb{E}_{(l, l')}[\phi_{(l, l')}] & \rho(l), \rho(l') = 1, \\ \phi_{(l, l')} - \mathbb{E}_l[\phi_{(l, l')}] & \rho(l) = 1, \rho(l') > 1, \\ \phi_{(l, l')} - \mathbb{E}_{l'}[\phi_{(l, l')}] & \rho(l) > 1, \rho(l') = 1, \\ \phi_{(l, l')} - \mathbb{E}_l[\phi_{(l, l')}] - \mathbb{E}_{l'}[\phi_{(l, l')}] + \mathbb{E}_{(l, l')}[\phi_{(l, l')}] & \rho(l) > 1, \rho(l') > 1, \end{cases} \quad (\text{A.10})$$

$$\begin{aligned} \text{and } \tilde{\phi}_q &= \sum_{q': (q, q') \in \mathcal{S}_2} (\mathbb{E}_{q'}[\phi_{(q, q')}] - \mathbb{E}_{(q, q')}[\phi_{(q, q')}] \\ &+ \sum_{q': (q', q) \in \mathcal{S}_2} (\mathbb{E}_{q'}[\phi_{(q', q)}] - \mathbb{E}_{(q', q)}[\phi_{(q', q)}]); \quad \forall q \in \mathcal{S}_2^{\text{var}}: \rho(q) > 1. \end{aligned} \quad (\text{A.11})$$

B Proofs

B.1 Proof of Theorem 1

The proof makes use of the following key theorem from [8], for stable approximation via ℓ_1 minimization: $\Delta(\mathbf{y}) = \underset{\mathbf{z} = \mathbf{V}\mathbf{z}}{\text{argmin}} \|\mathbf{z}\|_1$. While the first part is standard (see for example [37]), the second result was stated in [8] as a specialization of Theorem 1.2 from [20] to the case of Bernoulli measurement matrices.

Theorem 4 ([20, 8]). *Let \mathbf{V} be a $m_v \times d$ random matrix with all entries being Bernoulli i.i.d random variables scaled with $1/\sqrt{m_v}$. Then the following results hold.*

1. Let $0 < \kappa < 1$. Then there are two positive constants $c_1, c_2 > 0$, such that the matrix \mathbf{V} has the Restricted Isometry Property

$$(1 - \kappa) \|\mathbf{w}\|_2^2 \leq \|\mathbf{V}\mathbf{w}\|_2^2 \leq (1 + \kappa) \|\mathbf{w}\|_2^2 \quad (\text{B.1})$$

for all $\mathbf{w} \in \mathbb{R}^d$ such that $\#\text{supp}(\mathbf{w}) \leq c_2 m_v / \log(d/m_v)$ with probability at least $1 - e^{-c_1 m_v}$.

2. Let us suppose $d > (\log 6)^2 m_v$. Then there are positive constants $C, c'_1, c'_2 > 0$ such that with probability at least $1 - e^{-c'_1 m_v} - e^{-\sqrt{m_v d}}$ the matrix \mathbf{V} has the following property. For every $\mathbf{w} \in \mathbb{R}^d$, $\mathbf{n} \in \mathbb{R}^{m_v}$ and every natural number $k \leq c'_2 m_v / \log(d/m_v)$, we have

$$\|\Delta(\mathbf{V}\mathbf{w} + \mathbf{n}) - \mathbf{w}\|_2 \leq C \left(k^{-1/2} \sigma_k(\mathbf{w})_1 + \max \left\{ \|\mathbf{n}\|_2, \sqrt{\log d} \|\mathbf{n}\|_\infty \right\} \right), \quad (\text{B.2})$$

where

$$\sigma_k(\mathbf{w})_1 := \inf \{ \|\mathbf{w} - \mathbf{z}\|_1 : \#\text{supp}(\mathbf{z}) \leq k \}$$

is the best k -term approximation of \mathbf{w} .

Remark 1. *The proof of the second part of Theorem 4 requires (B.1) to hold, which is the case in our setting with high probability.*

Remark 2. *Since $m_v \geq K$ is necessary, note that $K \leq c'_2 m_v / \log(d/m_v)$ is satisfied if $m_v > (1/c'_2) K \log(d/K)$. Also note that $K \log(d/K) > \log d$ in the regime¹⁰ $K \ll d$.*

We can now prove Theorem 1. The proof is divided into the following steps.

¹⁰More precisely, if $d > K^{\frac{K}{K-1}}$.

Bounding the $\eta_{q,2}$ term. Since $\nabla f(\mathbf{x})$ is at most k sparse, therefore for any $\mathbf{x} \in \mathbb{R}^d$ we immediately have from Theorem 4, (B.2), the following. $\exists C_1, c'_4 > 0, c'_1 \geq 1$ such that for $c'_1 k \log(\frac{d}{k}) < m_v < \frac{d}{(\log 6)^2}$ we have with probability at least $1 - e^{-c'_4 m_v} - e^{-\sqrt{m_v d}}$ that

$$\|\widehat{\nabla} f(\mathbf{x}) - \nabla f(\mathbf{x})\|_2 \leq C_1 \max \left\{ \|\mathbf{n}\|_2, \sqrt{\log d} \|\mathbf{n}\|_\infty \right\}. \quad (\text{B.3})$$

Recall that $\mathbf{n} = [n_1 \dots n_{m_v}]$ where $n_j = \frac{R_3(\zeta_j) - R_3(\zeta'_j)}{2\mu}$, for some $\zeta_j, \zeta'_j \in \mathbb{R}^d$. Here $R_3(\zeta)$ denotes the third order Taylor remainder terms of f . By taking the structure of f into account, we can uniformly bound $|R_3(\zeta_j)|$ as follows (so the same bound holds for $|R_3(\zeta'_j)|$). Let us define $\alpha := |\{q \in \mathcal{S}_2^{\text{var}} : \rho(q) > 1\}|$, to be the number of variables in $\mathcal{S}_2^{\text{var}}$, with degree greater than one.

$$\begin{aligned} |R_3(\zeta_j)| &= \frac{\mu^3}{6} \left| \sum_{p \in \mathcal{S}_1} \partial_p^3 \phi_p(\zeta_{j,p}) v_p^3 + \sum_{(l,l') \in \mathcal{S}_2} (\partial_l^3 \phi_{(l,l')}(\zeta_{j,l}, \zeta_{j,l'}) v_l^3 + \partial_{l'}^3 \phi_{(l,l')}(\zeta_{j,l}, \zeta_{j,l'}) v_{l'}^3 \right. \\ &\quad + \sum_{(l,l') \in \mathcal{S}_2} (3\partial_l \partial_{l'}^2 \phi_{(l,l')}(\zeta_{j,l}, \zeta_{j,l'}) v_l v_{l'}^2 + 3\partial_{l'}^2 \partial_l \phi_{(l,l')}(\zeta_{j,l}, \zeta_{j,l'}) v_l^2 v_{l'}) \\ &\quad \left. + \sum_{q \in \mathcal{S}_2^{\text{var}}: \rho(q) > 1} \partial_q^3 \phi_q(\zeta_{j,q}) v_q^3 \right| \end{aligned} \quad (\text{B.4})$$

$$\leq \frac{\mu^3}{6} \left(\frac{k_1 B_3}{m_v^{3/2}} + \frac{2k_2 B_3}{m_v^{3/2}} + \frac{\alpha B_3}{m_v^{3/2}} + \frac{6k_2 B_3}{m_v^{3/2}} \right) \quad (\text{B.5})$$

$$= \frac{\mu^3}{6} \frac{(k_1 + \alpha + 8k_2) B_3}{m_v^{3/2}}. \quad (\text{B.6})$$

Using the fact $2k_2 = \sum_{l \in \mathcal{S}_2^{\text{var}}: \rho(l) > 1} \rho(l) + (|\mathcal{S}_2^{\text{var}}| - \alpha)$, we can observe that $2k_2 \leq \rho_m \alpha + (|\mathcal{S}_2^{\text{var}}| - \alpha) = |\mathcal{S}_2^{\text{var}}| + (\rho_m - 1)\alpha$. Plugging this in (B.6), and using the fact $\alpha \leq k$ (since we do not assume α to be known), we obtain

$$|R_3(\zeta_j)| \leq \frac{\mu^3}{6} \frac{(k_1 + \alpha + 4|\mathcal{S}_2^{\text{var}}| + 4(\rho_m - 1)\alpha) B_3}{m_v^{3/2}} \quad (\text{B.7})$$

$$\leq \frac{\mu^3 (4k + (4\rho_m - 3)\alpha) B_3}{6m_v^{3/2}} \leq \frac{\mu^3 ((4\rho_m + 1)k) B_3}{6m_v^{3/2}}. \quad (\text{B.8})$$

This in turn implies that $\|\mathbf{n}\|_\infty \leq \frac{\mu^2 ((4\rho_m + 1)k) B_3}{6m_v^{3/2}}$. Using the fact $\|\mathbf{n}\|_2 \leq \sqrt{m_v} \|\mathbf{n}\|_\infty$, we thus obtain for the stated choice of m_v (cf. Remark 2) that

$$\|\widehat{\nabla} f(\mathbf{x}) - \nabla f(\mathbf{x})\|_2 \leq \frac{C_1 \mu^2 ((4\rho_m + 1)k) B_3}{6m_v}, \quad \forall \mathbf{x} \in [-(1+r), 1+r]^d. \quad (\text{B.9})$$

Recall that $[-(1+r), 1+r]^d, r > 0$, denotes the enlargement around $[-1, 1]^d$, in which the smoothness properties of $\phi_p, \phi_{(l,l')}$ are defined in Section 2. Since $\widehat{\nabla} f(\mathbf{x}) = \nabla f(\mathbf{x}) + \mathbf{w}(\mathbf{x})$, therefore $\|\mathbf{w}(\mathbf{x})\|_\infty \leq \|\widehat{\nabla} f(\mathbf{x}) - \nabla f(\mathbf{x})\|_2$. Using the definition of $\eta_{q,2} \in \mathbb{R}^{m_v}$ from (3.2), we then have that $\|\eta_{q,2}\|_\infty \leq \frac{C_1 \mu^2 ((4\rho_m + 1)k) B_3}{3m_v \mu_1}$.

