



Fitting Tractable Convex Sets to Support Function Evaluations

Yong Sheng Soh¹  · Venkat Chandrasekaran²

Received: 9 May 2019 / Revised: 23 October 2019 / Accepted: 14 October 2020 /

Published online: 3 January 2021

© Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

The geometric problem of estimating an unknown compact convex set from evaluations of its support function arises in a range of scientific and engineering applications. Traditional approaches typically rely on estimators that minimize the error over all possible compact convex sets; in particular, these methods allow for limited incorporation of prior structural information about the underlying set and the resulting estimates become increasingly more complicated to describe as the number of measurements available grows. We address both of these shortcomings by describing a framework for estimating *tractably specified* convex sets from support function evaluations. Building on the literature in convex optimization, our approach is based on estimators that minimize the error over structured families of convex sets that are specified as linear images of concisely described sets—such as the simplex or the spectraplex—in a higher-dimensional space that is not much larger than the ambient space. Convex sets parametrized in this manner are significant from a computational perspective as one can optimize linear functionals over such sets efficiently; they serve a different purpose in the inferential context of the present paper, namely, that of incorporating *regularization* in the reconstruction while still offering considerable expressive power. We provide a geometric characterization of the asymptotic behavior of our estimators, and our analysis relies on the property that certain sets which admit semialgebraic descriptions are Vapnik–Chervonenkis classes. Our numerical experiments highlight the utility of our framework over previous approaches in settings in which the measurements available are noisy or small in number as well as those in which the underlying set to be reconstructed is non-polyhedral.

Editor in Charge: Kenneth Clarkson

The authors were supported in part by NSF grants CCF-1350590 and CCF-1637598, by Air Force Office of Scientific Research Grant FA9550-16-1-0210, by a Sloan research fellowship, and an A*STAR (Agency for Science, Technology and Research, Singapore) fellowship.

Extended author information available on the last page of the article

Keywords Constrained shape regression · Convex regression · Entropy of semialgebraic sets · K -means clustering · Simplicial polytopes · Stochastic equicontinuity

1 Introduction

We consider the problem of estimating a compact convex set given (possibly noisy) evaluations of its support function. Formally, let $K^* \subset \mathbb{R}^d$ be a set that is compact and convex. The support function $h_{K^*}(u)$ of the set K^* evaluated in the direction $u \in S^{d-1}$ is defined as

$$h_{K^*}(u) := \sup_{x \in K^*} \langle x, u \rangle.$$

Here $S^{d-1} := \{x \mid \|x\|_2 = 1\} \subset \mathbb{R}^d$ denotes the $(d - 1)$ -dimensional unit sphere. In words, the quantity $h_{K^*}(u)$ measures the maximum displacement in the direction u intersecting K^* . Given noisy support function evaluations $\{(u^{(i)}, y^{(i)}) \mid y^{(i)} = h_{K^*}(u^{(i)}) + \varepsilon^{(i)}, 1 \leq i \leq n\}$, where each $\varepsilon^{(i)}$ denotes additive noise, our goal is to reconstruct a convex set \hat{K} that is close to K^* .

The problem of estimating a convex set from support function evaluations arises in a wide range of problems such as computed tomography [24], target reconstruction from laser-radar measurements [18], and projection magnetic resonance imaging [13]. For example, in tomography the extent of the absorption of parallel rays projected onto an object provides support information [24,30], while in robotics applications support information can be obtained from an arm clamping onto an object in different orientations [24]. A natural approach to fit a compact convex set to support function data is the following least-squares estimator (LSE):

$$\hat{K}_n^{\text{LSE}} \in \underset{\text{compact, convex } K \subset \mathbb{R}^d}{\text{arg min}} \frac{1}{n} \sum_{i=1}^n (y^{(i)} - h_K(u^{(i)}))^2. \quad (1)$$

An LSE always exists and it is not defined uniquely, although it is always possible to select a polytope that is an LSE; this is the choice that is most commonly employed and analyzed in prior work. For example, the algorithm proposed by Prince and Willsky [24] for planar convex sets reconstructs a polygonal LSE described in terms of its facets, while the algorithm proposed by Gardner and Kiderlen [10] for convex sets in any dimension provides a polytopal LSE reconstruction described in terms of extreme points. The least-squares estimator \hat{K}_n^{LSE} is a consistent estimator of K^* , but it has a number of significant drawbacks. In particular, as the formulation (1) does not incorporate any additional structural information about K^* beyond convexity, the estimator \hat{K}_n^{LSE} can provide poor reconstructions when the measurements available are noisy or small in number. The situation is problematic even when the number of measurements available is large, as the complexity of the resulting estimate grows with the number of measurements in the absence of any regularization due to the regression problem (1) being non-parametric (the collection of all compact convex sets in \mathbb{R}^d is not

finitely parametrized); consequently, the facial structure of the reconstruction provides little information about the geometry of the underlying set.¹ Finally, if the underlying set K^* is not polyhedral, a polyhedral choice for the solution \hat{K}_n^{LSE} (as is the case with much of the literature on this topic) can provide poor reconstructions. Indeed, even for intrinsically “simple” convex bodies such as the Euclidean ball, one necessarily requires many vertices or facets to obtain accurate polyhedral approximations. Figure 1 provides an illustration of these points.

1.1 Our Contributions

To address the drawbacks underlying the least-squares estimator, we seek a framework that *regularizes* the complexity of the reconstruction in the formulation (1). A natural approach to developing such a framework is to design an estimator with the same objective as in (1) but in which the decision variable K is constrained to lie in a subclass \mathcal{F} of the collection of all compact, convex sets. For such a method to be useful, the subclass \mathcal{F} must balance several considerations. First, \mathcal{F} should be sufficiently expressive in order to faithfully model various attributes of convex sets that arise in applications (for example, sets consisting of both smooth and singular features in their boundary). Second, the elements of \mathcal{F} should be suitably structured so that incorporating the constraint $K \in \mathcal{F}$ leads to estimates that are more robust to noise; further, the type of structure underlying the sets in \mathcal{F} also informs the analysis of the statistical properties of the constrained analog of (1) as well as the development of efficient algorithms for computing the associated estimate. Building on the literature on lift-and-project methods in convex optimization [12,34], we consider families \mathcal{F} in which the elements are specified as images under linear maps of a fixed ‘concisely specified’ compact convex set; the choice of this set governs the expressivity of the family \mathcal{F} and we discuss this in greater detail in the sequel. Due to the availability of computationally tractable procedures for optimization over linear images of concisely described convex sets [21], the study of such descriptions constitutes a significant topic in optimization. We employ these ideas in a conceptually different context in the setting of the present paper, namely that of incorporating regularization in the reconstruction, which addresses many of the drawbacks with the LSE outlined previously. Formally, we consider the following regularized convex set regression problem:

$$\hat{K}_n^C \in \arg \min_{K: K=A(C), A \in L(\mathbb{R}^q, \mathbb{R}^d)} \frac{1}{n} \sum_{i=1}^n (y^{(i)} - h_K(u^{(i)}))^2. \quad (2)$$

Here $C \subset \mathbb{R}^q$ is a user-specified compact convex set and $L(\mathbb{R}^q, \mathbb{R}^d)$ denotes the set of linear maps from \mathbb{R}^q to \mathbb{R}^d . The set C governs the expressive power of the family $\{A(C) \mid A \in L(\mathbb{R}^q, \mathbb{R}^d)\}$. In addition to this consideration, our choices for C are also driven by statistical and computational aspects of the estimator (2). Our analysis

