
Lifting high-dimensional non-linear models with Gaussian regressors

Christos Thrampoulidis
University of California, Santa Barbara

Ankit Singh Rawat*
Massachusetts Institute of Technology

Abstract

We study the problem of recovering a structured signal \mathbf{x}_0 from high-dimensional data $\mathbf{y}_i = f(\mathbf{a}_i^T \mathbf{x}_0)$ for some nonlinear (and potentially unknown) link function f , when the regressors \mathbf{a}_i are iid Gaussian. Brillinger (1982) showed that ordinary least-squares estimates \mathbf{x}_0 up to a constant of proportionality μ_ℓ , which depends on f . Recently, Plan & Vershynin (2015) extended this result to the high-dimensional setting deriving sharp error bounds for the generalized Lasso. Unfortunately, both least-squares and the Lasso fail to recover \mathbf{x}_0 when $\mu_\ell = 0$. For example, this includes all even link functions. We resolve this issue by proposing and analyzing an alternative convex recovery method. In a nutshell, our method treats such link functions as if they were linear in a lifted space of higher-dimension. Interestingly, our error analysis captures the effect of both the nonlinearity and the problem’s geometry in a few simple summary parameters.

1 Introduction

We consider the problem of estimating an unknown signal $\mathbf{x}_0 \in \mathbb{R}^n$ from a vector $\mathbf{y} = (y_1, y_2, \dots, y_m)^T$ of m *generalized linear measurements* of the following form:

$$\mathbf{y}_i = f_i(\mathbf{a}_i^T \mathbf{x}_0), \quad i = 1, 2, \dots, m. \quad (1)$$

Here, each $\mathbf{a}_i \in \mathbb{R}^n$ represents a (known) measurement vector and the f_i ’s are independent copies of a (possibly random) link function f . A few examples include: $f_i(x) = x + z_i$, with say z_i being random noise, for standard linear regression setup; $f_i(x) = |x|^2 + z_i$, for quadratic (noisy) measurements; $f_i(x) = Q_b(|x|)$, where $Q_b(\cdot)$ denotes a quantizer on b number of bits. In the statistics and econometrics literature, (1) is popular under

the name *single-index model* (also regarded as a special case of the *sufficient dimension reduction* problem). As a slight generalization of (1), we allow measurements that are drawn independently according to a conditional distribution of some probability density function as follows:

$$y_i \sim p(y \mid \mathbf{a}_i^T \mathbf{x}_0), \quad i = 1, 2, \dots, m. \quad (2)$$

This expands the model by including instances such as logistic-regression and $y_i \sim \text{Poisson}(|\mathbf{a}_i^T \mathbf{x}_0|)$, $i \in [m]$.

We seek recovery methods that ensure the following favorable features: (i) computational efficiency; (ii) provable performance guarantees; (iii) flexibility to exploit various forms of possible prior structural knowledge about \mathbf{x}_0 (e.g., sparsity); and (iv) flexibility to the underlying link-function. Here, “flexible” refers to a method that can easily adapt to different structural information on \mathbf{x}_0 and different link-functions with minimal changes, such as appropriate tuning of involved regularization parameters. Lastly, we remark that information about the magnitude of \mathbf{x}_0 might in general be lost in the nonlinearity. Thus, for partial identifiability, we assume throughout that $\|\mathbf{x}_0\|_2 = 1$ and aim to recover an estimate that is highly correlated with the true signal (often referred to as *weak recovery*, e.g., [17]).

1.1 Prior art

In the simplest case with *linear* link function, i.e., $f_i(x) = x + z_i$, perhaps the most popular approach of estimating \mathbf{x}_0 is via solving the generalized Lasso:

$$\hat{\mathbf{x}} := \arg \min_{\mathbf{x}} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{a}_i^T \mathbf{x}_0)^2 \quad \text{s.t.} \quad \mathbf{x} \in \mathcal{K}_{\mathcal{R}}. \quad (3)$$

Here, for a properly chosen *regularizer* function $\mathcal{R} : \mathbb{R}^n \rightarrow \mathbb{R}$, $\mathcal{K}_{\mathcal{R}} \subset \mathbb{R}^n$ is a set that encodes the available information about \mathbf{x}_0 . For instance, $\mathcal{K}_{\mathcal{R}} = \{\mathbf{x} \in \mathbb{R}^n \mid \mathcal{R}(\mathbf{x}) \leq K\}$, where $K > 0$ is a (tuning) parameter. The generalized Lasso comes with provable performance guarantees under general assumptions on the choice of \mathcal{R} and on the measurement vectors. The Lasso objective is by nature tailored to linear measurement models, but one can always employ it as a candidate recovery algorithm for the more general model (1). This immediately raises the following question: *For such non-linear measurements, when (if ever) is the solution $\hat{\mathbf{x}}$ of the Lasso still a good estimate of \mathbf{x}_0 ?*

*Ankit Singh Rawat is now with Google Research, NY. Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

This question has been recently addressed in a quantitative way in [22], under the assumption that the measurement vectors are independent Gaussians. Naturally, the answer depends on both: (i) the specific non-linearity f in (1); and (ii) the structure of \mathbf{x}_0 and the associated choice of the regularizer. The dependence on these factors can be conveniently summarized in terms of a few easily computable model parameters. In particular, the effect of the non-linearity is entirely captured by the following:

$$\mu_\ell := \mathbb{E}[\gamma f(\gamma)] \quad \text{and} \quad \tau_\ell^2 := \mathbb{E}[(f(\gamma) - \mu_\ell \gamma)^2], \quad (4)$$

where, the expectations are over $\gamma \sim \mathcal{N}(0, 1)$ and the (possibly) random link function f . For simplicity, we focus on the case of sparse recovery, but the results of [22] are more generally applicable. When, \mathbf{x}_0 is k -sparse and $\hat{\mathbf{x}}$ is the solution to (3) with ℓ_1 -regularization, then [22, Thm. 1.4] shows the following. Provided that $m \gtrsim k \log(n/k)$ and n is large enough, it holds with high probability that¹:

$$\|\hat{\mathbf{x}} - \mu_\ell \cdot \mathbf{x}_0\|_2 \lesssim \tau_\ell \cdot \sqrt{k \log(n/k)} / \sqrt{m}. \quad (5)$$

What if the Lasso fails? It becomes clear from this discussion that (at least for Gaussian regressors) the generalized Lasso satisfies all the four aforementioned favorable properties that we seek in our recovery methods. However, a closer inspection of (5) reveals that the Lasso fails to produce a good estimate for all functions for which $\mu_\ell = 0$; e.g., this includes all *even* link functions!

1.2 Our contribution

In this paper, we provide an affirmative answer to the following natural questions that arise from the inadequacy of the Lasso for a large class of link functions: *Is there a generic convex program that can recover structured signals from even non-linear measurements? And if so, can we quantify its error performance?*

As we discuss next, our algorithm is motivated by a simple observation. It is well known that one can expand the link function f as a series in terms of the Hermite polynomials, i.e., $f(x) = \sum_{i=0}^{+\infty} \mu_i H_i(x)$, where $H_i(x)$ is the i^{th} -order Hermite polynomial with leading coefficient 1, and, $\mu_i = \frac{1}{n!} \mathbb{E}_\gamma[f(\gamma) H_i(\gamma)]$ for $\gamma \sim \mathcal{N}(0, 1)$. In particular, we may expand $y_i = f(\mathbf{a}_i^T \mathbf{x}_0)$, $i \in [m]$ as follows:

$$y_i = \mu_0 + \mu_1(\mathbf{a}_i^T \mathbf{x}_0) + \mu_2((\mathbf{a}_i^T \mathbf{x}_0)^2 - 1) + \dots \quad (6)$$

with $\mu_0 = \mathbb{E}[f(\gamma)]$, $\mu_1 = \mathbb{E}[\gamma f(\gamma)]$, and $\mu_2 = \frac{1}{2} \mathbb{E}[(\gamma^2 - 1)f(\gamma)]$. First, observe that $\mu_1 = \mu_\ell$ in (4); thus, we may interpret the objective function of the Lasso in (3) as one