Bounding the $\eta_{q,1}$ term. We will bound $\|\eta_{q,1}\|_\infty$. To this end, we see from (3.2) that it suffices to uniformly bound $|\mathbf{v}'^T \nabla^2 \partial_q f(\zeta) \mathbf{v}'|$, over all: $q \in \mathcal{S}_1 \cup \mathcal{S}_2^{\text{var}}, \mathbf{v}' \in \mathcal{V}', \zeta \in [-(1+r), (1+r)]^d$. Note that

$$\mathbf{v}'^T \nabla^2 \partial_q f(\zeta) \mathbf{v}' = \sum_{l=1}^d v_l'^2 (\nabla^2 \partial_q f(\zeta))_{l,l} + \sum_{i \neq j=1}^d v_i' v_j' (\nabla^2 \partial_q f(\zeta))_{i,j}. \quad (\text{B.10})$$

We have the following three cases, depending on the type of q .

1. $q \in \mathcal{S}_1$.

$$\mathbf{v}'^T \nabla^2 \partial_q f(\zeta) \mathbf{v}' = v_q'^2 \partial_q^3 \phi_q(\zeta_q) \Rightarrow |\mathbf{v}'^T \nabla^2 \partial_q f(\zeta) \mathbf{v}'| \leq \frac{B_3}{m_v'}. \quad (\text{B.11})$$

2. $(\mathbf{q}, \mathbf{q}') \in \mathcal{S}_2$, $\rho(\mathbf{q}) = \mathbf{1}$.

$$\mathbf{v}'^T \nabla^2 \partial_q f(\zeta) \mathbf{v}' = v_q'^2 \partial_q^3 \phi_{(q, q')}(\zeta_q, \zeta_{q'}) + v_{q'}'^2 \partial_{q'}^2 \partial_q \phi_{(q, q')}(\zeta_q, \zeta_{q'}) \quad (\text{B.12})$$

$$+ 2v_q' v_{q'}' \partial_{q'} \partial_q^2 \phi_{(q, q')}(\zeta_q, \zeta_{q'}), \quad (\text{B.13})$$

$$\Rightarrow |\mathbf{v}'^T \nabla^2 \partial_q f(\zeta) \mathbf{v}'| \leq \frac{4B_3}{m_{v'}}. \quad (\text{B.14})$$

3. $\mathbf{q} \in \mathcal{S}_2^{\text{var}}$, $\rho(\mathbf{q}) > \mathbf{1}$.

$$\begin{aligned} \mathbf{v}'^T \nabla^2 \partial_q f(\zeta) \mathbf{v}' &= v_q'^2 (\partial_q^3 \phi_q(\zeta_q)) + \sum_{(q, q') \in \mathcal{S}_2} \partial_q^3 \phi_{(q, q')}(\zeta_q, \zeta_{q'}) \\ &+ \sum_{(q', q) \in \mathcal{S}_2} \partial_{q'}^3 \phi_{(q', q)}(\zeta_{q'}, \zeta_q) + \sum_{(q, q') \in \mathcal{S}_2} v_{q'}'^2 \partial_{q'}^2 \partial_q \phi_{(q, q')}(\zeta_q, \zeta_{q'}) \\ &+ \sum_{(q', q) \in \mathcal{S}_2} v_{q'}'^2 \partial_{q'}^2 \partial_q \phi_{(q', q)}(\zeta_{q'}, \zeta_q) + 2 \sum_{(q, q') \in \mathcal{S}_2} v_q' v_{q'}' \partial_{q'} \partial_q^2 \phi_{(q, q')}(\zeta_q, \zeta_{q'}) \\ &+ 2 \sum_{(q', q) \in \mathcal{S}_2} v_q' v_{q'}' \partial_{q'} \partial_q^2 \phi_{(q', q)}(\zeta_{q'}, \zeta_q), \end{aligned} \quad (\text{B.15})$$

$$\Rightarrow |\mathbf{v}'^T \nabla^2 \partial_q f(\zeta) \mathbf{v}'| \leq \frac{1}{m_{v'}} ((\rho_m + 1)B_3 + \rho_m B_3 + 2\rho_m B_3) = \frac{(4\rho_m + 1)B_3}{m_{v'}}. \quad (\text{B.16})$$

We can now uniformly bound $\|\eta_{\mathbf{q}, \mathbf{1}}\|_\infty$ as follows.

$$\|\eta_{\mathbf{q}, \mathbf{1}}\|_\infty := \max_{j=1, \dots, m_{v'}} \frac{\mu_1}{2} |\mathbf{v}'^T \nabla^2 \partial_q f(\zeta_j) \mathbf{v}'| \leq \frac{\mu_1 (4\rho_m + 1)B_3}{2m_{v'}}. \quad (\text{B.17})$$

Estimating \mathcal{S}_2 . We now proceed towards estimating \mathcal{S}_2 . To this end, we estimate $\nabla \partial_q f(\mathbf{x})$ for each $q = 1, \dots, d$ and $\mathbf{x} \in \chi$. Since $\nabla \partial_q f(\mathbf{x})$ is at most $(\rho_m + 1)$ -sparse, therefore Theorem 4, (B.2), immediately yield the following. $\exists C_2, c'_5 > 0, c'_2 \geq 1$ such that for $c'_2 \rho_m \log(\frac{d}{\rho_m}) < m_{v'} < \frac{d}{(\log 6)^2}$ we have with probability at least $1 - e^{-c'_5 m_{v'}} - e^{-\sqrt{m_{v'} d}}$ that

$$\|\widehat{\nabla} \partial_q f(\mathbf{x}) - \nabla \partial_q f(\mathbf{x})\|_2 \leq C_2 \max \left\{ \|\eta_{\mathbf{q}, \mathbf{1}} + \eta_{\mathbf{q}, \mathbf{2}}\|_2, \sqrt{\log d} \|\eta_{\mathbf{q}, \mathbf{1}} + \eta_{\mathbf{q}, \mathbf{2}}\|_\infty \right\}. \quad (\text{B.18})$$

Since $\|\eta_{\mathbf{q}, \mathbf{1}} + \eta_{\mathbf{q}, \mathbf{2}}\|_\infty \leq \|\eta_{\mathbf{q}, \mathbf{1}}\|_\infty + \|\eta_{\mathbf{q}, \mathbf{2}}\|_\infty$, therefore using the bounds on $\|\eta_{\mathbf{q}, \mathbf{1}}\|_\infty, \|\eta_{\mathbf{q}, \mathbf{2}}\|_\infty$ and noting that $\|\eta_{\mathbf{q}, \mathbf{1}} + \eta_{\mathbf{q}, \mathbf{2}}\|_2 \leq \sqrt{m_{v'}} \|\eta_{\mathbf{q}, \mathbf{1}} + \eta_{\mathbf{q}, \mathbf{2}}\|_\infty$, we obtain for the stated choice of $m_{v'}$ (cf. Remark 2) that

$$\|\widehat{\nabla} \partial_q f(\mathbf{x}) - \nabla \partial_q f(\mathbf{x})\|_2 \leq C_2 \underbrace{\left(\frac{\mu_1 (4\rho_m + 1)B_3}{2\sqrt{m_{v'}}} + \frac{C_1 \sqrt{m_{v'}} \mu^2 ((4\rho_m + 1)k) B_3}{3m_v \mu_1} \right)}_{\tau'}. \quad (\text{B.19})$$

for $q = 1, \dots, d$, and $\forall \mathbf{x} \in [-1, 1]^d$. We next note that (B.19) trivially leads to the bound

$$\widehat{\partial_q \partial_{q'} f(\mathbf{x})} \in [\partial_q \partial_{q'} f(\mathbf{x}) - \tau', \partial_q \partial_{q'} f(\mathbf{x}) + \tau']; \quad q, q' = 1, \dots, d. \quad (\text{B.20})$$