¹ We note that this is the case even though the estimator \hat{K}_n^{LSE} is consistent; in particular, consistency simply refers to the convergence as $n \rightarrow \infty$ of \hat{K}_n^{LSE} to K^* in a topological sense (e.g., in Hausdorff distance) and it does not provide any information about the facial structure of \hat{K}_n^{LSE} relative to that of K^* .

of the statistical properties of the estimator (2) relies in part on the observation that sets \mathcal{F} that admit certain semialgebraic descriptions form VC classes; this fact serves as the foundation for our characterization based on stochastic equicontinuity of the asymptotic properties of the estimator \hat{K}_n^C as $n \rightarrow \infty$. On the computational front, the algorithms we propose for (2) require that the support function associated to C as well as its derivatives (when they exist) can be computed efficiently. Motivated by these issues, the choices of C that we primarily discuss in our examples and numerical illustrations are the simplex and the spectraplex:

Example The *simplex* in \mathbb{R}^q is the set

$$\Delta^q := \{x \in \mathbb{R}^q \mid x \geq 0, \langle x, 1 \rangle = 1\},$$

where $1 = (1, \dots, 1)^T$. Convex sets expressed as projections of Δ^q are precisely polytopes with at most q extreme points.

Example Let $\mathbb{S}^p \cong \mathbb{R}^{\binom{p+1}{2}}$ denote the space of $p \times p$ real symmetric matrices. The *spectraplex* $\mathcal{O}^p \subset \mathbb{S}^p$ (also called the *free spectrahedron*) is the set

$$\mathcal{O}^p := \{X \in \mathbb{S}^p \mid X \succeq 0, \langle X, I \rangle = 1\},$$

where $I \in \mathbb{S}^p$ is the identity matrix. The spectraplex is a semidefinite analog of the simplex, and it is especially useful if we seek non-polyhedral reconstructions, as can be seen in Fig. 1 and in Sect. 5; in particular, linear images of the spectraplex exhibit both smooth and singular features in their boundaries.

The specific selection of C from the families $\{\Delta^q\}_{q=1}^\infty$ and $\{\mathcal{O}^p\}_{p=1}^\infty$ is governed by the complexity of the reconstruction one seeks, which is typically based on prior information about K^* . Our analysis in Sect. 3 of the statistical properties of the estimator (2) relies on the availability of such additional knowledge about the complexity of K^* . In practice in the absence of such information, cross-validation may be employed to obtain a suitable reconstruction; see Sect. 6.

In Sect. 2 we discuss preliminary aspects of our technical setup such as properties of the set of minimizers of the problem (2) as well as a stylized probabilistic model for noisy support function evaluations. These serve as a basis for the subsequent development in the paper. In Sect. 3 we provide the main theoretical guarantees of our approach. In our first result, we show that the sequence of estimates $\{\hat{K}_n^C\}_{n=1}^\infty$ converges almost surely (as $n \rightarrow \infty$) in the Hausdorff metric to that linear image of C which is closest to the underlying set K^* (see Theorem 3.1). Under additional conditions, we also characterize certain asymptotic distributional aspects of the sequence $\{\hat{K}_n^C\}_{n=1}^\infty$ (see Theorem 3.4); this result is based on a functional central limit theorem, which requires the computation of appropriate entropy bounds for Vapnik–Chervonenkis (VC) classes of sets that admit semialgebraic descriptions, and it is here that our choice of C as either a simplex or a spectraplex plays a prominent role. Our third result describes the facial structure of $\{\hat{K}_n^C\}_{n=1}^\infty$ in relation to the underlying set K^* . We prove under appropriate conditions that if K^* is a polytope, our approach provides a reconstruction that recovers all the simplicial faces (for sufficiently large n); if K^* is

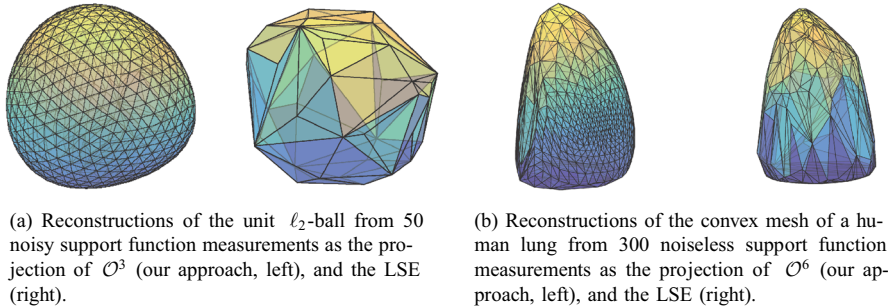


Fig. 1 Comparison between our approach and the LSE

a simplicial polytope, we recover a polytope that is combinatorially equivalent to K^* . This result also applies more generally to ‘rigid’ faces for non-polyhedral K^* (see Theorem 3.9).

In the sequel, we relate our formulation (2) (when C is a simplex) to the task of fitting piecewise affine convex functions to data (known as max-affine regression) as well as K -means clustering. Accordingly, the algorithm we propose in Sect. 4 for computing \hat{K}_n^C bears significant similarities with methods for max-affine regression [20] as well as Lloyd’s algorithm for clustering problems.

A restriction in the development in this paper is that the simplex and the spectraplex represent particular affine slices of the non-negative orthant and the cone of positive semidefinite matrices. In principle, one can further optimize these slices (both their orientation and their dimension) in (2) to obtain improved reconstructions. However, this additional degree of flexibility in (2) leads to technical complications in establishing asymptotic normality in Sect. 3.2 as well as to challenges in developing algorithms for solving (2) (even to obtain a local optimum). The root of these difficulties lies in the fact that it is hard to characterize the variation in the support function with respect to small changes in the slice. We remark on these challenges in greater detail in Sect. 6, and for the remainder of the paper we proceed with investigating the estimator (2).

1.2 Related Work

1.2.1 Consistency of Convex Set Regression

There is a well-developed body of prior work on the consistency of convex set regression (1). Gardner et al. [11] prove that the (polyhedral) estimates \hat{K}_n^{LSE} converge almost surely to the underlying set K^* in the Hausdorff metric as $n \rightarrow \infty$ provided the directions $\{u^{(i)}\}_{i=1}^n$ cover the sphere in a suitably uniform manner. Guntuboyina [14] analyzes rates of convergence in minimax settings, and also notes that constraining the growth of the number of vertices in the reconstruction as the number of measurements increases provides a form of robustness. Cai et al. [6] study the impact of choosing the directions $\{u^{(i)}\}_{i=1}^n$ adaptively in estimating planar convex sets. In contrast, the consistency result in the present paper corresponding to the constrained estimator (2) is qualitatively different. On the one hand, for a given compact convex set $C \subset \mathbb{R}^q$,

we prove that the sequence of estimates $\{\hat{K}_n^C\}_{n=1}^\infty$ converges to that linear image of C which is closest to the underlying set K^* ; in particular, $\{\hat{K}_n^C\}_{n=1}^\infty$ only converges to K^* if K^* can be represented as a linear image of C . On the other hand, there are several advantages to the framework presented in this paper in comparison with prior work. First, the constrained estimator (2) lends itself to a precise asymptotic distributional characterization which is unavailable in the unconstrained case (1). Second, under appropriate conditions, the constrained estimator (2) recovers the facial structure of the underlying set K^* unlike \hat{K}_{LSE} . More significantly, beyond these technical distinctions, the constrained estimator (2) also yields concisely described non-polyhedral reconstructions (as well as associated consistency and asymptotic distributional characterizations) based on linear images of the spectraplex, in contrast to the usual choice of a polyhedral LSE in the previous literature.