¹Here and in the rest of the paper, a statement is said to hold with high probability if it holds with probability at least 0.99 (say). Also, the symbol “ \lesssim ” is used to hide universal constants (in particular ones that do not depend on f).

that attempts to approximate (in ℓ_2 -sense) each y_i in (6) by only keeping the first-order (linear) term. Clearly, this approximation fails if $\mu_1 = 0$, and so does the Lasso. For such cases, we naturally propose keeping only the second-order (quadratic) term of the expansion in (6). Moreover, in order to obtain a favorable convex program, we apply the *lifting* technique. In particular, approximating (6), we write

$$\begin{aligned} y_i &\approx \mu_2((\mathbf{a}_i^T \mathbf{x}_0)^2 - 1) = \text{tr}((\mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}) \cdot \mu_2 \mathbf{x}_0 \mathbf{x}_0^T) \\ &= \text{tr}((\mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}) \cdot \tilde{X}_0), \end{aligned} \quad (7)$$

where \tilde{X}_0 denotes the rank-1 matrix $\mu_2 X_0 = \mu_2 \mathbf{x}_0 \mathbf{x}_0^T$, and the first equality follows as $\|\mathbf{x}_0\|_2 = 1$. Note that (after lifting) the quantity on the right-hand side (RHS) in (7) is now a linear function of the unknown rank-1 matrix \tilde{X}_0 .

We are now ready to describe our algorithm. From (7), one can attempt to reconstruct \tilde{X}_0 (i.e., a scaled version of $X_0 = \mathbf{x}_0 \mathbf{x}_0^T$) by searching for a positive-semidefinite and low-rank matrix X that minimizes the residual between y_i and $\text{tr}((\mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}) \cdot X)$. Further relaxing the rank constraint leads to the following convex method:

$$\begin{aligned} \hat{X} &= \arg \min_{X \succeq 0} \sum_{i=1}^m (y_i - \text{tr}((\mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}) \cdot X))^2 + \lambda \cdot \text{tr}(X), \\ &\text{subject to } X \in \mathcal{K}_{\mathcal{R}}, \end{aligned} \quad (8)$$

where $\lambda > 0$ is a regularization parameter. We have further added a constraint $X \in \mathcal{K}_{\mathcal{R}}$ to promote the structure of the signal-defined matrix X_0 (inherited by the structure of \mathbf{x}_0). For instance, if \mathbf{x}_0 is k -sparse, then X_0 is k^2 -sparse and a natural choice becomes $X \in \mathcal{K}_{\ell_1} := \{W \mid \|W\|_1 \leq K\}$, for $\|W\|_1 = \sum_{i,j=1}^n |W_{ij}|$ and regularizer $K > 0$. As a last step, our algorithm obtains a final estimate $\hat{\mathbf{x}}$ of (a scaled) \mathbf{x}_0 by the leading eigenvector of \hat{X} .

Error guarantees. We characterize the estimation performance of the following constrained version of (8):

$$\begin{aligned} \hat{X} &= \arg \min_{X \succeq 0} \sum_{i=1}^m (y_i - \text{tr}((\mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}) \cdot X))^2 \\ &\text{subject to } \text{tr}(X) \leq \tilde{\mu} \quad \text{and} \quad X \in \mathcal{K}_{\mathcal{R}}, \end{aligned} \quad (9)$$

where $\tilde{\mu} > 0$ is a tuning parameter. Here, we state a version of our result in the ideal case where $\tilde{\mu} = \mu_2$ (cf. (6) and recall that the goal of (9) is to recover $\tilde{X} = \mu_2 X_0$ for which $\text{tr}(\tilde{X}) = \mu_2$). In Sec. 3.4, we also analyze the performance of (9) when $\tilde{\mu}$ is not ideally tuned. Our result captures the effect of f via two simple parameters. For $\gamma \sim \mathcal{N}(0, 1)$:

$$\begin{aligned} \mu_q &:= \mu_2 = \frac{1}{2} \mathbb{E}[(\gamma^2 - 1)f(\gamma)] \quad \text{and} \\ \tau_q^2 &:= \mathbb{E}[(f(\gamma) - \mu \cdot (\gamma^2 - 1))^2]. \end{aligned} \quad (10)$$

For example, for recovery of a k -sparse signal we show the following about (9) with ℓ_1 -regularization (Thm. 2.1 for

details): For $m \gtrsim k^2 \log(n/k)$ and n large enough, it holds with high probability that

$$\|\hat{X} - \mu_q X_0\|_F \lesssim \tau_q \cdot k \log(n/k) / \sqrt{m}. \quad (11)$$

From eigenvalue perturbation theory, this further bounds the deviation of the leading eigenvector $\hat{\mathbf{x}}$ of \hat{X} from \mathbf{x}_0 (cf. Sec. 2.1). Observe the resemblance between (5) and our result in (11). The Lasso-related parameters μ_ℓ and τ_ℓ are replaced by μ_q and τ_q , respectively. It is evident from (11) that (9) is efficient for those link functions f for which μ_q is nonzero. In contrast with (5), in terms of sample complexity, (11) pays a penalty of $\mathcal{O}(k^2 \log(n/k))$ rather than $\mathcal{O}(k \log(n/k))$. This is not surprising since the same gap holds for the known algorithms for quadratic measurements (e.g., [20]), a special case of even link function.

We note that (11) is only an instantiation to sparse signals of our general result that characterizes the recovery performance of (9) for general convex regularizers \mathcal{R} (cf. Sec. 2.2). Also, inherent in the formulation of (9) is the condition $\mu_q > 0$ (since, $X \succeq 0$ implies that $\text{tr}(X) \geq 0$). Note that if $\mu_q < 0$, then the same method works only by replacing the constraint $X \succeq 0$ with $X \preceq 0$. Overall, the proposed algorithm requires that $\mu_q \neq 0$.

Remark 1 (Extensions). Our algorithm is motivated by the expansion of the link function in (6), which we attempt to approximate in (8) by solely keeping the second-order term. Naturally, one can imagine the possibility of developing extensions to (8) that achieve better performance by more accurate approximations in (6). For example, one may keep both the zeroth- and the second-order terms:

$$\begin{aligned} \hat{X} &= \arg \min_{X \succeq 0} \sum_{i=1}^m (\mathbf{y}_i - \hat{\mu}_0 - \text{tr}((\mathbf{a}_i \mathbf{a}_i^T - \mathbf{I}) \cdot X))^2 \\ &\text{subject to } \text{tr}(X) \leq \tilde{\mu} \text{ and } X \in \mathcal{K}_{\mathcal{R}}, \end{aligned} \quad (12)$$

where, $\hat{\mu}_0 = \frac{1}{m} \sum_{i=1}^m y_i$. Note that for large enough number of measurements m , it holds $\hat{\mu}_0 \approx \mathbb{E}[f(\gamma)] = \mu_0$. We remark that our analysis of (9) directly translates to guarantees about (12) when $\hat{\mu}_0$ is estimated by a fresh batch of measurements (for example, this is achieved in practice by sample splitting). Other interesting extensions of this flavor are certainly possible, albeit may require additional effort.