Now if $q \notin \mathcal{S}_2^{\text{var}}$ then clearly $\widehat{\partial_q \partial_{q'} f(\mathbf{x})} \in [-\tau', \tau']; \forall \mathbf{x} \in [-1, 1]^d, q \neq q'$. On the other hand, if $(q, q') \in \mathcal{S}_2$ then

$$\widehat{\partial_q \partial_{q'} f(\mathbf{x})} \in [\partial_q \partial_{q'} \phi_{(q, q')}(x_q, x_{q'}) - \tau', \partial_q \partial_{q'} \phi_{(q, q')}(x_q, x_{q'}) + \tau']. \quad (\text{B.21})$$

If furthermore $m_x \geq \lambda_2^{-1}$, then due to the construction of χ , $\exists \mathbf{x} \in \chi$ so that $|\widehat{\partial_q \partial_{q'} f(\mathbf{x})}| \geq D_2 - \tau'$. Hence if $\tau' < D_2/2$ holds, then we would have $|\widehat{\partial_q \partial_{q'} f(\mathbf{x})}| > D_2/2$, leading to the identification of (q, q') . Since this is true for each $(q, q') \in \mathcal{S}_2$, hence it follows that $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$. Now, $\tau' < D_2/2$ is equivalent to

$$\underbrace{\frac{(4\rho_m + 1)B_3}{2\sqrt{m_{v'}}}}_a \mu_1 + \underbrace{\left(\frac{C_1 \sqrt{m_{v'}} ((4\rho_m + 1)k) B_3}{3m_v} \right)}_b \frac{\mu^2}{\mu_1} < \frac{D_2}{2C_2} \quad (\text{B.22})$$

$$\Leftrightarrow a\mu_1^2 - \frac{D_2}{2C_2} \mu_1 + b\mu^2 < 0 \quad (\text{B.23})$$

$$\Leftrightarrow \mu_1 \in \left((D_2/(4aC_2)) - \sqrt{(D_2/(4aC_2))^2 - (b\mu^2/a)}, (D_2/(4aC_2)) + \sqrt{(D_2/(4aC_2))^2 - (b\mu^2/a)} \right). \quad (\text{B.24})$$

Lastly, we see that the bounds in (B.24) are valid if:

$$\mu^2 < \frac{D_2^2}{16abC_2^2} = \frac{3D_2^2 m_v}{8C_1 C_2^2 B_3^2 (4\rho_m + 1)((4\rho_m + 1)k)}. \quad (\text{B.25})$$

Estimating \mathcal{S}_1 . With $\mathcal{P} := [d] \setminus \widehat{\mathcal{S}}_2^{\text{var}}$, we have via Taylor's expansion of f at $j = 1, \dots, m_{v''}$:

$$\frac{f((\mathbf{x} + \mu' \mathbf{v}_j'')_{\mathcal{P}}) - f((\mathbf{x} - \mu' \mathbf{v}_j'')_{\mathcal{P}})}{2\mu'} = \underbrace{\langle (\mathbf{v}_j'')_{\mathcal{P}}, (\nabla f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}} \rangle}_{n_j} + \frac{R_3((\zeta_j)_{\mathcal{P}}) - R_3((\zeta_j')_{\mathcal{P}})}{2\mu'}. \quad (\text{B.26})$$

(B.26) corresponds to linear measurements of the $(k - |\widehat{\mathcal{S}}_2^{\text{var}}|)$ sparse vector: $(\nabla f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}}$. Note that we effectively perform ℓ_1 minimization over $\mathbb{R}^{|\mathcal{P}|}$. Therefore for any $\mathbf{x} \in \mathbb{R}^d$ we immediately have from Theorem 4, (B.2), the following. $\exists C_3, c'_6 > 0, c'_3 \geq 1$ such that for $c'_3(k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \log(\frac{|\mathcal{P}|}{k - |\widehat{\mathcal{S}}_2^{\text{var}}|}) < m_{v''} < \frac{|\mathcal{P}|}{(\log 6)^2}$, we have with probability at least $1 - e^{-c'_6 m_{v''}} - e^{-\sqrt{m_{v''}}|\mathcal{P}|}$ that

$$\|(\widehat{\nabla} f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}} - (\nabla f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}}\|_2 \leq C_3 \max \left\{ \|\mathbf{n}\|_2, \sqrt{\log |\mathcal{P}|} \|\mathbf{n}\|_{\infty} \right\}, \quad (\text{B.27})$$

where $\mathbf{n} = [n_1 \dots n_{m_{v''}}]$. We now uniformly bound $R_3((\zeta_j)_{\mathcal{P}})$ for all $j = 1, \dots, m_{v''}$ and $\zeta_j \in [-(1+r), 1+r]^d$ as follows.

$$R_3((\zeta_j)_{\mathcal{P}}) = \frac{\mu'^3}{6} \sum_{p \in \mathcal{S}_1 \cap \mathcal{P}} \partial_p^3 \phi_p(\zeta_{j,p}) v_{j,p}''^3 \Rightarrow |R_3((\zeta_j)_{\mathcal{P}})| \leq \frac{(k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \mu'^3 B_3}{6m_{v''}^{3/2}}. \quad (\text{B.28})$$

This in turn implies that $\|\mathbf{n}\|_{\infty} \leq \frac{(k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \mu'^2 B_3}{6m_{v''}^{3/2}}$ and $\|\mathbf{n}\|_2 \leq \sqrt{m_{v''}} \|\mathbf{n}\|_{\infty} \leq \frac{(k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \mu'^2 B_3}{6m_{v''}}$. Plugging these bounds in (B.27), we obtain for the stated choice of $m_{v''}$ (cf. Remark 2) that

$$\|(\widehat{\nabla} f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}} - (\nabla f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}}\|_2 \leq \underbrace{\frac{C_3 (k - |\widehat{\mathcal{S}}_2^{\text{var}}|) \mu'^2 B_3}{6m_{v''}}}_{\tau''}; \quad \mathbf{x} \in [-1, 1]^d. \quad (\text{B.29})$$

Finally, using the same arguments as before, we have that $\tau'' < D_1/2$ or equivalently $\mu'^2 < \frac{3m_{v''} D_1}{C_3 (k - |\widehat{\mathcal{S}}_2^{\text{var}}|) B_3}$ is sufficient to recover \mathcal{S}_1 . This completes the proof.

B.2 Proof of Theorem 2

We prove a more detailed version of Theorem 2, stated below.

Theorem 5. *Assuming notation in Theorem 1, let $m_x, m'_x, m_v, m_{v'}, m_{v''}$ be as defined in Theorem 1. Say $\varepsilon < \varepsilon_1 = \frac{D_2^3}{192\sqrt{3}C_1 C_2^3 \sqrt{a^3 b m_v m_v}}$. Denoting $\theta_1 = \cos^{-1}(-\varepsilon/\varepsilon_1)$, $b' = 2C_1 \sqrt{m_v m_{v'}}$, we have for $\mu \in (\sqrt{4a'^2 a/(3b)} \cos(\theta_1/3) - 2\pi/3, \sqrt{4a'^2 a/(3b)} \cos(\theta_1/3))$ and $\mu_1 \in (a' - \sqrt{a'^2 - ((b\mu^2 + b'\varepsilon)/a)}, a' + \sqrt{a'^2 - ((b\mu^2 + b'\varepsilon)/a)})$ that $\tau' = C_2 \left(a\mu_1 + \frac{b\mu^2}{\mu_1} + \frac{b'\varepsilon}{\mu_1} \right)$ implies $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$ with high probability. Given $\widehat{\mathcal{S}}_2 = \mathcal{S}_2$, denote $a_1 = \frac{(k - |\widehat{\mathcal{S}}_2^{\text{var}}|) B_3}{6m_{v''}}$, $b_1 = \sqrt{m_{v''}}$ and say $\varepsilon < \varepsilon_2 = \frac{D_1^{3/2}}{3\sqrt{6a_1 C_3^3 b_1^2}}$. For $\theta_2 = \cos^{-1}(-\varepsilon/\varepsilon_2)$, let $\mu' \in (2\sqrt{D_1/(6a_1 C_3)} \cos(\theta_2/3) - 2\pi/3, 2\sqrt{D_1/(6a_1 C_3)} \cos(\theta_2/3))$. Then $\tau'' = C_3(a_1 \mu'^2 + \frac{b_1 \varepsilon}{\mu'})$ implies $\widehat{\mathcal{S}}_1 = \mathcal{S}_1$ with high probability.*

Proof. We begin by establishing the conditions pertaining to the estimation of \mathcal{S}_2 . Then we prove the conditions for estimation of \mathcal{S}_1 .