1.2.2 Incorporating Prior Information and Fitting Smooth Boundaries

The problem of integrating prior information about the underlying convex set has been considered in [25], where the authors propose a method of incorporating certain structural or shape priors for fitting convex sets in two dimensions, with a particular focus on settings in which the underlying set is a disc or an ellipsoid. However, the reconstructions produced by the method in [25] are still polyhedral, and the method assumes that the support function evaluations are available at angles that are equally spaced. We are also aware of a line of work [8,15] on fitting convex sets in two dimensions with smooth boundaries to support function measurements. The first of these papers estimates a convex set with a smooth boundary without any vertices, while the second proposes a two-step method in which one initially estimates a set of vertices followed by a second step that connects these vertices via smooth boundaries. In both cases, splines are used to interpolate between the support function evaluations with a subsequent smoothing procedure using the von Mises kernel. The smoothing is done in a local fashion and the resulting reconstruction is increasingly complex to describe as the number of measurements available grows. In contrast, our approach to producing non-polyhedral estimates based on fitting linear images of spectraplices is more global in nature, and we explicitly regularize the complexity of our reconstruction based on the dimension of the spectraplex. Further, the approaches proposed in [8,15] estimate the singular and the smooth parts of the boundary separately, whereas our framework based on linear images of spectraplices estimates these features in a unified manner (for example, see the illustration in Fig. 12). Finally, the methods described in [8,15,25] are only applicable to two-dimensional reconstruction problems, while problems of a three-dimensional nature arise in many contexts (see Sect. 5.4 for an example that involves the reconstruction of a human lung).

1.2.3 Piecewise Affine Convex Regression

The formulation (2) when $C = \Delta^q$ may be viewed as a fitting a piecewise linear function (with at most q pieces) to the given data. This is a special case of the *max-affine regression* problem in which one is interested in fitting a piecewise affine function

(typically with a bound on the number of pieces) to data, which is a topic that has been studied previously [2,16,20]. In particular, our algorithm in Sect. 4 when specialized to the setting $C = \Delta^q$ is analogous to the methods described in [20]. However, our framework and the algorithm in Sect. 4 may also be employed to fit more general convex functions that are not piecewise linear, but that can still be specified in a tractable manner via linear images of the spectraplex.

1.3 Outline

In Sect. 2 we discuss the geometric, algebraic, and analytic aspects of the optimization problem (2); this section serves as the foundation for the subsequent statistical analysis in Sect. 3. Throughout both of these sections, we give several examples that provide additional insight into our mathematical development. We describe algorithms for solving (2) in Sect. 4, and we demonstrate the application of these methods in a range of numerical experiments in Sect. 5. We conclude with a discussion of future directions in Sect. 6.

Notation: Given a convex set $C \subset \mathbb{R}^q$, we denote the associated induced norm by $\|A\|_{C,2} := \sup_{x \in C} \|Ax\|_2$. We denote the unit $\|\cdot\|$ -ball centered at x by $B_{\|\cdot\|}(x) := \{y \mid \|y - x\| \leq 1\}$, and we denote the Frobenius norm by $\|\cdot\|_F$. Given a point $a \in \mathbb{R}^q$ and a subset $U \subseteq \mathbb{R}^q$, we define $\text{dist}(a, U) := \inf_{b \in U} \|a - b\|$, where the norm $\|\cdot\|$ is the Euclidean norm. Last, given any two subsets $U, V \subset \mathbb{R}^q$, the Hausdorff distance between U and V is denoted by

$$d_H(U, V) := \inf_{t \geq 0} \{t \mid U \subseteq V + tB_{\|\cdot\|_2}(0), V \subseteq U + tB_{\|\cdot\|_2}(0)\}.$$

2 Problem Setup and Other Preliminaries

In this section, we begin with a preliminary discussion of the geometric, algebraic, and analytical aspects of our procedure (2); these underpin our subsequent development in this paper. We make the following assumptions about our problem setup for the remainder of the paper:

- (A1) The set $K^\star \subset \mathbb{R}^d$ is compact and convex.
- (A2) The set $C \subset \mathbb{R}^q$ is compact and convex.
- (A3) *Probabilistic Model for Support Function Measurements:* We assume that we are given n independent and identically distributed support function evaluations $\{(u^{(i)}, y^{(i)})\}_{i=1}^n \subset S^{d-1} \times \mathbb{R}$ from the following probabilistic model:

$$P_{K^\star}: \quad y = h_{K^\star}(u) + \varepsilon. \quad (3)$$

Here $u \in S^{d-1}$ is a vector distributed uniformly at random (u.a.r.) over the unit sphere, ε is a centered random variable with variance σ^2 (i.e., $\mathbb{E}[\varepsilon] = 0$ and $\mathbb{E}[\varepsilon^2] = \sigma^2$), and u and ε are independent.

In our analysis, we quantify dissimilarity between convex sets in terms of a metric applied to their respective support functions. Let K_1, K_2 be compact convex sets in \mathbb{R}^d , and let $h_{K_1}(\cdot), h_{K_2}(\cdot)$ be the corresponding support functions. We define the L_p metric to be

$$\rho_p(K_1, K_2) := \left(\int_{S^{d-1}} |h_{K_1}(u) - h_{K_2}(u)|^p \mu(du) \right)^{1/p}, \quad 1 \leq p < \infty, \quad (4)$$

where the integral is with respect to the Lebesgue measure over S^{d-1} ; as usual, we denote $\rho_\infty(K_1, K_2) = \max_u |h_{K_1}(u) - h_{K_2}(u)|$. We prove our convergence guarantees in Sect. 3.1 in terms of the ρ_p -metric. This metric represents an important class of distance measures over convex sets. For instance, it features prominently in the literature on approximating convex sets as polytopes [5]. In addition, the specific case of $p = \infty$ coincides with the Hausdorff distance [27, p. 66].