Remark 2 (Relation to PhaseLift). The lifting technique is by now well established in the literature, the most prominent example being its use in phase retrieval (i.e., quadratic measurements). In fact, the popular PhaseLift method [5] is very similar to (8), the only difference being in the loss function: the PhaseLift penalizes $\sum_{i=1}^m (\mathbf{y}_i - \text{tr}(\mathbf{a}_i \mathbf{a}_i^T \cdot X))^2$, instead. In reference to (6), note that in the special case of quadratic measurements, it holds $\mu_0 = \mu_2 = 1$, which explains the choice of that particular loss function in the PhaseLift. However, our arguments show the need for modifications to be able to treat a wide class of functions beyond quadratics. This leads to (12), which can

be viewed as a canonical extension or robust version of the PhaseLift to more general link-functions.

1.3 Relevant literature and outlook

The vast majority of the literature on structured signal recovery in the high-dimensional regime assumes a linear measurement model. In this vanilla setting, convex methods offer many desired features: computational efficiency, tractable analysis, flexibility to adjust to different problem instances (such as, sparse recovery with ℓ_1 -regularization and low-rank matrix recovery with nuclear-norm minimization). By now, performance guarantees for these methods are well-established ([9] and references therein).

Extensions to nonlinear link functions, have been only more recently considered in high dimensions. While, the key observation that for Gaussian measurement vectors it holds $\mu_\ell \mathbf{x}_0 = \arg \min_{\mathbf{x}} \mathbb{E}(f(\langle \mathbf{a}, \mathbf{x}_0 \rangle) - \langle \mathbf{a}, \mathbf{x} \rangle)^2$ goes back to [4] (also, [14]), Plan and Vershynin [22, 23] were the first to show how that idea extends to the high-dimensional setting by applying it to the generalized Lasso. Subsequently [29] extended [22] to the regularized Lasso, obtained the exact constants in the analysis, and used the results to design optimal quantization schemes. More recent works involve extensions to other loss-functions beyond least squares [10], to elliptically symmetric distributions [11], to projected gradient-descent [21], to signal-demixing applications [24], and to demonstrating computational speedups compared to maximum-likelihood estimation [8]. Finally, Yang et al. [32] have appropriately modified Brillinger’s original observation to sub-Gaussian vectors, based on which they propose and analyze generic convex solvers that work in this more general setting.

Unfortunately, all these works, starting with the original result by Brillinger, assume that the link function satisfies $\mu_\ell \neq 0$. Instead, our method and analysis works for these; and more generally for link functions satisfying $\mu_q \neq 0$. In that sense, our paper is a direct counterpart of [22] for “even-like” nonlinearities. Given the aforementioned extensions that followed [22], it is natural to expect analogous extensions of our results as part of future research. Also, perhaps an interesting research question that is raised is related to unifying the efficacy of the Lasso and of (8) in a single generic algorithm that would combine the best of the two worlds and would apply to nonlinearities satisfying either $\mu_\ell \neq 0$ or $\mu_q \neq 0$ (cf. Remark 1). On the one hand, since $\mu_\ell \neq 0$, one can obtain an estimate of \mathbf{x}_0 by the Lasso; however, this entirely ignores the hidden “quadratic part” in the measurements. On the other hand, seemingly natural lifting procedure leads to a semidefinite-optimization estimator that successfully accounts for both “linear and quadratic parts”, but suffers from worse sample complexity. We point out the interesting and relevant work by Yi et al. [35] in this direction. Their algorithm is appli-

cable to such a wide class of link-functions, but, it imposes other restrictions compared to our work, such as requiring binary measurements and a sparse signal \mathbf{x}_0 .

Out of all the nonlinear link functions, the quadratic case (cf. phase-retrieval) deserves special attention; there has been a surge in the provable recovery methods for this case over the past few years. Among the most well-established and the first chronologically to enjoy rigorous recovery guarantees are the methods based on semidefinite relaxation [3, 5]. Such methods operate by lifting the original n -dimensional natural parameter space to a higher dimensional matrix space. Here, we extend the lifting idea and the recovery guarantees to general link functions beyond quadratics. Naturally, the increase in dimensionality often introduces challenges in computational complexity and memory requirements. To overcome these issues, subsequent works on phase retrieval develop non-convex formulations that start with a careful spectral initialization (see [25, Sec. 6] for references), which is then iteratively refined by a gradient-descent-like scheme of low computational complexity. See also [2, 12] for an alternative *convex* formulation of the problem in the natural parameter space. Naturally, many of these solution methods can be combined with the ideas introduced in this paper to extend their reach to non-quadratic link functions. In fact, while preparing our paper, we became aware of the recent work by Yang et al. [34] exactly along these lines, which proposes an iterative non-convex method for sparse-signal recovery from link functions beyond quadratics. Interestingly, in order to arrive at their algorithm, Yang et al. identify necessary modifications to the vanilla non-convex methods for quadratics, which are in exact agreement with our modification to the PhaseLift needed to arrive at (8). Thus, (8) and (12) can be interpreted as the convex counterparts to the algorithm of [34]. Despite the similarities, (a) the two papers arrive at the proposed algorithms through different ideas (cf. Remark 2); (b) our algorithm and analysis support structures beyond sparsity; (c) our performance bounds precisely capture the effect of the link-function.

Another closely related work [33] takes a somewhat different view than ours on the measurement model in (1) and arrives at a different semidefinite optimization algorithm closer to the algorithms for sparse PCA. Notably, [33] also addresses the design matrices with sub-gaussian entries. For Gaussian regressors, a parallel work to ours [28] extends the analysis of [33] beyond sparse recovery. Finally, [18, 15, 17] study the performance of spectral initialization for measurements as in (1). In contrast to our work, the results in [15, 17] are asymptotic and do not exploit structural information on \mathbf{x}_0 . Also, compared to [18, 33, 28] our analysis is sharp with respect to the link function f . Such sharp results are directly useful in various applications where one has control over some parameters of f , and also, to the design of pre-processing functions h (see [29, 15, 17] for pre-

cursors of this idea). For example, by inspecting (11), the estimation error is minimized by choosing h so that μ_q and τ_q result in the effective noise parameter τ_q/μ_q being as small as possible (see Sec. 4). We postpone further investigations of such implications to future work.

Notation. We use \mathbb{S}_n (\mathbb{S}_n^+) to denote the sets of real $n \times n$ symmetric (resp., positive semidefinite) matrices. For $U, V \in \mathbb{S}_n$, $\langle U, V \rangle = \text{tr}(UV)$ denotes the standard inner product in \mathbb{S}_n , $\|V\|_2$ ($\|V\|_F$) denotes the spectral (resp., Frobenius) norm, and $\|V\|_1 = \sum_{j \in [m]} \sum_{i \in [n]} |V_{ji}|$. For $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_p$ denotes its ℓ_p norm, $p \geq 1$. We denote by $\mathcal{S}_{F,n}$ the unit sphere in \mathbb{S}_n , i.e., $\mathcal{S}_{F,n} := \{V \in \mathbb{S}_n \mid \|V\|_F = 1\}$.

Organization. We formally state our results along with some technical background in Sec. 2. The case of sparse recovery is formally treated in Sec. 2.1, while error guarantees for the general setting are provided in Sec. 2.2. We highlight the key steps of the proofs in Sec. 3 (full details are relegated to the supplementary material). Also, in Sec. 3.4, we study the performance of (9) under imperfect tuning. Numerical simulations are presented in Sec. 4.

2 Error bounds for nonlinear measurements

2.1 An example: Sparse recovery

Assume that the true signal \mathbf{x}_0 is k -sparse. Then, $X_0 = \mathbf{x}_0 \mathbf{x}_0^T$ is at most k^2 -sparse. Thus, we solve (9) with an ℓ_1 -norm constraint. Thm. 2.1 below characterizes the performance of the algorithm. Before that, recall the definitions of μ_q and τ_q in (10). Here onwards, for ease of exposition, we refer to μ_q and τ_q simply as μ and τ , respectively. If f is twice differentiable with f'' being its second derivative, then by integration by parts we have $\mu = \mathbb{E}[f''(\gamma)]/2$. Moreover, it is easily shown that $\tau^2 = \mathbb{E}[f^2(\gamma)] - 2\mu^2$.