Estimation of \mathcal{S}_2 . We first note that the linear system (3.4) now has the form: $\mathbf{y} = \mathbf{V}\nabla f(\mathbf{x}) + \mathbf{n} + \mathbf{z}$ where $z_j = (z'_{j,1} - z'_{j,2})/(2\mu)$ represents the external noise component, for $j = 1, \dots, m_v$. Observe that $\|\mathbf{z}\|_\infty \leq \varepsilon/\mu$. Using the bounds on $\|\mathbf{n}\|_\infty, \|\mathbf{n}\|_2$ from Section B.1, we then observe that (B.9) changes to:

$$\|\widehat{\nabla} f(\mathbf{x}) - \nabla f(\mathbf{x})\|_2 \leq C_1 \left(\frac{\mu^2((4\rho_m + 1)k)B_3}{6m_v} + \frac{\varepsilon\sqrt{m_v}}{\mu} \right), \quad \forall \mathbf{x} \in [-(1+r), 1+r]^d. \quad (\text{B.30})$$

As a result, we then have that

$$\|\eta_{\mathbf{q},2}\|_\infty \leq C_1 \left(\frac{\mu^2((4\rho_m + 1)k)B_3}{3m_v\mu_1} + \frac{2\varepsilon\sqrt{m_v}}{\mu\mu_1} \right). \quad (\text{B.31})$$

Now note that the bound on $\|\eta_{\mathbf{q},1}\|_\infty$ is unchanged from Section B.1 i.e., $\|\eta_{\mathbf{q},1}\|_\infty \leq \frac{\mu_1(4\rho_m+1)B_3}{2m_{v'}}$. As a consequence, we see that (B.19) changes to:

$$\|\widehat{\nabla} \partial_q f(\mathbf{x}) - \nabla \partial_q f(\mathbf{x})\|_2 \leq C_2 \underbrace{\left(\frac{\mu_1(4\rho_m + 1)B_3}{2\sqrt{m_{v'}}} + C_1 \frac{\sqrt{m_{v'}}\mu^2((4\rho_m + 1)k)B_3}{3m_v\mu_1} + \frac{2C_1\varepsilon\sqrt{m_v m_{v'}}}{\mu\mu_1} \right)}_{\tau'}. \quad (\text{B.32})$$

With a and b as stated in the Theorem, we then see that $\tau' < D_2/2$ is equivalent to

$$a\mu_1^2 - \frac{D_2}{2C_2}\mu_1 + \left(b\mu^2 + \frac{2C_1\varepsilon\sqrt{m_v m_{v'}}}{\mu} \right) < 0. \quad (\text{B.33})$$

which in turn is equivalent to

$$\mu_1 \in \left(\frac{D_2}{4aC_2} - \sqrt{\left(\frac{D_2}{4aC_2} \right)^2 - \left(\frac{b\mu^3 + 2C_1\varepsilon\sqrt{m_v m_{v'}}}{a\mu} \right)}, \frac{D_2}{4aC_2} + \sqrt{\left(\frac{D_2}{4aC_2} \right)^2 - \left(\frac{b\mu^3 + 2C_1\varepsilon\sqrt{m_v m_{v'}}}{a\mu} \right)} \right). \quad (\text{B.34})$$

For the above bound to be valid, we require

$$\frac{b\mu^2}{a} + \frac{2C_1\varepsilon\sqrt{m_v m_{v'}}}{a\mu} < \frac{D_2^2}{16a^2C_2^2}, \quad (\text{B.35})$$

$$\Leftrightarrow \mu^3 - \frac{D_2^2}{16abC_2^2}\mu + \frac{2C_1\varepsilon\sqrt{m_v m_{v'}}}{b} < 0 \quad (\text{B.36})$$

to hold. (B.36) is a cubic inequality. A cubic equation of the form: $y^3 + py + q = 0$, has 3 distinct real roots if its discriminant $\frac{p^3}{27} + \frac{q^2}{4} < 0$. Note that for this to be possible, p must be negative, which is the case in (B.36). Applying this to (B.36) leads to the condition: $\varepsilon < \frac{D_2^3}{192\sqrt{3}C_1C_2^3\sqrt{a^3b m_{v'} m_v}} = \varepsilon_1$. Furthermore, the 3 distinct real roots are given by:

$$y_1 = 2\sqrt{-p/3} \cos(\theta/3), \quad y_2 = -2\sqrt{-p/3} \cos(\theta/3 + \pi/3), \quad y_3 = -2\sqrt{-p/3} \cos(\theta/3 - \pi/3) \quad (\text{B.37})$$

where $\theta = \cos^{-1}\left(\frac{-q/2}{\sqrt{-p^3/27}}\right)$. Applying this to (B.36) then leads to $\theta_1 = \cos^{-1}(-\varepsilon/\varepsilon_1)$. For $0 < \varepsilon < \varepsilon_1$ we have $\pi/2 < \theta_1 < \pi$ which implies $0 < y_2 < y_1$ and $y_3 < 0$. In particular if $q > 0$, then one can verify that $y^3 + py + q < 0$ holds if $y \in (y_2, y_1)$. Applying this to (B.36), we consequently obtain:

$$\mu \in \left(\sqrt{\frac{D_2^2}{12abC_2^2}} \cos(\theta_1/3 - 2\pi/3), \sqrt{\frac{D_2^2}{12abC_2^2}} \cos(\theta_1/3) \right). \quad (\text{B.38})$$

Estimation of \mathcal{S}_1 . We now prove the conditions for estimation of \mathcal{S}_1 . First note that (B.26) now changes to:

$$\frac{f((\mathbf{x} + \mu' \mathbf{v}''_j)_{\mathcal{P}}) - f((\mathbf{x} - \mu' \mathbf{v}''_j)_{\mathcal{P}})}{2\mu'} = \langle (\mathbf{v}''_j)_{\mathcal{P}}, (\nabla f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}} \rangle + \underbrace{\frac{R_3((\zeta_j)_{\mathcal{P}}) - R_3((\zeta'_j)_{\mathcal{P}})}{2\mu'}}_{n_j} + \underbrace{\frac{z'_{j,1} - z'_{j,2}}{2\mu'}}_{z_j}, \quad (\text{B.39})$$

for $j = 1, \dots, m_{v''}$. Denoting $\mathbf{z} = [z_1 \cdots z_{m_{v''}}]$, we have $\|\mathbf{z}\|_\infty \leq \varepsilon/\mu'$. As the bounds on $\|\mathbf{n}\|_2, \|\mathbf{n}\|_\infty$ are unchanged, therefore (B.40) now changes to:

$$\|(\widehat{\nabla} f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}} - (\nabla f((\mathbf{x})_{\mathcal{P}}))_{\mathcal{P}}\|_2 \leq C_3 \underbrace{\left(\frac{(k - |\widehat{\mathcal{S}}_2^{\text{var}}|)\mu'^2 B_3}{6m_{v''}} + \frac{\varepsilon\sqrt{m_{v''}}}{\mu'} \right)}_{\tau''}; \quad \mathbf{x} \in [-1, 1]^d. \quad (\text{B.40})$$

Denoting $a_1 = \frac{(k - |\widehat{\mathcal{S}}_2^{\text{var}}|)B_3}{6m_{v''}}$, $b_1 = \sqrt{m_{v''}}$, we then see from (B.40) that the condition $\tau'' < D_1/2$ is equivalent to

$$\mu'^3 - \frac{D_1}{2a_1 C_3} \mu' + \frac{b_1 \varepsilon}{a_1} < 0. \quad (\text{B.41})$$

As discussed earlier for estimation of \mathcal{S}_2 , the cubic equation corresponding to (B.41) has 3 distinct real roots if its discriminant is negative. This then leads to the condition $\varepsilon < \frac{D_1^{3/2}}{3\sqrt{6a_1 C_3^3 b_1^2}} = \varepsilon_2$. Then by using the expressions for the roots of the cubic from (B.37), one can verify that (B.41) holds if

$$\mu' \in (2\sqrt{D_1/(6a_1 C_3)} \cos(\theta_2/3 - 2\pi/3), 2\sqrt{D_1/(6a_1 C_3)} \cos(\theta_2/3)) \quad (\text{B.42})$$

with $\theta_2 = \cos^{-1}(-\varepsilon/\varepsilon_2)$. This completes the proof. \square

B.3 Proof of Theorem 3

We first derive conditions for estimating \mathcal{S}_2 , and then for \mathcal{S}_1 .