Due to the form of the estimator (2), one may reparametrize the optimization problem in terms of the linear map A . In particular, by noting that $h_{A(C)}(u) = h_C(A^T u)$, the problem (2) can be reformulated as follows:

$$\hat{A}_n \in \arg \min_{A \in L(\mathbb{R}^q, \mathbb{R}^d)} \frac{1}{n} \sum_{i=1}^n (y^{(i)} - h_C(A^T u^{(i)}))^2. \quad (5)$$

Based on this observation, we analyze some properties of the set of minimizers of (2) via an analysis of (5). (The reformulation (5) is also more conducive to the development of numerical algorithms for solving (2).) In turn, a basic strategy for investigating the asymptotic properties of the estimator (5) is to analyze the minimizers of the loss function at the *population level*. Concretely, for any probability measure P over pairs $(u, y) \in S^{d-1} \times \mathbb{R}$, the loss function with respect to P is defined as

$$\Phi_C(A, P) := \mathbb{E}_P[(h_C(A^T u) - y)^2]. \quad (6)$$

Thus, the focus of our analysis is on studying the set of minimizers of the population loss function $\Phi_C(\cdot, P_{K^*})$:

$$M_{K^*, C} := \arg \min_A \Phi_C(A, P_{K^*}). \quad (7)$$

2.1 Geometric Aspects

In this subsection, we focus on the *convex sets* defined by the elements of the set of minimizers $M_{K^*, C}$. In the next subsection, we consider the elements of $M_{K^*, C}$ as *linear maps*. To begin with, we state a simple lemma on the continuity of $\Phi_C(\cdot, P)$:

Proposition 2.1 *Let P be any probability distribution over measurement pairs $(u, y) \in S^{d-1} \times \mathbb{R}$ satisfying $\mathbb{E}_P[|y|] < \infty$. Then the function $\Phi_C(\cdot, P)$ defined in (6) is continuous.*

The proof uses a simple bound. We state the result formally as we require it at a later part.

Lemma 2.2 *Given any pair of linear maps $A_1, A_2 \in L(\mathbb{R}^q, \mathbb{R}^d)$, any unit-norm vector u , and any scalar y , we have*

$$||h_C(A_1^T u) - y| - |h_C(A_2^T u) - y|| \leq |h_C(A_1^T u) - h_C(A_2^T u)| \leq \|A_1 - A_2\|_{C,2}. \tag{8}$$

Proof First note that $h_C(A_1^T u) \leq h_C(A_2^T u) + h_C((A_1 - A_2)^T u)$. Since u is unit-norm, we have $h_C(A_1^T u) - h_C(A_2^T u) \leq h_C((A_1 - A_2)^T u) \leq \|A_1 - A_2\|_{C,2}$. Similarly we have $h_C(A_2^T u) - h_C(A_1^T u) \leq \|A_1 - A_2\|_{C,2}$. Hence (8) follows. \square

Proof of Proposition 2.1 Let $\epsilon > 0$ be arbitrary. Let $r = 1 + \|A\|_{C,2}$ and pick $\delta = \min \{\epsilon/(3\mathbb{E}[r + |y|]), r\}$. Then for any A_0 satisfying $\|A - A_0\|_{C,2} < \delta$, we have

$$|\Phi(A, P) - \Phi(A_0, P)| = |\mathbb{E}_P[(h_C(A^T u) - y)^2 - (h_C(A_0^T u) - y)^2]| \leq \mathbb{E}_P[|h_C(A^T u) - h_C(A_0^T u)| \cdot |h_C(A^T u) + h_C(A_0^T u) - 2y|].$$

We apply Lemma 2.2 to obtain the bound

$$|h_C(A^T u) + h_C(A_0^T u) - 2y| \leq 2|h_C(A^T u)| + 2|y| + \|A_0 - A\|_{C,2} < 2r + 2|y| + \delta.$$

Combining this bound with our earlier expression, it follows that $|\Phi(A, P) - \Phi(A_0, P)| < \mathbb{E}_P[\delta(2r + \delta + 2|y|)] \leq \epsilon$. \square

The following result gives a series of properties about the set $M_{K^*,C}$. Crucially, it shows that $M_{K^*,C}$ characterizes the optimal approximations of K^* as linear images of C :

Proposition 2.3 *Suppose that the assumptions (A1), (A2), (A3) hold. Then the set of minimizers $M_{K^*,C}$ defined in (7) is compact and non-empty. Moreover, we have*

$$\hat{A} \in M_{K^*,C} \iff \hat{A} \in \arg \min_{A \in L(\mathbb{R}^q, \mathbb{R}^d)} \rho_2(A(C), K^*).$$

Proof Define the event $G_{r,v} := \{(u, y) \mid \langle v, u \rangle \geq 1/2, |y| \leq r/4\}$ over $v \in S^{d-1}$. In addition, define the function $s(v) := \mathbb{P}[(u, y) \mid \langle v, u \rangle \geq 1/2]$. For every $r \geq 0$, consider the function $g_r(v) := \mathbb{P}[G_{r,v}]$. By noting that $g_r \leq g_{r'}$ whenever $r \leq r'$ (i.e., the sequence $\{g_r\}_{r \geq 0}$ is monotone increasing), $g_r(\cdot) \uparrow s(\cdot)$, and that $g_r(\cdot)$ is a continuous function over the compact domain S^{d-1} , we conclude that $g_r(\cdot)$ converges to $s(\cdot)$ uniformly. Thus there exists \hat{r} sufficiently large so that $\hat{r}^2 \mathbb{P}[G_{\hat{r},v}]/16 > \Phi_C(0, P_{K^*})$ for all $v \in S^{d-1}$.

Next, we show that $M_{K^*,C} \subseteq \hat{r}B_{\|\cdot\|_{C,2}}(0)$. Let $A \notin \hat{r}B_{\|\cdot\|_{C,2}}(0)$. Then, for such an A , there exists $\hat{x} \in C$ such that $\|A\hat{x}\|_2 > \hat{r}$. Define $\hat{v} = A\hat{x}/\|A\hat{x}\|_2$. We have

$$\Phi_C(A, P_{K^*}) \geq \mathbb{E}[\mathbf{1}(G_{\hat{r},\hat{v}})(h_C(A^T u) - y)^2] \geq \hat{r}^2 \mathbb{P}[G_{\hat{r},\hat{v}}]/16 > \Phi_C(0, P_{K^*}).$$

Here, $\mathbf{1}(G_{\hat{r}, \hat{\nu}})$ denotes the indicator function for the event $G_{\hat{r}, \hat{\nu}}$. As such, the above inequality implies that $A \notin M_{K^*, C}$. Therefore $M_{K^*, C} \subseteq \hat{r}B_{\|\cdot\|_{C,2}}(0)$, and hence $M_{K^*, C}$ is bounded.

By Proposition 2.1, the function $A \mapsto \Phi_C(A, P_{K^*})$ is continuous. As the minimizers of $\Phi_C(\cdot, P_{K^*})$, if they exist, must be contained in $\hat{r}B_{\|\cdot\|_{C,2}}(0)$, we can view $M_{K^*, C}$ as the set of minimizers of a continuous function restricted to the compact set $\hat{r}B_{\|\cdot\|_{C,2}}(0)$, and hence it is non-empty. Moreover, since $M_{K^*, C}$ is the pre-image of a closed set under a continuous function, it is also closed and thus compact.

By Fubini’s theorem, we have $\mathbb{E}[\varepsilon(h_C(K^*) - h_C(A^T u))] = \mathbb{E}_u[\mathbb{E}_\varepsilon[\varepsilon(h_C(K^*) - h_C(A^T u))]] = 0$. Hence $\Phi_C(A, P_{K^*}) = \mathbb{E}[(h_C(K^*) + \varepsilon - h_C(A^T u))^2] = \mathbb{E}[(h_C(K^*) - h_C(A^T u))^2] + \mathbb{E}[\varepsilon^2]$, from which the last assertion follows. \square

It follows from Proposition 2.3 that an optimal approximation of K^* as the projection of C always exists. In Sect. 3.1, we show that the estimators obtained using our method converge to an optimal approximation of K^* as a linear image of C if such an approximation is unique. While this is often the case in applications one encounters in practice, the following examples demonstrate that the uniqueness condition need not always hold:

Example Suppose K^* is the regular q -gon in \mathbb{R}^2 , and C is the spectraplex \mathcal{O}^2 . Then $M_{K^*, C}$ uniquely specifies an ℓ_2 -ball.