For all the theorems that follow: *a statement is said to hold with high probability if it holds with probability at least 0.99 (say)*. Also, the appearing constants $c, C > 0$ may only depend on the probability of success.

Theorem 2.1 (Sparse recovery). *Suppose that \mathbf{x}_0 is k -sparse, $\mathbf{a}_i \sim \mathcal{N}(0, \mathbf{I})$, and that \mathbf{y} follows the generalized linear model of (1). Assume that $\mu > 0$, and let \hat{X} be the solution to (9) with $\tilde{\mu} = \mu$ and $\mathcal{K}_{\mathcal{R}} = \{X : \|X\|_1 \leq \mu \|X_0\|_1\}$. There exist universal constants $c, C > 0$ such that, if the number of observations obeys*

$$m \geq c \cdot k^2 \log(n/k) \quad (13)$$

for sufficiently large n , then, with high probability, we have

$$\|\hat{X} - \mu X_0\|_F \leq C \cdot \tau \cdot k \log(n/k) / \sqrt{m}. \quad (14)$$

We defer the proof of Theorem 2.1 to Sec. A.1 in the supplementary. The theorem does not claim that \hat{X} is rank one.

As usual, we obtain an estimate $\hat{\mathbf{x}}$ of \mathbf{x}_0 by extracting the rank-one component (e.g. [5]). In particular, if λ_1 and \mathbf{v}_1 denote the maximum eigenvalue and the principal eigenvector of \hat{X} , respectively, then we obtain $\hat{\mathbf{x}} = \sqrt{\lambda_1} \mathbf{v}_1$. It follows from (19) that $\|\hat{\mathbf{x}} - \sqrt{\mu} \mathbf{x}_0\|_2 \leq C \cdot \min\{\sqrt{\mu}, \frac{E}{\sqrt{\mu}}\}$, where E denotes the expression in the RHS of (14). The proof is based on the Davis-Kahan $\sin(\theta)$ -theorem and is the same as in [5, Sec. 6]; thus it is omitted for brevity.

Thm. 2.1 implies the sample complexity of $\mathcal{O}(k^2 \log(n/k))$. Notably, for quadratic measurements, this is the same as the guarantees of [5]. In fact, the same k^2 -barrier appears in most of the algorithms that have been proposed for sparse recovery from quadratic measurements (e.g., [25, Sec. 6] and references therein). However, given more information about the link-function (e.g., f having sub-exponential moments), the numerator of the RHS in (14) can be improved to $k^2 \sqrt{\log(n/k)}$.

2.2 General result

Thm. 2.2 below characterizes the performance of (9) for general signal structure and $\mathcal{K}_{\mathcal{R}}$. The bounds are given in terms of specific summary parameters. We distinguish between: (i) *Geometric parameters* that capture the efficacy of the imposed geometric constraints in (9) in promoting solutions of desired structure (positive semidefinite, low-rank, sparse, etc.); (ii) *Model parameters* that capture f .

2.2.1 Geometric parameters

First, we introduce the notions of the tangent cone and the (local) Gaussian width.

Definition 1 (Tangent cone). *The tangent cone of a subset $\mathcal{K} \subset \mathbb{S}_n$ at $X \in \mathbb{S}_n$ is defined as $\mathcal{D}(\mathcal{K}, X) := \{\tau V : \tau \geq 0, V \in \mathcal{K} - X\}$.*

Definition 2 (Local) Gaussian width). *The local Gaussian width of a set $\mathcal{C} \subset \mathbb{S}_n$ is a function of a scale parameter $t > 0$, which is defined as*

$$\omega_{g,t}(\mathcal{C}) := \mathbb{E}_G \left[\sup_{V \in \mathcal{C} \cap t\mathcal{S}_{F,n}} \langle G, V \rangle \right], \quad (15)$$

where G is a matrix from the Gaussian orthogonal ensemble (GOE), i.e. $G = G^T$, $G_{ii} \stackrel{iid}{\sim} \mathcal{N}(0, 1)$ for $i \in [n]$, and $G_{ij} \stackrel{iid}{\sim} \mathcal{N}(0, 1/2)$ for $i > j \in [n]$.

For $t = 1$, we refer to $\omega_{g,1}(\mathcal{C})$ simply as the Gaussian width, and, we use the shorthand $\omega_g(\mathcal{C})$.

The results of this section only involve the Gaussian width $\omega_{g,1}(\mathcal{C}) = \omega_g(\mathcal{C})$ ($t = 1$, above). The general definition of local Gaussian width becomes useful when we study the performance of (9) under imperfect parameter tuning in Sec. 3.4. The Gaussian width plays a central role in asymptotic convex geometry. It also appears as a key quantity in the study of random linear inverse problems [7, 1]: for a

cone $\mathcal{C} \subset \mathbb{S}_n$, $\omega_g(\mathcal{C})^2$ can be formally described as a measure of the effective dimension of the cone \mathcal{C} [31, 1]. Importantly, while it is an abstract geometric quantity, it is possible in many instances to derive sharp numerical bounds that are explicit in terms of the parameters of interest (such as sparsity level, rank) [26, 7, 1]. We make use of these ideas in the proof of Thm. 2.2 (see Sec. 3.3.2).

Finally, we need two more geometric parameters: Talagrand's γ_1 and γ_2 -functionals. To streamline the presentation, we defer the formal definitions of these parameters to Sec. C in the supplementary. For a set $\mathcal{K} \subset \mathbb{S}_n$ we write $\gamma_1(\mathcal{K}, \|\cdot\|_2)$ and $\gamma_2(\mathcal{K}, \|\cdot\|_F)$ for the γ_1 and γ_2 -functionals with respect to the spectral and the Frobenius norms, respectively. The γ -functionals are fundamental in the study of suprema of random processes and specifically in the theory of generic chaining [27]. In general, explicit calculation of the γ -functionals can be challenging depending on the specific set \mathcal{K} ; however, it is often possible to control them in a sufficient way. Specifically, for the term $\gamma_2(\mathcal{C}; \|\cdot\|_F)$, one can appeal to Talagrand's majorizing measure theorem that establishes a tight (up to constants) relations to the Gaussian width [27, Thm. 2.1.1], which can in turn be often well approximated: $\gamma_2(\mathcal{K}; \|\cdot\|_F) \leq C \cdot \omega_g(\mathcal{K})$. More generally, Dudley's integral produces bounds on γ_1, γ_2 in terms of the metric entropy of the set (see (66) in the supplementary)².

2.2.2 Model parameters

First, recall the definition of μ and τ in (10). Moreover, for $\gamma \sim \mathcal{N}(0, 1)$ and all expectations take over γ and f :

$$\begin{aligned} v^2 &:= \mathbb{E}[\gamma^2 \cdot (f(\gamma) - \mu \cdot (\gamma^2 - 1))^2] \quad \text{and} \\ \chi^2 &:= \mathbb{E}[(\gamma^2 - 1)^2 \cdot (f(\gamma) - \mu \cdot (\gamma^2 - 1))^2]. \end{aligned} \quad (16)$$

Remark 3. The results extend to the more general model of (2) with a natural modification in (10) and (16). In particular, $f(\gamma)$ is substituted by a (random variable) $y \sim p(y|\gamma)$, and the expectation is over the conditional distribution $p(y|\gamma)$ and γ , e.g., $\mu = \frac{1}{2} \int y \mathbb{E}_\gamma[(\gamma^2 - 1)p(y|\gamma)] dy$.

2.2.3 Main result

We are now ready to state the main result of this section.