Estimating \mathcal{S}_2 . Upon resampling N_1 times and averaging, we have for the noise vector $\mathbf{z} \in \mathbb{R}^{m_v}$ where

$$\mathbf{z} = \left[\frac{(z'_{1,1} - z'_{1,2})}{2\mu} \cdots \frac{(z'_{m_v,1} - z'_{m_v,2})}{2\mu} \right], \quad (\text{B.43})$$

that $z'_{j,1}, z'_{j,2} \sim \mathcal{N}(0, \sigma^2/N_1)$ are i.i.d. Note that it is in fact sufficient to guarantee that $|z'_{j,1} - z'_{j,2}| < 2\varepsilon$ holds $\forall j = 1, \dots, m_v$, and across all points where ∇f is estimated. Indeed, we can then simply use the proof in Section B.2, for the setting of arbitrary bounded noise. To this end, note that $z'_{j,1} - z'_{j,2} \sim \mathcal{N}(0, \frac{2\sigma^2}{N_1})$. It can be shown for $X \sim \mathcal{N}(0, 1)$ that:

$$\mathbb{P}(|X| > t) \leq \frac{2e^{-t^2/2}}{t}, \quad \forall t > 0. \quad (\text{B.44})$$

Since $z'_{j,1} - z'_{j,2} = \sigma\sqrt{\frac{2}{N_1}}X$ therefore for any $\varepsilon > 0$ we have that:

$$\mathbb{P}(|z'_{j,1} - z'_{j,2}| > 2\varepsilon) = \mathbb{P}\left(|X| > \frac{2\varepsilon}{\sigma}\sqrt{\frac{N_1}{2}}\right) \quad (\text{B.45})$$

$$\leq \frac{\sigma}{\varepsilon}\sqrt{\frac{2}{N_1}} \exp\left(-\frac{\varepsilon^2 N_1}{\sigma^2}\right) \quad (\text{B.46})$$

$$\leq \frac{\sqrt{2}\sigma}{\varepsilon} \exp\left(-\frac{\varepsilon^2 N_1}{\sigma^2}\right). \quad (\text{B.47})$$

Now to estimate $\nabla f(\mathbf{x})$ we have m_v many ‘‘difference’’ terms: $z'_{j,1} - z'_{j,2}$. We additionally estimate $m_{v'}$ many gradients at each \mathbf{x} implying a total of $m_v(m_{v'} + 1)$ difference terms. As this is done for each $\mathbf{x} \in \mathcal{X}$, therefore we have a total of $m_v(m_{v'} + 1)(2m_x + 1)^2 |\mathcal{H}_2^d|$ many difference terms. Taking a union bound over all of them, we have for any $p_1 \in (0, 1), \varepsilon > 0$ that the choice $N_1 > \frac{\sigma^2}{\varepsilon^2} \log(\frac{\sqrt{2}\sigma}{\varepsilon p_1} m_v(m_{v'} + 1)(2m_x + 1)^2 |\mathcal{H}_2^d|)$ implies that the magnitudes of all difference terms are bounded by 2ε , with probability at least $1 - p_1$. Thereafter, we can simply follow the proof in Section B.2, for estimating \mathcal{S}_2 in the presence of arbitrary bounded noise.

Estimating \mathcal{S}_1 . In this case, we resample each query N_2 times and average – therefore the variance of the noise terms gets scaled by N_2 . We now have $|\chi_{\text{diag}}|m_{v''} = (2m'_x + 1)m_{v''}$ many “difference” terms corresponding to Gaussian noise. Therefore, taking a union bound over all of them, we have for any $p_2 \in (0, 1), \varepsilon' > 0$ that the choice $N_2 > \frac{\sigma^2}{\varepsilon'^2} \log\left(\frac{\sqrt{2}\sigma(2m'_x+1)m_{v''}}{\varepsilon'p_2}\right)$ implies that the magnitudes of all difference terms are bounded by $2\varepsilon'$, with probability at least $1 - p_2$. Thereafter, we can simply follow the proof in Section B.2, for estimating \mathcal{S}_1 in the presence of arbitrary bounded noise. The only change there would be to replace ε by ε' .

C Learning individual components of model

Recall from (2.2) the unique representation of the model:

$$f(x_1, \dots, x_d) = c + \sum_{p \in \mathcal{S}_1} \phi_p(x_p) + \sum_{(l, l') \in \mathcal{S}_2} \phi_{(l, l')}(x_l, x_{l'}) + \sum_{q \in \mathcal{S}_2^{\text{var}}: \rho(q) > 1} \phi_q(x_q), \quad (\text{C.1})$$

where $\mathcal{S}_1 \cap \mathcal{S}_2^{\text{var}} = \emptyset$. Having estimated the sets \mathcal{S}_1 and \mathcal{S}_2 , we now show how the individual univariate and bivariate functions in the model can be estimated. We will see this for the settings of noiseless, as well as noisy (arbitrary, bounded noise and stochastic noises) point queries.

C.1 Noiseless queries

In this scenario, we obtain the exact value $f(\mathbf{x})$ at each query $\mathbf{x} \in \mathbb{R}^d$. Let us first see how each ϕ_p ; $p \in \mathcal{S}_1$ can be estimated. For some $-1 = t_1 < t_2 < \dots < t_n = 1$, consider the set

$$\chi_p := \left\{ \mathbf{x}_i \in \mathbb{R}^d : (\mathbf{x}_i)_j = \begin{cases} t_i; & j = p, \\ 0; & j \neq p \end{cases}; 1 \leq i \leq n; 1 \leq j \leq d \right\}; \quad p \in \mathcal{S}_1. \quad (\text{C.2})$$

We obtain the samples $\{f(\mathbf{x}_i)\}_{i=1}^n$; $\mathbf{x}_i \in \chi_p$. Here $f(\mathbf{x}_i) = \phi_p(t_i) + C$ with C being a constant that depends on the other components in the model. Given the samples, one can then employ spline based “quasi interpolant operators” [24], to obtain an estimate $\tilde{\phi}_p : [-1, 1] \rightarrow \mathbb{R}$, to $\phi_p + C$. Construction of such operators can be found for instance in [24] (see also [38]). One can suitably choose the t_i 's and construct quasi interpolants that approximate any C^m smooth univariate function with optimal $L_\infty[-1, 1]$ error rate $O(n^{-m})$ [24, 38]. Having obtained $\tilde{\phi}_p$, we then define

$$\hat{\phi}_p := \tilde{\phi}_p - \mathbb{E}_p[\tilde{\phi}_p]; \quad p \in \mathcal{S}_1, \quad (\text{C.3})$$

to be the estimate of ϕ_p . The bivariate components corresponding to each $(l, l') \in \mathcal{S}_2$ can be estimated in a similar manner as above. To this end, for some strictly increasing sequences: $(-1 = t'_1, t'_2, \dots, t'_{n_1} = 1)$, $(-1 = t_1, t_2, \dots, t_{n_1} = 1)$, consider the set

$$\chi_{(l, l')} := \left\{ \mathbf{x}_{i, j} \in \mathbb{R}^d : (\mathbf{x}_{i, j})_q = \begin{cases} t'_i; & q = l, \\ t_j; & q = l', \\ 0; & q \neq l, l' \end{cases}; 1 \leq i, j \leq n_1; 1 \leq q \leq d \right\}; \quad (l, l') \in \mathcal{S}_2. \quad (\text{C.4})$$

We then obtain the samples $\{f(\mathbf{x}_{i, j})\}_{i, j=1}^{n_1}$; $\mathbf{x}_{i, j} \in \chi_{(l, l')}$ where

$$\begin{aligned} f(\mathbf{x}_{i, j}) &= \phi_{(l, l')}(t'_i, t_j) + \sum_{\substack{l_1: (l, l_1) \in \mathcal{S}_2 \\ l_1 \neq l'}} \phi_{(l, l_1)}(t'_i, 0) + \sum_{\substack{l_1: (l_1, l) \in \mathcal{S}_2 \\ l_1 \neq l'}} \phi_{(l_1, l)}(0, t'_i) \\ &+ \sum_{\substack{l'_1: (l', l'_1) \in \mathcal{S}_2 \\ l'_1 \neq l'}} \phi_{(l', l'_1)}(t_j, 0) + \sum_{\substack{l'_1: (l'_1, l') \in \mathcal{S}_2 \\ l'_1 \neq l'}} \phi_{(l'_1, l')}(0, t_j) + \phi_l(t'_i) + \phi_{l'}(t_j) + C, \end{aligned} \quad (\text{C.5})$$

$$= g_{(l, l')}(t'_i, t_j) + C, \quad (\text{C.6})$$

with C being a constant. (C.5) is a general expression – if for example $\rho(l) = 1$, then the terms $\phi_l, \phi_{(l, l_1)}, \phi_{(l_1, l)}$ will be zero. Given this, we can again obtain estimates $\tilde{\phi}_{(l, l')} : [-1, 1]^2 \rightarrow \mathbb{R}$ to $g_{(l, l')} + C$, via spline based quasi interpolants. Let us denote $n = n_1^2$ to be the total number of samples of f . For an appropriate choice of (t'_i, t_j) 's, one can construct bivariate

quasi interpolants that approximate any C^m smooth bivariate function, with optimal $L_\infty[-1, 1]^2$ error rate $O(n^{-m/2})$ [24, 38]. Subsequently, we define the final estimates $\widehat{\phi}_{(l, \nu)}$ to $\phi_{(l, \nu)}$ as follows.

$$\widehat{\phi}_{(l, \nu)} := \begin{cases} \tilde{\phi}_{(l, \nu)} - \mathbb{E}_{(l, \nu)}[\tilde{\phi}_{(l, \nu)}]; & \rho(l), \rho(l') = 1, \\ \tilde{\phi}_{(l, \nu)} - \mathbb{E}_l[\tilde{\phi}_{(l, \nu)}]; & \rho(l) = 1, \rho(l') > 1, \\ \tilde{\phi}_{(l, \nu)} - \mathbb{E}_{\nu'}[\tilde{\phi}_{(l, \nu)}]; & \rho(l) > 1, \rho(l') = 1, \\ \tilde{\phi}_{(l, \nu)} - \mathbb{E}_l[\tilde{\phi}_{(l, \nu)}] - \mathbb{E}_{\nu'}[\tilde{\phi}_{(l, \nu)}] + \mathbb{E}_{(l, \nu)}[\tilde{\phi}_{(l, \nu)}]; & \rho(l) > 1, \rho(l') > 1. \end{cases} \quad (\text{C.7})$$