Example Suppose K^* is the unit ℓ_2 -ball in \mathbb{R}^2 , and C is the simplex Δ^q . Then the sets specified by the elements $M_{K^*, C}$ are not unique; they all correspond to a centered regular q -gon, but with an unspecified rotation.

A natural question then is to identify settings in which $M_{K^*, C}$ defines a unique set. Unfortunately, obtaining a complete characterization of this uniqueness property appears to be difficult due to the interplay between the invariances underlying the sets K^* and C . However, based on Proposition 2.3, we can provide a simple sufficient condition under which $M_{K^*, C}$ defines a unique set:

Corollary 2.4 *Assume that the conditions of Proposition 2.3 hold. Suppose further that we have $K^* = A^*(C)$ for some $A^* \in L(\mathbb{R}^q, \mathbb{R}^d)$. Then the set of minimizers $M_{K^*, C}$ described in (7) uniquely defines K^* ; i.e., $K^* = A(C)$ for all $A \in M_{K^*, C}$.*

Proof It is clear that $A^* \in M_{K^*, C}$. Note that $h_C(A^T u)$ is a continuous function of u over a compact domain for every $A \in L(\mathbb{R}^q, \mathbb{R}^d)$. Hence it follows that $\hat{A} \in M_{K^*, C}$ if and only if $h_C(A^{*T} u) = h_C(\hat{A}^T u)$ everywhere. By applying Proposition 2.3 and using the fact that a pair of compact convex sets that have the same support function must be equal, it follows that $K^* = \hat{A}(C)$ for all $\hat{A} \in M_{K^*, C}$. \square

2.2 Algebraic Aspects of Our Method

While the preceding subsection focused on conditions under which the set of minimizers $M_{K^*, C}$ specifies a unique convex set, the aim of the present section is to obtain

a more refined picture of the collection of linear maps in $M_{K^*,C}$. We begin by discussing the identifiability issues that arise in reconstructing a convex set by estimating a linear map via (5). Given a compact convex set C , let g be a linear transformation that preserves C ; i.e., $g(C) = C$. Then the linear map defined by Ag specifies the same convex set as A because $Ag(C) = A(g(C)) = A(C)$. As such, every linear map $A \in L(\mathbb{R}^q, \mathbb{R}^d)$ is a member of the equivalence class defined by

$$A \sim Ag, \quad g \in \text{Aut } C. \quad (9)$$

Here $\text{Aut } C$ denotes the subset of linear transformations that preserve C . When C is non-degenerate, the elements of $\text{Aut } C$ are invertible matrices and form a subgroup of $\text{GL}(q, \mathbb{R})$. As a result, the equivalence class $A \cdot \text{Aut } C := \{Ag \mid g \in \text{Aut } C\}$ specified by (9) can be viewed as the orbit of $A \in L(\mathbb{R}^q, \mathbb{R}^d)$ under (right) action of the group $\text{Aut } C$. In the sequel, we focus our attention on convex sets C for which the associated automorphism group $\text{Aut } C$ consists of isometries:

(A4) The automorphism group of C is a subgroup of the orthogonal group, i.e., $\text{Aut } C \triangleleft O(q, \mathbb{R})$.

This assumption leads to structural consequences that are useful in our analysis. In particular, as $\text{Aut } C$ can be viewed as a compact matrix Lie group, the orbit $A \cdot \text{Aut } C$ inherits structure as a smooth manifold of the ambient space $L(\mathbb{R}^q, \mathbb{R}^d)$. The assumption (A4) is satisfied for the choices of C that are primarily considered in this paper—the automorphism group of the simplex is the set of permutation matrices, and the automorphism group of the spectraplex is the set of linear operators specified as conjugation by an orthogonal matrix.

Based on this discussion, it follows that the space of linear maps $L(\mathbb{R}^q, \mathbb{R}^d)$ can be partitioned into orbits $A \cdot \text{Aut } C$. Further, the population loss $\Phi_C(\cdot, P_{K^*})$ is also invariant over orbits of A : for every $g \in \text{Aut } C$ we have that $h_C(A^T u) = h_C((Ag)^T u)$. Thus, the set of minimizers $M_{K^*,C}$ can also be partitioned into a union of orbits. Consequently, in our analysis in Sect. 3 we view the problem (5) as one of recovering an orbit rather than a particular linear map. The convergence results we obtain in Sect. 3 depend on the number of orbits in $M_{K^*,C}$, with sharper asymptotic distributional characterizations available when $M_{K^*,C}$ consists of a single orbit.

When $M_{K^*,C}$ specifies multiple convex sets, then $M_{K^*,C}$ clearly consists of multiple orbits; as an illustration, in the example in the previous subsection in which K^* is the unit ℓ_2 -ball in \mathbb{R}^2 and $C = \Delta^q$, the corresponding set $M_{K^*,C}$ is a union of multiple orbits in which each orbit specifies a unique rotation of the centered regular q -gon. However, even when $M_{K^*,C}$ specifies a unique convex set, it may still be the case that it consists of more than one orbit:

Example Suppose K^* is the interval $[-1, 1] \subset \mathbb{R}$ and $C = \Delta^3$. Then $M_{K^*,C}$ is a union of orbits, with an orbit specified as the set of all permutations of the vector $(-1, 1, \epsilon)$ for each $\epsilon \in [-1, 1]$. Nonetheless, $M_{K^*,C}$ specifies a unique convex set, namely, K^* .

More generally, it is straightforward to check that $M_{K^*,C}$ consists of a single orbit if K^* is a polytope with q extreme points and $C = \Delta^q$. The situation for linear images of non-polyhedral sets such as the spectraplices is much more delicate. One simple

instance in which $M_{K^*,C}$ consists of a single orbit is when K^* is the image under a bijective linear map A of \mathcal{O}^q . Our next result states an extension to convex sets that are representable as linear images of an appropriate slice of the outer product of cones of positive semidefinite matrices.

Proposition 2.5 *Let $C = \{X_1 \times \cdots \times X_k \mid X_i \in \mathbb{S}^{q_i}, X_i \geq 0, \sum_{i=1}^k \text{trace } X_i = 1\}$, and let $K^* = A^*(C) \subset \mathbb{R}^d$. Suppose that there is a collection of disjoint exposed faces $F_i \subseteq K^*$ such that: (i) $(A^*)^{-1}(F_i) \cap C$ is the i -th block $\{0 \times \cdots \times 0 \times X_i \times 0 \times \cdots \times 0 \mid X_i \in \mathbb{S}^{q_i}, X_i \geq 0, \text{trace } X_i = 1\} \subset C$, and (ii) $\dim F_i = \dim \mathcal{O}^{q_i}$. Then $M_{K^*,C}$ consists of a single orbit.*

Example By expressing $C = \Delta^q = \{X_1 \times \cdots \times X_q \mid X_i \in \mathbb{S}^1, X_i \geq 0, \sum_{i=1}^q \text{trace } X_i = 1\}$ and by considering K^* to be a polytope with q extreme points, Proposition 2.5 simplifies to our earlier remark noting that $M_{K^*,C}$ consists of a single orbit.