Theorem 2.2 (General result). *Suppose that $\mathbf{a}_i \sim \mathcal{N}(0, \mathbf{I})$, and that \mathbf{y} follows the model in (1). Recall the definitions of μ, τ, v, χ in (10) and (16). Assume that $\mu > 0$ and that $\tilde{\mu} = \mu$ and $\mu X_0 \in \mathcal{K}_{\mathcal{R}}$ in (9), where $X_0 = \mathbf{x}_0 \mathbf{x}_0^T$. Denote*

$$\Gamma := \min \left\{ \sqrt{n}, \gamma_2(\mathcal{D}(\mathcal{K}_{\mathcal{R}}, \mu X_0) \cap \mathcal{S}_{F,n}, \|\cdot\|_F) + \gamma_1(\mathcal{D}(\mathcal{K}_{\mathcal{R}}, \mu X_0) \cap \mathcal{S}_{F,n}, \|\cdot\|_2) \right\}. \quad (17)$$

²See also [19] for recent progress on bounding the γ_1 functional of constraint sets with ‘‘good’’ covering number in terms of their Gaussian width. In particular, [19, Lemma D.19] can be combined with Thm. 2.2 to obtain specific results for other structures beyond sparsity.

There exist universal constants $c, C > 0$ such that, if

$$m \geq c \cdot \Gamma^2, \quad (18)$$

then, with high-probability, the solution \widehat{X} of (9) satisfies

$$\|\widehat{X} - \mu X_0\|_F \leq \frac{C}{\sqrt{m}} \cdot \left(\tau \cdot \Gamma + \chi + v \cdot \min \left\{ \sqrt{n}, \omega_g(\mathcal{D}(\mathcal{K}_{\mathcal{R}}, \mu X_0)) \right\} \right), \quad (19)$$

Note that Thm. 2.2 holds for all values of n . Typically, for large enough n , the first term in the right hand side (RHS) of (19), i.e., $\tau \cdot \Gamma / \sqrt{m}$, becomes the dominant term. For a simple illustration of the theorem, consider the case of a generic true signal \mathbf{x}_0 without any prior structural information. In this case, we solve (9) with no additional constraints other than $X \succeq 0$ and $\text{tr}(X) \leq \mu$. Hence, $\Gamma \leq \sqrt{n}$ and from (18) the sample requirement is $m \geq c' \cdot n$.

3 Technical results and proofs

We now outline the proof of Thm. 2.2, which has two main steps. First, in Thm. 3.1, we upper bound the error $\|\widehat{X} - \mu X_0\|_F$ of (9) in terms of an appropriate geometric quantity, namely the *weighted empirical width*. We state the theorem and outline its proof in Sec. 3.1 and 3.2, respectively. Next, in Sec. 3.3, we show how to control the weighted empirical width to finally arrive at Thm. 2.2.

3.1 General upper bound

We begin the section by defining the local weighted empirical width. Similar to the definition of the local Gaussian width, this geometric is also a function of a scale parameter t . In this section, we only make use of the case $t = 1$, but the generality of the definition will prove useful in Sec. 3.4.

Definition 3 (Local) empirical width). Let $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^n$ be independent copies of a standard normal vector $\mathcal{N}(0, \mathbf{I}_n)$ and $\varepsilon_1, \dots, \varepsilon_m$ be independent Rademacher random variables. For a set $\mathcal{C} \subset \mathbb{S}_n$ and a vector $\mathbf{p} := (p_1, \dots, p_m)$ the local weighted empirical width $\omega_{e,t}(\mathcal{C}; \mathbf{p})$ is a function of a scale parameter $t > 0$ defined as follows:

$$\omega_{e,t}(\mathcal{C}; \mathbf{p}) := \mathbb{E} \left[\sup_{V \in \mathcal{C} \cap t\mathcal{S}_{F,n}} \langle V, H_{\mathbf{p}} \rangle \right], \quad (20)$$

where $H_{\mathbf{p}} := \frac{1}{\sqrt{m}} \sum_{i=1}^m p_i \cdot \varepsilon_i \cdot \mathbf{a}_i^T$, and, the expectation is over the randomness of $\{\mathbf{a}_i\}$ and of $\{\varepsilon_i\}$. In particular, when $\mathbf{p} = \mathbf{1}$ we write $\omega_{e,t}(\mathcal{C}) := \omega_{e,t}(\mathcal{C}; \mathbf{1})$, and call this the local empirical width.

If $t = 1$, we simply call $\omega_{e,1}(\mathcal{C}; \mathbf{p})$ the weighted empirical width, and, we use the shorthand $\omega_e(\mathcal{C}; \mathbf{p})$.

We study further the weighted empirical width in Sec. 3.3. Now, we are ready to state a general upper bound on the

error of the estimate obtained by (9). For convenience, let:

$$\mathcal{K}_0 := \{X : X \succeq 0 \text{ and } \text{tr}(X) \leq \tilde{\mu}\}. \quad (21)$$

Theorem 3.1. Let the same assumptions as in Thm. 2.2 hold (including $\tilde{\mu} = \mu$). Further let

$$\mathcal{C}_0 := \mathcal{D}(\mathcal{K}_0, \mu X_0) \cap \mathcal{D}(\mathcal{K}_{\mathcal{R}}, \mu X_0). \quad (22)$$

Finally, define $\boldsymbol{\eta} := (\eta_1, \dots, \eta_m)$ with $\eta_i := (f(\gamma_i) - \mu \cdot (\gamma_i^2 - 1))$ and $\gamma_i \stackrel{iid}{\sim} \mathcal{N}(0, 1)$, $i \in [m]$. There exist constants $c, C > 0$ such that, if the number of observations obeys

$$m \geq c \cdot (\omega_e(\mathcal{C}_0))^2, \quad (23)$$

then, the solution \widehat{X} of (9) satisfies with high probability:

$$\|\widehat{X} - \mu X_0\|_F \leq C \cdot (\mathbb{E}_{\boldsymbol{\eta}}[\omega_e(\mathcal{C}_0; \boldsymbol{\eta})] + v\sqrt{2} \cdot \omega_g(\mathcal{C}_0) + \chi) / \sqrt{m}. \quad (24)$$

3.2 Proof outline of Theorem 3.1

Let $\mathbf{y} = (y_1, \dots, y_m)$ consist of m observations: $y_i = f_i(\mathbf{a}_i^T \mathbf{x}_0)$, $i \in [m]$. We use X_0 and \widehat{X} to denote $\mathbf{x}_0 \mathbf{x}_0^T$ and the solution to (9), respectively. For convenience, we write the loss function in (9) as $\mathcal{L}(X) := \|\mathbf{y} - \mathcal{A}(X)\|_2^2$, where the operator $\mathcal{A} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ returns $\mathcal{A}(X) := (\mathbf{a}_1^T X \mathbf{a}_1 - \text{tr}(X), \dots, \mathbf{a}_m^T X \mathbf{a}_m - \text{tr}(X))^T$. We define the error matrix $\widehat{V} = \widehat{X} - \mu X_0$. We need to upper bound $\|\widehat{V}\|_F$ to establish Thm. 2.2. Towards this direction, let us consider the excess loss function

$$\begin{aligned} 0 &\leq \mathcal{L}(\mu X_0) - \mathcal{L}(\widehat{X}) \\ &= \|\mathbf{y} - \mathcal{A}(\mu X_0)\|_2^2 - \|\mathbf{y} - \mathcal{A}(\mu X_0 + \widehat{V})\|_2^2. \end{aligned} \quad (25)$$

The nonnegativity follows by optimality of \widehat{X} and feasibility of μX_0 (recall that $\mu > 0$ and $\tilde{\mu} = \mu$). Therefore,

$$\|\mathcal{A}(\widehat{V})\|_2^2 \leq 2 \langle \mathbf{y} - \mu \cdot \mathcal{A}(X_0), \mathcal{A}(\widehat{V}) \rangle. \quad (26)$$