Lastly, we require to estimate the univariate's : ϕ_l for each $l \in \mathcal{S}_2^{\text{var}}$ such that $\rho(l) > 1$. As above, for some strictly increasing sequences: $(-1 = t'_1, t'_2, \dots, t'_{n_1} = 1)$, $(-1 = t_1, t_2, \dots, t_{n_1} = 1)$, consider the set

$$\chi_l := \left\{ \mathbf{x}_{i,j} \in \mathbb{R}^d : (\mathbf{x}_{i,j})_q = \begin{cases} t'_i; & q = l, \\ t_j; & q \neq l \text{ \& } q \in \mathcal{S}_2^{\text{var}}, \\ 0; & q \notin \mathcal{S}_2^{\text{var}}, \end{cases} ; 1 \leq i, j \leq n_1; 1 \leq q \leq d \right\}; \quad l \in \mathcal{S}_2^{\text{var}} : \rho(l) > 1. \quad (\text{C.8})$$

We obtain $\{f(\mathbf{x}_{i,j})\}_{i,j=1}^{n_1}$; $\mathbf{x}_{i,j} \in \chi_l$ where this time

$$f(\mathbf{x}_{i,j}) = \phi_l(t'_i) + \sum_{\rho(l') > 1, l' \neq l} \phi_{l'}(t_j) + \sum_{\nu': (l, \nu') \in \mathcal{S}_2} \phi_{(l, \nu')}(t'_i, t_j) \quad (\text{C.9})$$

$$+ \sum_{\nu': (l, \nu') \in \mathcal{S}_2} \phi_{(\nu', l)}(t_j, t'_i) + \sum_{(q, q') \in \mathcal{S}_2 : q, q' \neq l} \phi_{(q, q')}(t_j, t'_i) + C \quad (\text{C.10})$$

$$= g_l(t'_i, t_j) + C \quad (\text{C.11})$$

for a constant, C . Denoting $n = n_1^2$ to be the total number of samples of f , we can again obtain an estimate $\tilde{\phi}_l(x_l, x)$ to $g_l(x_l, x) + C$, with $L_\infty[-1, 1]^2$ error rate $O(n^{-3/2})$. Then with $\tilde{\phi}_l$ at hand, we define the estimate $\widehat{\phi}_l : [-1, 1] \rightarrow \mathbb{R}$ as

$$\widehat{\phi}_l := \mathbb{E}_x[\tilde{\phi}_l] - \mathbb{E}_{(l, x)}[\tilde{\phi}_l]; \quad l \in \mathcal{S}_2^{\text{var}} : \rho(l) > 1. \quad (\text{C.12})$$

The following proposition formally describes the error rates for the aforementioned estimates.

Proposition 1. For C^3 smooth components $\phi_p, \phi_{(l, \nu)}, \phi_l$, let $\widehat{\phi}_p, \widehat{\phi}_{(l, \nu)}, \widehat{\phi}_l$ be the respective estimates as defined in (C.3), (C.7) and (C.12) respectively. Also, let n denote the number of queries (of f) made per component. We then have that:

1. $\|\widehat{\phi}_p - \phi_p\|_{L_\infty[-1, 1]} = O(n^{-3}); \forall p \in \mathcal{S}_1$,
2. $\|\widehat{\phi}_{(l, \nu)} - \phi_{(l, \nu)}\|_{L_\infty[-1, 1]^2} = O(n^{-3/2}); \forall (l, \nu') \in \mathcal{S}_2$, and
3. $\|\widehat{\phi}_l - \phi_l\|_{L_\infty[-1, 1]} = O(n^{-3/2}); \forall l \in \mathcal{S}_2^{\text{var}} : \rho(l) > 1$.

Proof. 1. $\mathbf{p} \in \mathcal{S}_1$.

We have for $\tilde{\phi}_p$ that $\|\tilde{\phi}_p - (\phi_p + C)\|_{L_\infty[-1, 1]} = O(n^{-3})$. Denoting $\tilde{\phi}_p(x_p) - (\phi_p(x_p) + C) = z_p(x_p)$, this means $|z_p(x_p)| = O(n^{-3}), \forall x_p \in [-1, 1]$. Now $|\mathbb{E}_p[\tilde{\phi}_p - (\phi_p + C)]| = |\mathbb{E}_p[\tilde{\phi}_p] - C| = |\mathbb{E}_p[z_p]| \leq \mathbb{E}_p[|z_p|] = O(n^{-3})$.

Lastly, we have that:

$$\|\widehat{\phi}_p - \phi_p\|_{L_\infty[-1, 1]} = \|\tilde{\phi}_p - \mathbb{E}_p[\tilde{\phi}_p] - \phi_p\|_{L_\infty[-1, 1]} \quad (\text{C.13})$$

$$= \|\tilde{\phi}_p - (\phi_p + C) - (\mathbb{E}_p[\tilde{\phi}_p] - C)\|_{L_\infty[-1, 1]} \quad (\text{C.14})$$

$$= O(n^{-3}). \quad (\text{C.15})$$

2. $(l, \nu') \in \mathcal{S}_2$.

We only consider the case where $\rho(l), \rho(l') > 1$ as proofs for the other cases are similar. Now for $\tilde{\phi}_{(l, \nu)}$ we have that $\|\tilde{\phi}_{(l, \nu)} - (g_{(l, \nu)} + C)\|_{L_\infty[-1, 1]^2} = O(n^{-3/2})$. Denoting $\tilde{\phi}_{(l, \nu)}(x_l, x_{\nu'}) - (g_{(l, \nu)}(x_l, x_{\nu'}) + C) = z_{(l, \nu)}(x_l, x_{\nu'})$, this means $|z_{(l, \nu)}(x_l, x_{\nu'})| = O(n^{-3/2}), \forall (x_l, x_{\nu'}) \in [-1, 1]^2$. Consequently, one can easily verify that:

$$\|\mathbb{E}_l[\tilde{\phi}_{(l,\nu)}] - (\mathbb{E}_l[g_{(l,\nu)}] + C)\|_{L_\infty[-1,1]} = O(n^{-3/2}), \quad (\text{C.16})$$

$$\|\mathbb{E}_{\nu'}[\tilde{\phi}_{(l,\nu)}] - (\mathbb{E}_{\nu'}[g_{(l,\nu)}] + C)\|_{L_\infty[-1,1]} = O(n^{-3/2}), \quad (\text{C.17})$$

$$\|\mathbb{E}_{(l,\nu)}[\tilde{\phi}_{(l,\nu)}] - (\mathbb{E}_{(l,\nu)}[g_{(l,\nu)}] + C)\|_{L_\infty} = O(n^{-3/2}). \quad (\text{C.18})$$

Now note that using the form for $g_{(l,\nu)}$ from (C.5), we have that

$$\begin{aligned} \mathbb{E}_l[g_{(l,\nu)}] &= \sum_{\substack{l_1:(l,l_1) \in \mathcal{S}_2 \\ l_1 \neq l'}} \mathbb{E}_l[\phi_{(l,l_1)}(x_l, 0)] + \sum_{\substack{l_1:(l_1,l) \in \mathcal{S}_2 \\ l_1 \neq l'}} \mathbb{E}_l[\phi_{(l_1,l)}(0, x_l)] + \sum_{\substack{l'_1:(l',l'_1) \in \mathcal{S}_2 \\ l'_1 \neq l'}} \phi_{(l',l'_1)}(x_{l'}, 0) \\ &+ \sum_{\substack{l'_1:(l'_1,l') \in \mathcal{S}_2 \\ l'_1 \neq l'}} \phi_{(l'_1,l')}(0, x_{l'}) + \phi_{\nu'}(x_{\nu'}) + C, \quad \text{and} \end{aligned} \quad (\text{C.19})$$

$$\begin{aligned} \mathbb{E}_{\nu'}[g_{(l,\nu)}] &= \sum_{\substack{l_1:(l,l_1) \in \mathcal{S}_2 \\ l_1 \neq l'}} \phi_{(l,l_1)}(x_l, 0) + \sum_{\substack{l_1:(l_1,l) \in \mathcal{S}_2 \\ l_1 \neq l'}} \phi_{(l_1,l)}(0, x_l) + \sum_{\substack{l'_1:(l',l'_1) \in \mathcal{S}_2 \\ l'_1 \neq l'}} \mathbb{E}_{\nu'}[\phi_{(l',l'_1)}(x_{l'}, 0)] \\ &+ \sum_{\substack{l'_1:(l'_1,l') \in \mathcal{S}_2 \\ l'_1 \neq l'}} \mathbb{E}_{\nu'}\phi_{(l'_1,l')}(0, x_{l'}) + \phi_l(x_l) + C, \quad \text{and} \end{aligned} \quad (\text{C.20})$$

$$\begin{aligned} \mathbb{E}_{(l,\nu)}[g_{(l,\nu)}] &= \sum_{\substack{l_1:(l,l_1) \in \mathcal{S}_2 \\ l_1 \neq l'}} \mathbb{E}_l[\phi_{(l,l_1)}(x_l, 0)] + \sum_{\substack{l_1:(l_1,l) \in \mathcal{S}_2 \\ l_1 \neq l'}} \mathbb{E}_l[\phi_{(l_1,l)}(0, x_l)] \\ &+ \sum_{\substack{l'_1:(l',l'_1) \in \mathcal{S}_2 \\ l'_1 \neq l'}} \mathbb{E}_{\nu'}[\phi_{(l',l'_1)}(x_{l'}, 0)] + \sum_{\substack{l'_1:(l'_1,l') \in \mathcal{S}_2 \\ l'_1 \neq l'}} \mathbb{E}_{\nu'}\phi_{(l'_1,l')}(0, x_{l'}) + C. \end{aligned} \quad (\text{C.21})$$