Example The nuclear norm ball $B_{\text{nuc}} := \{X \in \mathbb{S}^2 \mid \|X\|_{\text{nuc}} \leq 1\}$ is expressible as the linear image of $\mathcal{O}^2 \times \mathcal{O}^2$. The extreme points of B_{nuc} comprise two connected components of unit-norm rank-one matrices specified by $\{U^T E_{11} U \mid U \in \text{SO}(2, \mathbb{R})\}$ and $\{-U^T E_{11} U \mid U \in \text{SO}(2, \mathbb{R})\}$, where E_{11} is the 2×2 matrix with $(1, 1)$ -entry equal to one and other entries equal to zero. Furthermore, each connected component is isomorphic to \mathcal{O}^2 . It is straightforward to verify that the conditions of Proposition 2.5 hold for this instance, and thus $M_{B_{\text{nuc}}, \mathcal{O}^2 \times \mathcal{O}^2}$ consists of a single orbit.

The proof of Proposition 2.5 requires an impossibility result showing that a spectraplex cannot be expressed as the linear image of the outer product of finitely many smaller-sized spectraplices. The result follows as a consequence of a related result stated in terms of the cone of positive semidefinite matrices [1,26]. In the following, \mathbb{S}_+^q denotes the cone of $q \times q$ dimensional positive semidefinite matrices.

Proposition 2.6 *Suppose that $\mathbb{S}_+^q = A(\mathbb{S}_+^{q_1} \times \cdots \times \mathbb{S}_+^{q_k} \cap L)$ for some linear map A and some affine subspace L . Then $q \leq \max q_i$.*

Proposition 2.7 *Let $C = \{X_1 \times \cdots \times X_k \mid X_i \in \mathbb{S}_+^{q_i}, \sum_{i=1}^k \text{trace } X_i = 1\}$. Suppose $\mathcal{O}^q = A(C)$ for some A . Then $q \leq \max q_i$.*

Proof Express \mathbb{S}_+^q as

$$A \circ \Pi \left(\left\{ X_1 \times \cdots \times X_k \times t \mid X_i \in \mathbb{S}_+^{q_i}, t \in \mathbb{S}_+^1, \sum_{i=1}^k \text{trace } X_i = t \right\} \right),$$

where Π projects out the coordinate t . The result follows from Proposition 2.6. □

Lemma 2.8 *Let $K = A(C) \subset \mathbb{R}^d$ where $C \subset \mathbb{R}^q$ is compact convex, and suppose that $\dim K = \dim C$. If $K = \tilde{A}(C)$ for some $\tilde{A} \in L(\mathbb{R}^q, \mathbb{R}^d)$, then $\tilde{A} = Ag$ for some $g \in \text{Aut } C$.*

Proof Suppose that $0 \notin \text{aff } K$. By applying a suitable rotation, we may assume that K is contained in the first $\dim K$ dimensions. Then the maps A and \tilde{A} are of the form

$$A = \begin{pmatrix} A_1 \\ 0 \end{pmatrix}, \quad \tilde{A} = \begin{pmatrix} \tilde{A}_1 \\ 0 \end{pmatrix}, \quad A_1, \tilde{A}_1 \in L(\mathbb{R}^q, \mathbb{R}^q).$$

Since $\dim K = \dim C$, the map A_1 is invertible. Subsequently $A_1^{-1}A_1 \in \text{Aut } C$, and thus $\tilde{A}_1 = A_1g$ for some $g \in \text{Aut } C$.

The proof is similar for the case where $0 \in \text{aff } K$. The only necessary modification is that we embed K into \mathbb{R}^{d+1} via the set $\tilde{K} := \{(x, 1) \mid x \in K\}$, and we repeat the same sequence of steps with \tilde{K} in place of K . We omit the necessary details as they follow in a straightforward fashion from the previous case. \square

Proof of Proposition 2.5 Let $\tilde{A} \in M_{K^*, C}$. We show that \tilde{A} defines a one-to-one correspondence between the collection of faces $\{F_i\}_{i=1}^k$ and the collection of blocks $\{0 \times \cdots \times X_j \times \cdots \times 0 \mid X_j \in \mathcal{O}^{q_j}\}_{j=1}^k$ subject to the condition $\dim F_i = \dim \mathcal{O}^{q_i}$. We prove such a correspondence via an inductive argument beginning with the faces of largest dimensions.

We assume (without loss of generality) that $\dim F_1 = \dots = \dim F_{k'} > \dots$ and that $\dim \mathcal{O}^{q_1} = \dots = \dim \mathcal{O}^{q_{k'}} > \dots$. We further denote $q = q_1 = \dots = q_{k'}$. As F_i is an exposed face, the pre-image $\tilde{A}^{-1}(F_i) \cap C$ must be an exposed face of C , and thus is of the form $U_{i,1}X_{i,1}U'_{i,1} \times \cdots \times U_{i,k}X_{i,k}U'_{i,k}$, where $X_{i,j} \in \mathcal{O}^{q_{i,j}}$ for some $q_{i,j} \leq q_j$, and where $U_{i,j} \in \mathbb{R}^{q_{i,j} \times q_j}$ are partial orthogonal matrices. By Proposition 2.7, we have $\max_j q_{i,j} \geq q_i = q$. Subsequently, by noting that there are k' blocks with dimensions $q \times q$, that there are also k' faces F_i with $\dim F_i = q$, and that the faces F_i are disjoint, we conclude that each block in the collection $\{0 \times \cdots \times X_j \times \cdots \times 0 \mid X_j \in \mathcal{O}^{q_j}\}_{j=1}^k$ lies in the pre-image of a unique face F_i , $1 \leq i \leq k'$. By repeating the same sequence of arguments for the remaining faces of smaller dimensions, we establish a one-to-one correspondence between faces and blocks. Finally, we apply Lemma 2.8 to each face-block pair to conclude that, after accounting for permutations among blocks of the same size, the maps A_i^* and \tilde{A}_i are equivalent up to conjugation by an orthogonal matrix. The final assertion that $M_{K^*, C}$ consists of a single orbit is straightforward to establish. \square

2.3 Analytical Aspects of Our Method

In this third subsection, we describe some of the derivative computations that repeatedly play a role in our paper in our analysis, examples, and numerical experiments. Given a compact convex set C , the support function $h_C(\cdot)$ is differentiable at u if and only if

$$\arg \max_{x \in C} \langle x, u \rangle \tag{10}$$

is a singleton; the derivative in these cases is given by (10) (see [27, p. 47]). We denote the derivative of h_C at a differentiable u by $e_C(u) := \nabla_u(h_C(u))$.

Example Suppose $C = \Delta^q \subset \mathbb{R}^q$ is the simplex. The function $h_C(\cdot)$ is the maximum entry of the input vector, and it is differentiable at this point if and only if the maximum is unique with the derivative $e_C(\cdot)$ equal to the corresponding standard basis vector.

Example Suppose $C = \mathcal{O}^p \subset \mathbb{S}^p$ is the spectraplex. The function $h_C(\cdot)$ is the largest eigenvalue of the input matrix, and it is differentiable at this point if and only if the largest eigenvalue has multiplicity one with the derivative $e_C(\cdot)$ equal to the projector onto the corresponding one-dimensional eigenspace.