On the one hand, recall from (9) that \widehat{V} satisfies $\text{tr}(\widehat{V}) \leq 0$, $\mu X_0 + \widehat{V} \succeq 0$, and $\mu X_0 + \widehat{V} \in \mathcal{K}_{\mathcal{R}}$, i.e., $\widehat{V} \in \mathcal{C}_0$. Thus,

$$\widehat{V} / \|\widehat{V}\|_F \in \mathcal{E} := \mathcal{C}_0 \cap \mathcal{S}_{F,n}. \quad (27)$$

On the other hand, observe in (26) that the LHS (resp., RHS) is homogeneous of degree 2 (resp., 1). With these, it follows from (26) that

$$\begin{aligned} \|\widehat{V}\|_F \cdot \inf_{V \in \mathcal{E}} \|\mathcal{A}(V)\|_2^2 &\leq \\ &2 \sup_{V \in \mathcal{E}} \langle \mathcal{A}(V), \mathbf{y} - \mu \cdot \mathcal{A}(X_0) \rangle. \end{aligned} \quad (28)$$

The strategy is now clear: we will obtain high-probability lower and upper bounds on the LHS and on the RHS, respectively. This will immediately upper bound $\|\widehat{V}\|_F$.

Lower bound: We lower bound the LHS of (28) by employing Mendelson's *Small Ball method* [16]. We defer details to Sec. B.1 of the supplementary material.

Lemma 3.1 (Lower bound). *There exists positive absolute constant $c > 0$ such that for any $s > 0$ it holds with probability at least $1 - e^{-s^2/4}$:*

$$\inf_{V \in \mathcal{E}} \|\mathcal{A}(V)\|_2 \geq c\sqrt{m} - 4\omega_e(\mathcal{C}_0) - s. \quad (29)$$

Upper bound: The key observation for upper bounding the RHS of (28) is that we have set up the objective function of (9) and we have chosen the value of μ in such a way that the following holds. Let $z_i = f(\mathbf{a}_i^T \mathbf{x}_0) - \mu \cdot ((\mathbf{a}_i^T \mathbf{x}_0)^2 - 1)$, $i \in [m]$. Then, setting $\gamma_i := \mathbf{a}_i^T \mathbf{x}_0$, $i \in [m]$, observing that $\gamma_i \sim \mathcal{N}(0, 1)$, and recalling the definition of μ , we obtain

$$\mathbb{E}[z_i \cdot (\gamma_i^2 - 1)] = 0. \quad (30)$$

Therefore, z_i is uncorrelated with $((\mathbf{a}_i^T \mathbf{x}_0)^2 - 1)$, and (by definition) it only depends on the measurement vector \mathbf{a}_i through $\mathbf{a}_i^T \mathbf{x}_0$. Using these, it can be shown that $\mathbb{E}\langle \mathbf{y} - \mu \mathcal{A}(X_0), \mathcal{A}(V) \rangle = 0$ for each fixed V (see (59) in Sec. B.3 of the supplementary material). On the other hand, in (28), we need to uniformly bound the deviation over all V . To do that while exploiting the key observation above, we decompose the measurement vectors as follows:

$$\mathbf{a}_i = \mathbf{x}_0 \mathbf{x}_0^T \mathbf{a}_i + (\mathbf{I} - \mathbf{x}_0 \mathbf{x}_0^T) \mathbf{a}_i =: P \mathbf{a}_i + P^\perp \mathbf{a}_i,$$

where P, P^\perp denote the projection operators to the direction of \mathbf{x}_0 and to its orthogonal subspace, respectively. With this representation, it can be shown that

$$\langle \mathbf{y} - \mu \cdot \mathcal{A}(X_0), \mathcal{A}(V) \rangle = \text{Term I} + \text{Term II} + \text{Term III}$$

where,

$$\text{Term I} := \sum_{i \in [m]} \text{tr}(V X_0) (\gamma_i^2 - 1) z_i, \quad (31a)$$

$$\text{Term II} := \sum_{i \in [m]} (\mathbf{a}_{i,\perp}^T V \mathbf{a}_{i,\perp} + \text{tr}(V(X_0 - \mathbf{I}))) z_i, \quad (31b)$$

$$\text{Term III} := \sum_{i \in [m]} (\mathbf{x}_0^T V \mathbf{a}_{i,\perp} + \mathbf{a}_{i,\perp}^T V \mathbf{x}_0) \gamma_i z_i. \quad (31c)$$

and, we denote $\mathbf{a}_{i,\perp} = P^\perp \mathbf{a}_i$ and $z_i = f(\gamma_i) - \mu(\gamma_i^2 - 1)$. In Sec. B.2 we upper bound $\mathbb{E}[\sup_{V \in \mathcal{E}} \text{Term } \star]$ for all three terms, i.e., $\star = I, II, III$; each one giving rise to one of the terms in the final upper bound of the following lemma.

Lemma 3.2 (Upper bound). *Let v, χ and $\boldsymbol{\eta}$ be defined as in the statement of Thm. 3.1. Then,*

$$\mathbb{E} \sup_{V \in \mathcal{E}} \langle \mathbf{y} - \mu \mathcal{A}(X_0), \mathcal{A}(V) \rangle \leq \sqrt{m} (\chi + 2 \mathbb{E}[\omega_e(\mathcal{C}_0; \boldsymbol{\eta})] + \sqrt{2} v \omega_g(\mathcal{C}_0)). \quad (32)$$

While we defer the proof to the supplementary material, it is not hard to see that (30) is key in bounding the supremum of Term I by $\chi\sqrt{m}$, rather than a trivial bound of order m .

Putting things together: First, Lemma 3.1 and (13) imply that there exists constant $c > 0$ such that, with probability at least 0.995, we have: $\inf_{V \in \mathcal{E}} \frac{1}{m} \|\mathcal{A}(V)\|_2^2 \geq c$. Second, Lemma 3.2 combined with Markov's inequality imply that, with probability at least 0.995, we have: $\sup_{V \in \mathcal{E}} \langle \mathbf{y} - \mu \cdot \mathcal{A}(X_0), \mathcal{A}(V) \rangle \leq C \cdot \sqrt{m} \cdot (\chi + \sqrt{2} \cdot v \cdot \omega_g(\mathcal{C}_0) + 2\mathbb{E}[\omega_e(\mathcal{E}; \boldsymbol{\eta})])$. Then, by union bound, (28) holds with probability 0.99.

3.3 Controlling the weighted empirical width

Note that $\omega_e(\mathcal{C}; \mathbf{p})$ depends both on the geometry of the cone \mathcal{C} and on the weights p_i . We present two ways of controlling $\omega_e(\mathcal{C}; \mathbf{p})$ in terms of simpler geometric quantities.

3.3.1 First bound: generic chaining

This general bound captures the geometry by the γ_1 and γ_2 -functionals with respect to appropriate metrics, and quantifies the role of the weights by ℓ_∞ and ℓ_2 -norms of \mathbf{p} .