We then have from (C.5), (C.19), (C.20), (C.21) that

$$g_{(l,\nu)} - \mathbb{E}_l[g_{(l,\nu)}] - \mathbb{E}_{\nu'}[g_{(l,\nu)}] + \mathbb{E}_{(l,\nu)}[g_{(l,\nu)}] = \phi_{(l,\nu)}. \quad (\text{C.22})$$

Using (C.16), (C.17), (C.18), (C.22), and (C.7) it then follows that:

$$\|\widehat{\phi}_{(l,\nu)} - \phi_{(l,\nu)}\|_{L_\infty[-1,1]^2} = O(n^{-3/2}). \quad (\text{C.23})$$

3. $\mathbf{1} \in \mathcal{S}_2^{\text{var}} : \rho(\mathbf{1}) > \mathbf{1}$.

In this case, for $\tilde{\phi}_l : [-1, 1]^2 \rightarrow \mathbb{R}$, we have that $\|\tilde{\phi}_l - (g_l + C)\|_{L_\infty[-1,1]^2} = O(n^{-3/2})$, with

$$\begin{aligned} g_l(x_l, x) &= \phi_l(x_l) + \sum_{\rho(l') > 1, l' \neq l} \phi_{\nu'}(x) + \sum_{\nu':(l,\nu') \in \mathcal{S}_2} \phi_{(l,\nu')}(x_l, x) \\ &+ \sum_{\nu':(l',l) \in \mathcal{S}_2} \phi_{(l',l)}(x, x_l) + \sum_{(q,q') \in \mathcal{S}_2: q, q' \neq l} \phi_{(q,q')}(x, x). \end{aligned} \quad (\text{C.24})$$

From (C.24), we see that:

$$\mathbb{E}_x[g_l(x_l, x)] = \phi_l(x_l) + \sum_{(q,q') \in \mathcal{S}_2: q, q' \neq l} \mathbb{E}_x[\phi_{(q,q')}(x, x)], \quad (\text{C.25})$$

$$\text{and } \mathbb{E}_{(l,x)}[g_l(x_l, x)] = \sum_{(q,q') \in \mathcal{S}_2: q, q' \neq l} \mathbb{E}_x[\phi_{(q,q')}(x, x)]. \quad (\text{C.26})$$

Hence clearly, $\mathbb{E}_x[g_l(x_l, x)] - \mathbb{E}_{(l,x)}[g_l(x_l, x)] = \phi_l(x_l)$. One can also easily verify that

$$\|\mathbb{E}_x[\tilde{\phi}_l] - (\mathbb{E}_x[g_l] + C)\|_{L_\infty[-1,1]} = O(n^{-3/2}), \quad (\text{C.27})$$

$$\|\mathbb{E}_{(l,x)}[\tilde{\phi}_l] - (\mathbb{E}_{(l,x)}[g_l] + C)\|_{L_\infty} = O(n^{-3/2}). \quad (\text{C.28})$$

Therefore it follows that

$$\|\widehat{\phi}_l - \phi_l\|_{L_\infty[-1,1]} = \|(\mathbb{E}_x[\tilde{\phi}_l] - \mathbb{E}_{(l,x)}[\tilde{\phi}_l]) - (\mathbb{E}_x[g_l] - \mathbb{E}_{(l,x)}[g_l])\|_{L_\infty[-1,1]} \quad (\text{C.29})$$

$$\leq \|\mathbb{E}_x[\tilde{\phi}_l] - (\mathbb{E}_x[g_l] + C)\|_{L_\infty[-1,1]} + \|\mathbb{E}_{(l,x)}[\tilde{\phi}_l] - (\mathbb{E}_{(l,x)}[g_l] + C)\|_{L_\infty} \quad (\text{C.30})$$

$$= O(n^{-3/2}). \quad (\text{C.31})$$

This completes the proof. \square

C.2 Noisy queries

We now look at the case where for each query $\mathbf{x} \in \mathbb{R}^d$, we obtain a noisy value $f(\mathbf{x}) + z'$.

Arbitrary bounded noise. We begin with the scenario where z'_i is arbitrary and bounded with $|z'_i| < \varepsilon$; $\forall i$. Since the noise is arbitrary in nature, therefore we simply proceed *as in the noiseless case*, i.e., by approximating each component via a quasi-interpolant. As the magnitude of the noise is bounded by ε , it results in an additional $O(\varepsilon)$ term in the approximation error rates of Proposition 1.

To see this for the univariate case, let us denote $Q : C(\mathbb{R}) \rightarrow \mathcal{H}$ to be a quasi-interpolant operator. This a linear operator, with $C(\mathbb{R})$ denoting the space of continuous functions defined over \mathbb{R} and \mathcal{H} denoting a univariate spline space. Consider $u \in C^m[-1, 1]$ for some positive integer m , and let $g : [-1, 1] \rightarrow \mathbb{R}$ be an arbitrary continuous function with $\|g\|_{L_\infty[-1,1]} < \varepsilon$. Denote $\widehat{u} = u + g$ to be the ‘‘corrupted’’ version of u , and let n be the number of samples of \widehat{u} used by Q . We then have by linearity of Q that:

$$\|Q(\widehat{u}) - u\|_{L_\infty[-1,1]} = \|Q(u) + Q(g) - u\|_{L_\infty[-1,1]} \leq \underbrace{\|Q(u) - u\|_{L_\infty[-1,1]}}_{=O(n^{-m})} + \|Q\| \underbrace{\|g\|_{L_\infty[-1,1]}}_{\leq \|Q\|\varepsilon}, \quad (\text{C.32})$$

with $\|Q\|$ being the operator norm of Q . One can construct Q with $\|Q\|$ bounded¹¹ from above by a constant depending only on m . The above argument can be extended easily to the multivariate case. We state this for the bivariate case for completeness. Denote $Q_1 : C(\mathbb{R}^2) \rightarrow \mathcal{H}$ to be a quasi-interpolant operator, with \mathcal{H} denoting a bivariate spline space. Consider $u_1 \in C^m[-1, 1]^2$ for some positive integer m , and let $g_1 : [-1, 1] \rightarrow \mathbb{R}$ be an arbitrary continuous function with $\|g_1\|_{L_\infty[-1,1]^2} < \varepsilon$. Let $\widehat{u}_1 = u_1 + g_1$ and let n be the number of samples of \widehat{u}_1 used by Q_1 . We then have by linearity of Q_1 that:

$$\|Q_1(\widehat{u}_1) - u_1\|_{L_\infty[-1,1]^2} = \|Q_1(u_1) + Q_1(g_1) - u_1\|_{L_\infty[-1,1]^2} \leq \underbrace{\|Q_1(u_1) - u_1\|_{L_\infty[-1,1]^2}}_{=O(n^{-m/2})} + \|Q_1\| \underbrace{\|g_1\|_{L_\infty[-1,1]^2}}_{\leq \|Q_1\|\varepsilon}, \quad (\text{C.33})$$

with $\|Q_1\|$ being the operator norm of Q_1 . As for the univariate case, one can construct Q_1 with $\|Q_1\|$ bounded¹¹ from above by a constant depending only on m .

Let us define our final estimates $\widehat{\phi}_p$, $\widehat{\phi}_{(l,l')}$ and $\widehat{\phi}_l$ as in (C.3), (C.7) and (C.12), respectively. The following proposition formally states the error bounds, for this particular noise model.

Proposition 2 (Arbitrary bounded noise). *For C^3 smooth components $\phi_p, \phi_{(l,l')}, \phi_l$, let $\widehat{\phi}_p, \widehat{\phi}_{(l,l')}, \widehat{\phi}_l$ be the respective estimates as defined in (C.3), (C.7) and (C.12) respectively. Also, let n denote the number of noisy queries (of f) made per component with the external noise magnitude being bounded by ε . We then have that*

1. $\|\widehat{\phi}_p - \phi_p\|_{L_\infty[-1,1]} = O(n^{-3}) + O(\varepsilon); \forall p \in \mathcal{S}_1$,
2. $\|\widehat{\phi}_{(l,l')} - \phi_{(l,l')}\|_{L_\infty[-1,1]^2} = O(n^{-3/2}) + O(\varepsilon); \forall (l, l') \in \mathcal{S}_2$, and
3. $\|\widehat{\phi}_l - \phi_l\|_{L_\infty[-1,1]} = O(n^{-3/2}) + O(\varepsilon); \forall l \in \mathcal{S}_2^{\text{var}} : \rho(l) > 1$.