The following result gives a formula for the derivative of $\Phi_C(\cdot, P_{K^*})$.

Proposition 2.9 *Let P be a probability distribution over the measurement pairs (u, y) , and suppose that $\mathbb{E}_P[y^2] < \infty$. Let $A \in L(\mathbb{R}^q, \mathbb{R}^d)$ be a linear map such that $h_C(\cdot)$ is differentiable at $A^T u$ for P -a.e. u . Then the function $\Phi_C(\cdot, P)$ is differentiable with derivative*

$$2\mathbb{E}_P[(h_C(A^T u) - y)u \otimes e_C(A^T u)].$$

Proof Let $\lambda_C(\cdot, A, D)$ denote the remainder term satisfying

$$(h_C((A + D)^T u) - y)^2 = (h_C(A^T u) - y)^2 + \langle \nabla_A((h_C(A^T u) - y)^2), D \rangle + \lambda_C(\cdot, A, D)\|D\|_{C,2}.$$

Since the function $h_C(\cdot)$ is differentiable at $A^T u$ for P -a.e. u , we have $\lambda_C(\cdot, A, D) \rightarrow 0$ as $\|D\|_{C,2} \rightarrow 0$, P -a.e. Suppose D is in a bounded set. First, we can bound $|h_C((A + D)^T u) + h_C(A^T u) - 2y| \leq c_1(1 + |y|)$ for some constant c_1 . Second, by Lemma 2.2, we have the inequality $|h_C((A + D)^T u) - h_C(A^T u)| \leq \|D\|_{C,2}$. Third, by noting that $|h_C(A^T u) - y|$ can be bounded by $c_2(1 + |y|)$ for some constant c_2 , and by noting that the entries of the linear map $u \otimes e_C(A^T u)$ are uniformly bounded, we may bound $\|(h_C(A^T u) - y)u \otimes e_C(A^T u)\|_{C,2}$ by a function of the form $c_3(1 + |y|)$ for some constant c_3 . Subsequently we may bound

$$\begin{aligned} |\lambda_C(\cdot, A, D)| &\leq \|2(h_C(A^T u) - y)u \otimes e_C(A^T u)\|_{C,2} \\ &\quad + \frac{|h_C((A + D)^T u) - h_C(A^T u)| \cdot |h_C((A + D)^T u) + h_C(A^T u) - 2y|}{\|D\|_{C,2}} \\ &\leq c(1 + |y|) \end{aligned}$$

for some constant c . Since $\mathbb{E}_P[y^2] < \infty$, we have $\lambda_C(\cdot, A, D) \in \mathcal{L}^2(P)$, and hence $\lambda_C(\cdot, A, D) \in \mathcal{L}^1(P)$. The result follows from an application of the Dominated Convergence Theorem. □

It turns out to be considerably more difficult to compute an explicit expression of the second derivative of $\Phi_C(\cdot, P_{K^*})$. For this reason, our next result applies in a much more restrictive setting in comparison to Proposition 2.9.

Proposition 2.10 *Suppose that the underlying set $K^\star = A^\star(C)$ for some $A^\star \in L(\mathbb{R}^q, \mathbb{R}^d)$. In addition, suppose that the function $h_C(\cdot)$ is continuously differentiable at $A^{\star T}u$ for P_{K^\star} -a.e. u . Then the map $A \mapsto \Phi_C(A, P_{K^\star})$ is twice differentiable at A^\star and its second derivative is the operator $\Gamma \in L(L(\mathbb{R}^q, \mathbb{R}^d), L^\star(\mathbb{R}^q, \mathbb{R}^d))$ defined by*

$$\Gamma(D) = 2\mathbb{E}[\langle u \otimes e_C(A^{\star T}u), D \rangle u \otimes e_C(A^{\star T}u)]. \tag{11}$$

Proof To simplify notation, we denote the operator norm $\|\cdot\|_{C,2}$ by $\|\cdot\|$ in the remainder of the proof. By Proposition 2.9, the map $A \mapsto \Phi(A, P)$ is differentiable in an open neighborhood around A^\star with derivative $2(h_C(A^T u) - y)u \otimes e_C(A^T u)$. Hence to show that the map is twice differentiable with second derivative Γ , it suffices to show that

$$\begin{aligned} \lim_{\|D\| \rightarrow 0} \frac{1}{\|D\|} &\| \mathbb{E}[2(h_C((A^\star + D)^T u) - y)u \otimes e_C((A^\star + D)^T u)] \\ &- \mathbb{E}[2(h_C(A^{\star T}u) - y)u \otimes e_C(A^{\star T}u)] - \Gamma(D) \| = 0. \end{aligned}$$

First we note that every component of $\varepsilon(u)u \otimes e_C((A^\star + D)^T u)$ is integrable because $\mathbb{E}[\varepsilon(u)^2] < \infty$, and $u \otimes e_C((A^\star + D)^T u)$ is uniformly bounded. Hence by Fubini’s Theorem we have

$$\mathbb{E}[(h_C(A^{\star T}u) - y)u \otimes e_C((A^\star + D)^T u)] = \mathbb{E}_u[\mathbb{E}_{\varepsilon(u)}[-\varepsilon(u)u \otimes e_C((A^\star + D)^T u)]] = 0. \tag{12}$$

Similarly,

$$\mathbb{E}[(h_C(A^{\star T}u) - y)u \otimes e_C(A^{\star T}u)] = 0. \tag{13}$$

Second, by differentiability of the map $A \mapsto \Phi(A, P)$ at A^\star the limit

$$\lim_{\|D\| \rightarrow 0} \frac{\| \mathbb{E}[2(h_C((A^\star + D)^T u) - y) - 2(h_C(A^{\star T}u) - y) - 2\langle u \otimes e_C(A^{\star T}u), D \rangle] \|}{\|D\|}$$

is 0. By noting that every component of $u \otimes e_C((A^\star + D)^T u)$ is uniformly bounded, and an application of the Dominated Convergence Theorem, we have

$$\begin{aligned} \lim_{\|D\| \rightarrow 0} \frac{1}{\|D\|} &\| \mathbb{E}[(2(h_C((A^\star + D)^T u) - y) \\ &- 2(h_C(A^{\star T}u) - y) - 2\langle u \otimes e_C(A^{\star T}u), D \rangle)u \otimes e_C((A^\star + D)^T u)] \| = 0. \end{aligned} \tag{14}$$

Third, since $h_C(\cdot)$ is continuously differentiable at $A^{\star T}u$ for P -a.e. u , we have $e_C((A^\star + D)^T u) \rightarrow e_C(A^{\star T}u)$ as $\|D\| \rightarrow 0$, for P -a.e. u . By the Dominated Convergence Theorem we have $\mathbb{E}[e_C((A^\star + D)^T u)] \rightarrow \mathbb{E}[e_C(A^{\star T}u)]$ as $\|D\| \rightarrow 0$. It follows that

$$\lim_{\|D\| \rightarrow 0} \frac{1}{\|D\|} \left\| 2\mathbb{E}[\langle u \otimes e_C(A^*T u), D \rangle u \otimes e_C((A^* + D)^T u)] - 2\mathbb{E}[\langle u \otimes e_C(A^*T u), D \rangle u \otimes e_C(A^*T u)] \right\| = 0. \tag{15}$$

The result follows by summing the contributions from (14) and (15), as well as noting that the expressions in (12) and (13) vanish. \square

3 Main Results

In this section, we investigate the statistical aspects of minimizers of the optimization problem (5). Our objective in this section is to relate a sequence of minimizers $\{\hat{A}_n\}_{n=1}^\infty$ of $\Phi_C(\cdot, P_{n,K^*})$ to minimizers of $\Phi_C(\cdot, P_{K^*})$. Based on this analysis, we draw conclusions about properties of sequences of minimizers $\{\hat{K}_n^C\}_{n=1}^\infty$ of the problem (2). In establishing various convergence results, we rely on an important property of the Hausdorff distance, namely that it defines a metric over collections of non-empty compact sets; therefore, the collection of all orbits $\{A \cdot \text{Aut } C \mid A \in L(\mathbb{R}^q, \mathbb{R}^d)\}$ endowed with the Hausdorff distance defines a metric space.