Lemma 3.3 (Generic chaining bound). *For a cone $\mathcal{C} \subset \mathbb{S}_n$ and $\mathbf{p} \in \mathbb{R}^m$ there exists universal constant $C > 0$ such that $\omega_e(\mathcal{C}; \mathbf{p}) \leq (C/\sqrt{m}) \cdot (\|\mathbf{p}\|_2 \cdot \gamma_2(\mathcal{C}; \|\cdot\|_F) + \|\mathbf{p}\|_\infty \cdot \gamma_1(\mathcal{C}; \|\cdot\|_2))$. In particular, $\omega_e(\mathcal{C}) = \omega_e(\mathcal{C}; \mathbf{1}) \leq C \cdot (\gamma_2(\mathcal{C}; \|\cdot\|_F) + \frac{\gamma_1(\mathcal{C}; \|\cdot\|_2)}{\sqrt{m}})$.*

See Sec. C of the supplementary for a proof. Note that further using the crude bound $\|\mathbf{p}\|_\infty \leq \|\mathbf{p}\|_2$ implies:

$$\omega_e(\mathcal{C}; \mathbf{p}) \leq \frac{C\|\mathbf{p}\|_2}{\sqrt{m}} \cdot (\gamma_2(\mathcal{C}; \|\cdot\|_F) + \gamma_1(\mathcal{C}; \|\cdot\|_2)). \quad (33)$$

3.3.2 Second bound: polarity

Alternatively, one can apply polarity arguments, by extending recent results regarding the Gaussian width to more general notions such as the weighted empirical width. The idea is as follows. By using polarity it can be shown that $\omega_e(\mathcal{C}; \mathbf{p})$ is upper bounded by the expected distance of $H_{\mathbf{p}}$ (cf. Definition 3) to the polar cone \mathcal{C}° . When \mathcal{C} is the tangent cone of some convex proper function (say) \mathcal{R} , then, \mathcal{C}° is the cone of subdifferential of \mathcal{R} . Thus, we arrive at the result below that is based on techniques from [26, 7].

Proposition 3.1 (Polarity bound). *Let $\mathcal{R} : \mathbb{S}_n \rightarrow \mathbb{R}$ be a proper convex function, fix $X_0 \in \mathbb{S}_n$, and let $\tilde{\mathcal{K}}_{\mathcal{R}} := \{V : \mathcal{R}(X_0 + V) \leq \mathcal{R}(X_0)\}$. Assume that the subdifferential $\partial\mathcal{R}(X_0)$ is non-empty and does not contain the origin. Then, for $H_{\mathbf{p}} = \frac{1}{\sqrt{m}} \sum_{i \in [m]} p_i \cdot \varepsilon_i \mathbf{a}_i \mathbf{a}_i^T$ it holds:*

$$\omega_e^2(\mathcal{D}(\tilde{\mathcal{K}}_{\mathcal{R}}, X_0); \mathbf{p}) \leq \mathbb{E} \left[\inf_{\lambda \geq 0} \inf_{V \in \partial\mathcal{R}(X_0)} \|H_{\mathbf{p}} - \lambda \cdot V\|_F^2 \right].$$

For illustration, in Lemmas 3.4 and 3.5 below we apply Proposition 3.1 to obtain bounds for the following two cones: (a) $\mathcal{D}(\mathcal{K}_0, \mu X_0)$; and, (b) $\mathcal{D}(\mathcal{K}_{\mathcal{R}}, \mu X_0)$. We defer the proofs to Sec. D in the supplementary material.

Lemma 3.4 (Weighted empirical width of \mathcal{C}_+). For $\mathbf{p} := (p_1, \dots, p_m)$ and $\mathcal{C}_+ = \{V : \mu X_0 + V \succeq 0 \text{ and } \text{tr}(V) \leq 0\}$, there exists constant $C > 0$ such that $\omega_e(\mathcal{C}_+; \mathbf{p}) \leq \frac{C}{\sqrt{m}} (\|\mathbf{p}\|_2 \sqrt{n} + \|\mathbf{p}\|_\infty n)$. In particular, $\omega_e(\mathcal{C}_+; \mathbf{1}) \leq C\sqrt{n}$, provided that $m \geq cn$ for some constant $c > 0$.

Lemma 3.5 (Weighted empirical width of $\mathcal{C}_{\text{sparse}}$). Let $X_0 = \mathbf{x}_0 \mathbf{x}_0^T$ where \mathbf{x}_0 is k -sparse. For $\mathbf{p} := (p_1, \dots, p_m)$ and the cone $\mathcal{C}_{\text{sparse}} := \{V : \|\mu X_0 + V\|_1 \leq \|\mu X_0\|_1\}$, there exists constant $C > 0$ such that $\omega_e(\mathcal{C}_{\text{sparse}}; \mathbf{p}) \leq \frac{C}{\sqrt{m}} k \sqrt{\log(n/k)} (\|\mathbf{p}\|_2 + \|\mathbf{p}\|_\infty \sqrt{2 \log(n/k)})$. In particular, $\omega_e(\mathcal{C}_{\text{sparse}}; \mathbf{1}) \leq Ck \sqrt{\log(n/k)}$, provided that $m \geq ck^2 \log(n/k)$, for $c > 0$.

3.4 Signal recovery with imperfect tuning

Thm. 3.1 assumes ideal tuning $\tilde{\mu} = \mu$ in (9), which guarantees that X_0 belongs to the boundary of the constraint set \mathcal{K}_0 (cf. (21)). Moreover, its bounds in terms of the Gaussian/empirical widths of $\mathcal{D}(\mathcal{K}_{\mathcal{R}}, \mathbf{x}_0)$ are most informative when X_0 is at the boundary of $\mathcal{K}_{\mathcal{R}}$. Thm. 3.2 below relaxes these assumptions, via a local analysis that gives rise to the local versions of the Gaussian/empirical widths. The proof is deferred to the supplementary material. For example, the theorem shows that the error performance of (9) is controlled as long as $\tilde{\mu} \geq \mu$ (but, not necessarily equal).

Theorem 3.2 (Imperfect tuning). Suppose that $\mathbf{a}_i \sim \mathcal{N}(0, \mathbf{I}_n)$, and that \mathbf{y} follows the model in (1). Recall the definitions of μ and of τ, v, χ . Further let

$$\mathcal{K}_{\mathbf{x}_0} := \{\mathcal{K}_0 - \mu X_0\} \cap \{\mathcal{K}_{\mathcal{R}} - \mu X_0\}. \quad (34)$$

Finally, define $\boldsymbol{\eta} := (\eta_1, \dots, \eta_m)$, where $\eta_i := (f(\gamma_i) - \mu \cdot (\gamma_i^2 - 1))$, $i \in [m]$ for $\gamma_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. For any $t > 0$, there exist constants $c, C > 0$ such that, if the number of observations obeys

$$m \geq c \cdot (\omega_{e,t}(\mathcal{K}_{\mathbf{x}_0})/t)^2, \quad (35)$$

then, with high probability, the solution \hat{X} of (9) satisfies:

$$\|\hat{X} - \mu X_0\|_F \leq \frac{C}{\sqrt{m}} \cdot \left(\mathbb{E}_{\boldsymbol{\eta}} \left[\frac{\omega_{e,t}(\mathcal{K}_{\mathbf{x}_0}; \boldsymbol{\eta})}{t} \right] + v\sqrt{2} \frac{\omega_{g,t}(\mathcal{K}_{\mathbf{x}_0})}{t} + \chi \right) + t.$$

4 Simulations

In this section, we show numerical results regarding the performance of the proposed method. In order to solve the optimization in (9), we use PhasePack library [6] (designed for quadratic measurements) with appropriate modifications to account for the different objective function compared to the PhaseLift [5] (cf. Remark 2), as well as for the regularization in case of structured signal recovery.

In our setup, the unknown unit-norm signal \mathbf{x}_0 has dimension $n = 50$. m measurements are generated according to




Figure 1: Performance for unstructured signal recovery.




Figure 2: Performance of (9) for sparse signal recovery.