The proof is similar to that of Proposition 1 and hence skipped.

¹¹For instance, see Theorems 14.4, 15.2 in [38]

Stochastic noise. We now consider the setting where $z'_i \sim \mathcal{N}(0, \sigma^2)$ are i.i.d Gaussian random variables. Similar to the noiseless case, estimating the individual components again involves sampling f along the subspaces corresponding to $\mathcal{S}_1, \mathcal{S}_2$. Due to the presence of stochastic noise however, we now make use of *nonparametric regression* techniques to compute the estimates. While there exist a number of methods that could be used for this purpose (cf. [39]), we only discuss a specific one for clarity of exposition.

To elaborate, we again construct the sets defined in (C.2),(C.4) and(C.8). In particular, we uniformly discretize the domains $[-1, 1]$ and $[-1, 1]^2$, by choosing the respective t_i 's and (t'_i, t'_j) 's accordingly. This is the so called ‘‘fixed design’’ setting in nonparametric statistics. Upon collecting the samples $\{f(\mathbf{x}_i) + z'_i\}_{i=1}^n$ one can then derive estimates $\tilde{\phi}_p, \tilde{\phi}_{(l,l')}, \tilde{\phi}_l$, to $\phi_p + C, g_{(l,l')} + C$ and $g_l + C$ respectively, by using *local polynomial estimators* (cf. [39, 40] and references within). It is known that these estimators achieve the (minimax optimal) L_∞ error rate: $\Omega((n^{-1} \log n)^{\frac{m}{2m+d}})$, for estimating d -variate, C^m smooth functions over compact domains¹². Translated to our setting, we then have that the functions: $\phi_p + C, g_{(l,l')} + C$ and $g_l + C$ are estimated at the rates: $O((n^{-1} \log n)^{\frac{3}{7}})$ and $O((n^{-1} \log n)^{\frac{3}{8}})$ respectively.

Denoting the above intermediate estimates by $\tilde{\phi}_p, \tilde{\phi}_{(l,l')}, \tilde{\phi}_l$, we define our final estimates $\hat{\phi}_p, \hat{\phi}_{(l,l')}$ and $\hat{\phi}_l$ as in (C.3), (C.7) and (C.12), respectively. The following Proposition describes the error rates of these estimates.

Proposition 3 (i.i.d Gaussian noise). *For C^3 smooth components $\phi_p, \phi_{(l,l')}, \phi_l$, let $\hat{\phi}_p, \hat{\phi}_{(l,l')}, \hat{\phi}_l$ be the respective estimates as defined in (C.3), (C.7) and (C.12) respectively. Let n denote the number of noisy queries (of f) made per component, with noise samples z'_1, z'_2, \dots, z'_n being i.i.d Gaussian. Furthermore, let $\mathbb{E}_z[\cdot]$ denote expectation w.r.t the joint distribution of z'_1, z'_2, \dots, z'_n . We then have that*

1. $\mathbb{E}_z[\|\hat{\phi}_p - \phi_p\|_{L_\infty[-1,1]}] = O((n^{-1} \log n)^{\frac{3}{7}}); \forall p \in \mathcal{S}_1,$
2. $\mathbb{E}_z[\|\hat{\phi}_{(l,l')} - \phi_{(l,l')}\|_{L_\infty[-1,1]^2}] = O((n^{-1} \log n)^{\frac{3}{8}}); \forall (l, l') \in \mathcal{S}_2, \text{ and}$
3. $\mathbb{E}_z[\|\hat{\phi}_l - \phi_l\|_{L_\infty[-1,1]}] = O((n^{-1} \log n)^{\frac{3}{8}}); \forall l \in \mathcal{S}_2^{var} : \rho(l) > 1.$

Although the proof is again very similar to that of Proposition 1, there are some technical differences. Hence we provide a brief sketch of the proof, avoiding details already highlighted in the proof of Proposition 1.

Proof. 1. $\mathbf{p} \in \mathcal{S}_1$.

We have for $\tilde{\phi}_p$ that $\mathbb{E}_z[\|\tilde{\phi}_p - (\phi_p + C)\|_{L_\infty[-1,1]}] = O((n^{-1} \log n)^{\frac{3}{7}})$. Denoting $\tilde{\phi}_p(x_p) - (\phi_p(x_p) + C) = b_p(x_p)$, this means $\mathbb{E}_z[|b_p(x_p)|] = O((n^{-1} \log n)^{\frac{3}{7}})$. Now,

$$\mathbb{E}_z[\|\mathbb{E}_p[\tilde{\phi}_p - (\phi_p + C)]\|] = \mathbb{E}_z[\|\mathbb{E}_p[b_p]\|] \leq \mathbb{E}_z[\mathbb{E}_p[|b_p|]] = \mathbb{E}_p[\mathbb{E}_z[|b_p(x_p)|]] = O((n^{-1} \log n)^{\frac{3}{7}}). \quad (\text{C.34})$$

The penultimate equality above involves swapping the order of expectations, which is possible by Tonelli's theorem (since $|b_p| > 0$). Then using triangle inequality, it follows that $\mathbb{E}_z[\|\hat{\phi}_p - \phi_p\|_{L_\infty[-1,1]}] = O((n^{-1} \log n)^{\frac{3}{7}})$.

2. $(l, l') \in \mathcal{S}_2$.

We only consider the case where $\rho(l), \rho(l') > 1$ as proofs for the cases are similar. For $\tilde{\phi}_{(l,l')}$, we have that $\mathbb{E}_z[\|\tilde{\phi}_{(l,l')} - (g_{(l,l')} + C)\|_{L_\infty[-1,1]^2}] = O((n^{-1} \log n)^{\frac{3}{8}})$. Denoting $\tilde{\phi}_{(l,l')}(x_l, x_{l'}) - (g_{(l,l')}(x_l, x_{l'}) + C) = b_{(l,l')}(x_l, x_{l'})$, this means $\mathbb{E}_z[|b_{(l,l')}(x_l, x_{l'})|] = O((n^{-1} \log n)^{\frac{3}{8}}), \forall (x_l, x_{l'}) \in [-1, 1]^2$. Using Tonelli's theorem as earlier, one can next verify that:

$$\mathbb{E}_z[\|\mathbb{E}_l[\tilde{\phi}_{(l,l')}] - (\mathbb{E}_l[g_{(l,l')} + C])\|_{L_\infty[-1,1]}] = O((n^{-1} \log n)^{\frac{3}{8}}), \quad (\text{C.35})$$

$$\mathbb{E}_z[\|\mathbb{E}_{l'}[\tilde{\phi}_{(l,l')}] - (\mathbb{E}_{l'}[g_{(l,l')} + C])\|_{L_\infty[-1,1]}] = O((n^{-1} \log n)^{\frac{3}{8}}), \quad (\text{C.36})$$

$$\mathbb{E}_z[\|\mathbb{E}_{(l,l')}[\tilde{\phi}_{(l,l')}] - (\mathbb{E}_{(l,l')}[g_{(l,l')} + C])\|] = O((n^{-1} \log n)^{\frac{3}{8}}). \quad (\text{C.37})$$

As in the proof of Proposition 1, we obtain from (C.35), (C.36), (C.37), (C.22), (C.7) (via triangle inequality):

$$\mathbb{E}_z[\|\hat{\phi}_{(l,l')} - \phi_{(l,l')}\|_{L_\infty[-1,1]^2}] = O((n^{-1} \log n)^{\frac{3}{8}}). \quad (\text{C.38})$$

¹²See [39] for $d = 1$, and [41] for $d \geq 1$

3. $\mathbf{l} \in \mathcal{S}_2^{\text{var}} : \rho(\mathbf{l}) > \mathbf{1}$.

In this case, for $\tilde{\phi}_l : [-1, 1]^2 \rightarrow \mathbb{R}$, we have that $\mathbb{E}_z[\|\tilde{\phi}_l - (g_l + C)\|_{L_\infty[-1,1]^2}] = O((n^{-1} \log n)^{\frac{3}{8}})$, with $g_l(x_l, x)$ as defined in (C.24). Using Tonelli's theorem as earlier, one can verify that

$$\mathbb{E}_z[\|\mathbb{E}_x[\tilde{\phi}_l] - (\mathbb{E}_x[g_l] + C)\|_{L_\infty[-1,1]}] = O((n^{-1} \log n)^{\frac{3}{8}}), \quad (\text{C.39})$$

$$\mathbb{E}_z[\|\mathbb{E}_{(l,x)}[\tilde{\phi}_l] - (\mathbb{E}_{(l,x)}[g_l] + C)\|] = O((n^{-1} \log n)^{\frac{3}{8}}). \quad (\text{C.40})$$

Then using the fact $\mathbb{E}_x[g_l(x_l, x)] - \mathbb{E}_{(l,x)}[g_l(x_l, x)] = \phi_l(x_l)$, we obtain via triangle inequality the bound: $\mathbb{E}_z[\|\hat{\phi}_l - \phi_l\|_{L_\infty[-1,1]}] = O((n^{-1} \log n)^{\frac{3}{8}})$. This completes the proof. □