(1) with $\mathbf{a}_i \sim \mathcal{N}(0, \mathbf{I})$, for the following link functions: (1) $f^{(1)}(u) = (|u| + \frac{1}{2}) \cdot \mathbb{1}_{\{|u| < 3\}} + \frac{7}{2} \cdot \mathbb{1}_{\{|u| \geq 3\}}$; and (2) $f^{(2)}(u) = \frac{1}{2} \cdot (|2|u|^2| + \frac{1}{2}) \cdot \mathbb{1}_{\{|u|^2 < 4\}} + \frac{17}{4} \cdot \mathbb{1}_{\{|u|^2 \geq 4\}}$. Note that $f^{(1)}$ and $f^{(2)}$ correspond to quantized measurements with alphabets sizes 4 and 9, respectively. We also experiment with random link functions as per (2); in particular, $y_i \sim \text{Poisson}(|\mathbf{a}_i^T \mathbf{x}_0|^2)$, for all $i \in [m]$. In all cases, we solve (9) with $\tilde{\mu} = \mu$, with appropriate μ for each link function. The results shown are averages over 40 trials.

Figure 1 shows the estimation error of (9) as a function of m for a generic \mathbf{x}_0 (aka no constraint $X \in \mathcal{K}_{\mathcal{R}}$). We also compare the proposed method with the PCA-based method in [33, 28]; we empirically observe that the former method outperforms the latter for all tested link functions. Similarly, Figure 2 illustrates the performance of (9) in a sparse-signal recovery setting with sparsity levels $k \in \{5, 10\}$.

Our simulations also confirm the point briefly made in Sec. 1.3. The precise nature of our results can be useful in optimal design of system parameters. In particular, our analysis suggests that link functions with smaller effective noise parameter τ/μ result in smaller error. Indeed, the ratio τ/μ of $f^{(1)}$ and $f^{(2)}$ is equal to 2.25 and 1.51, respectively, which is in agreement to the better recovery performance of $f^{(2)}$ in our simulations.

Acknowledgements

C.T. would like to thank Pablo Parrilo and Joel Tropp for helpful discussions.

References

- [1] Dennis Amelunxen, Martin Lotz, Michael B. McCoy, and Joel A. Tropp. Living on the edge: Phase transitions in convex programs with random data. *Information and Inference*, 3(3):224–294, Sep. 2014.
- [2] Sohaib Bahmani and Justin Romberg. Phase Retrieval Meets Statistical Learning Theory: A Flexible Convex Relaxation. *arXiv preprint arXiv:1610.04210*, 2016.
- [3] Radu Balan, Bernhard G Bodmann, Peter G Casazza, and Dan Edidin. Painless reconstruction from magnitudes of frame coefficients. *Journal of Fourier Analysis and Applications*, 15(4):488–501, 2009.
- [4] David R Brillinger. A generalized linear model with “gaussian” regressor variables. *A Festschrift For Erich L. Lehmann*, page 97, 1982.
- [5] Emmanuel J Candes, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [6] Rohan Chandra, Ziyuan Zhong, Justin Hontz, Val McCulloch, Christoph Studer, and Tom Goldstein. PhasePack: A phase retrieval library. *Asilomar Conference on Signals, Systems, and Computers*, 2017.
- [7] Venkat Chandrasekaran, Benjamin Recht, Pablo A Parrilo, and Alan S Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.
- [8] Murat A Erdogdu, Lee H Dicker, and Mohsen Bayati. Scaled least squares estimator for glms in large-scale problems. In *Advances in Neural Information Processing Systems*, pages 3324–3332, 2016.
- [9] Simon Foucart and Holger Rauhut. *A Mathematical Introduction to Compressive Sensing*. Springer Science & Business Media, 2013.
- [10] Martin Genzel. High-dimensional estimation of structured signals from non-linear observations with general convex loss functions. *IEEE Transactions on Information Theory*, 63(3):1601–1619, 2017.
- [11] Larry Goldstein, Stanislav Minsker, and Xiaohan Wei. Structured signal recovery from non-linear and heavy-tailed measurements. *arXiv preprint arXiv:1609.01025*, 2016.
- [12] Tom Goldstein and Christoph Studer. PhaseMax: Convex Phase Retrieval via Basis Pursuit. *arXiv preprint arXiv:1610.07531*, 2016.
- [13] Daniel Hsu, Sham Kakade, Tong Zhang, et al. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability*, 17, 2012.
- [14] Ker-Chau Li and Naihua Duan. Regression analysis under link violation. *The Annals of Statistics*, pages 1009–1052, 1989.
- [15] Yue M. Lu and Gen Li. Phase transitions of spectral initialization for high-dimensional nonconvex estimation. *arXiv preprint arXiv:1702.06435*, 2017.
- [16] Shahar Mendelson. Learning without concentration. *Journal of the ACM (JACM)*, 62(3):21, 2015.
- [17] Marco Mondelli and Andrea Montanari. Fundamental limits of weak recovery with applications to phase retrieval. *arXiv preprint arXiv:1708.05932*, 2017.
- [18] Matej Neykov, Zhaoran Wang, and Han Liu. Agnostic estimation for misspecified phase retrieval models. In *Advances in Neural Information Processing Systems*, pages 4089–4097, 2016.
- [19] Samet Oymak. Learning compact neural networks with regularization. *arXiv preprint arXiv:1802.01223*, 2018.
- [20] Samet Oymak, Amin Jalali, Maryam Fazel, Yonina C Eldar, and Babak Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Transactions on Information Theory*, 61(5):2886–2908, 2015.
- [21] Samet Oymak and Mahdi Soltanolkotabi. Fast and reliable parameter estimation from nonlinear observations. *arXiv preprint arXiv:1610.07108*, 2016.
- [22] Yaniv Plan and Roman Vershynin. The generalized lasso with non-linear observations. *arXiv preprint arXiv:1502.04071*, 2015.
- [23] Yaniv Plan, Roman Vershynin, and Elena Yudovina. High-dimensional estimation with geometric constraints. *arXiv preprint arXiv:1404.3749*, 2014.
- [24] Mohammadreza Soltani and Chinmay Hegde. Fast algorithms for demixing sparse signals from nonlinear observations. *IEEE Transactions on Signal Processing*, 2017.
- [25] Mahdi Soltanolkotabi. Structured signal recovery from quadratic measurements: Breaking sample complexity barriers via nonconvex optimization. *arXiv preprint arXiv:1702.06175*, 2017.

- [26] Mihailo Stojnic. Various thresholds for ℓ_1 -optimization in compressed sensing. *arXiv preprint arXiv:0907.3666*, 2009.
- [27] Michel Talagrand. *The generic chaining: upper and lower bounds of stochastic processes*. Springer Science & Business Media, 2006.
- [28] Yan Shuo Tan. Sparse phase retrieval via sparse pca despite model misspecification: A simplified and extended analysis. *arXiv preprint arXiv:1712.04106*, 2017.
- [29] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. LASSO with non-linear measurements is equivalent to one with linear measurements. In *Advances in Neural Information Processing Systems*, pages 3420–3428, 2015.
- [30] Joel A Tropp. Convex recovery of a structured signal from independent random linear measurements. In *Sampling Theory, a Renaissance*, pages 67–101. Springer, 2015.
- [31] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2017.
- [32] Zhuoran Yang, Krishnakumar Balasubramanian, and Han Liu. High-dimensional non-gaussian single index models via thresholded score function estimation. In *International Conference on Machine Learning*, pages 3851–3860, 2017.
- [33] Zhuoran Yang, Krishnakumar Balasubramanian, Zhaoran Wang, and Han Liu. Estimating high-dimensional non-gaussian multiple index models via stein’s lemma. In *Advances in Neural Information Processing Systems*, pages 6099–6108, 2017.
- [34] Zhuoran Yang, Lin F Yang, Ethan X Fang, Tuo Zhao, Zhaoran Wang, and Matey Neykov. Misspecified nonconvex statistical optimization for phase retrieval. *arXiv preprint arXiv:1712.06245*, 2017.
- [35] Xinyang Yi, Zhaoran Wang, Constantine Caramanis, and Han Liu. Optimal linear estimation under unknown nonlinear transform. In *Advances in neural information processing systems*, pages 1549–1557, 2015.