Our results provide progressively sharper recovery guarantees based on increasingly stronger assumptions. Specifically, Sect. 3.1 focuses on conditions under which a sequence of minimizers $\{\hat{K}_n^C\}_{n=1}^\infty$ of (2) converges to K^* ; this result relies only on the fact that the optimal approximation of K^* by a convex set specified by an element of $M_{K^*,C}$ is unique (see Sect. 2.1 for the relevant discussion). Next, Sect. 3.2 gives a limiting distributional characterization of the sequence $\{\hat{K}_n^C\}_{n=1}^\infty$ based on an asymptotic normality analysis of the sequence $\{\hat{A}_n\}_{n=1}^\infty$; among other assumptions, this analysis relies on the stronger requirement that $M_{K^*,C}$ consists of a single orbit. Finally, based on additional conditions on the facial structure of K^* , we describe in Sect. 3.3 how the sequence $\{\hat{K}_n^C\}_{n=1}^\infty$ preserves various attributes of the face structure of K^* .

3.1 Strong Consistency

We describe conditions for convergence of a sequence of minimizers $\{\hat{K}_n^C\}_{n=1}^\infty$ of (2). Our main result essentially states that such a sequence converges to an optimal ρ_2 approximation of K^* as a linear image of C , provided that such an approximation is unique:

Theorem 3.1 *Suppose that the assumptions (A1)–(A4) hold. Let $\{\hat{A}_n\}_{n=1}^\infty$ be a sequence of minimizers of the empirical loss function $\Phi_C(\cdot, P_{n,K^*})$ with the corresponding reconstructions given by $\hat{K}_n^C = \hat{A}_n(C)$. We have that $\text{dist}(\hat{A}_n, M_{K^*,C}) \rightarrow 0$ a.s. and*

$$\inf_{A \in M_{K^*,C}} d_H(\hat{A}_n \cdot \text{Aut } C, A \cdot \text{Aut } C) \rightarrow 0 \quad \text{a.s.}$$

As a consequence, if $M_{K^*,C}$ specifies a unique set—there exists $\hat{K} \subset \mathbb{R}^d$ such that $\hat{K} = A(C)$ for all $A \in M_{K^*,C}$ —then $\rho_p(\hat{K}_n^C, \hat{K}) \rightarrow 0$ a.s. for $1 \leq p \leq \infty$. Further, if $K^* = A^*(C)$ for some linear map $A^* \in L(\mathbb{R}^q, \mathbb{R}^d)$, then $\rho_p(\hat{K}_n^C, K^*) \rightarrow 0$ a.s. for $1 \leq p \leq \infty$.

When $M_{K^*,C}$ defines multiple sets, our result does not imply convergence of the sequence $\{\hat{K}_n^C\}_{n=1}^\infty$. Rather, we obtain the weaker consequence that the sequence $\{\hat{K}_n^C\}_{n=1}^\infty$ eventually becomes arbitrarily close to the collection $\{A(C) \mid A \in M_{K^*,C}\}$.

Example Suppose K^* is the unit ℓ_2 -ball in \mathbb{R}^2 , and $C = \Delta^q$. The optimal ρ_2 approximation is the regular q -gon with an unspecified rotation. The sequence $\{\hat{K}_n^C\}_{n=1}^\infty$ does not have a limit; rather, there is a sequence $\{g_n\}_{n=1}^\infty \subset \text{SO}(2, \mathbb{R})$ such that $g_n \hat{K}_n^C$ converges to a centered regular q -gon (with fixed orientation) a.s.

The proof of Theorem 3.1 comprises two parts. First, we show that there exists a ball in $L(\mathbb{R}^q, \mathbb{R}^d)$ such that \hat{A}_n belongs to this ball for all sufficiently large n a.s. Second, we appeal to the following uniform convergence result. The structure of our proof is similar to that of a corresponding result for K -means clustering (see the main theorem in [22]).

Lemma 3.2 *Let $U \subset L(\mathbb{R}^q, \mathbb{R}^d)$ be bounded and suppose the set $C \subset \mathbb{R}^q$ satisfies assumption (A3). Let P be a probability distribution over the measurement pairs $(u, y) \subset S^{d-1} \times \mathbb{R}$ satisfying $\mathbb{E}_P[y^2] < \infty$, and let P_n be the empirical measure corresponding to drawing n i.i.d. observations from the distribution P . Consider the collection of functions $\mathcal{G} := \{(h_C(A^T u) - y)^2 \mid A \in U\}$ in the variables (u, y) . Then $\sup_{g \in \mathcal{G}} |\mathbb{E}_{P_n}[g] - \mathbb{E}_P[g]| \rightarrow 0$ as $n \rightarrow \infty$ a.s.*

The proof of Lemma 3.2 follows from an application of the following uniform strong law of large numbers (SLLN) [23, Thm. 3, p. 8].

Theorem 3.3 (Uniform SLLN) *Let Q be a probability measure and let Q_n be the corresponding empirical measure. Let \mathcal{G} be a collection of Q -integrable functions. Suppose that for every $\epsilon > 0$ there exists a finite collection of functions \mathcal{G}_ϵ such that for every $g \in \mathcal{G}$ there exist $\bar{g}, \underline{g} \in \mathcal{G}_\epsilon$ satisfying*

- (i) $g \leq \bar{g}$, and
- (ii) $\mathbb{E}_Q[\bar{g} - \underline{g}] < \epsilon$.

Then $\sup_{g \in \mathcal{G}} |\mathbb{E}_{Q_n}[g] - \mathbb{E}[g]| \rightarrow 0$ a.s.

Proof of Lemma 3.2 Based on Theorem 3.3, it suffices to construct the finite collection of functions \mathcal{G}_ϵ . Pick r sufficiently large so that $U \subset rB_{\|\cdot\|_{C,2}}(0)$. Let \mathcal{D}_δ be a δ -cover for U in the $\|\cdot\|_{C,2}$ -norm, where δ is chosen so that $4\delta\mathbb{E}_P[r + |y|] \leq \epsilon$. We define $\mathcal{G}_\epsilon := \{(|h_C(A^T u) - y| - \delta)_+^2\}_{A \in \mathcal{D}_\delta} \cup \{(|h_C(A^T u) - y| + \delta)^2\}_{A \in \mathcal{D}_\delta}$.

We proceed to verify (i) and (ii). Let $g = (h_C(A^T u) - y)^2 \in \mathcal{G}$ be arbitrary. Let $A_0 \in \mathcal{D}_\delta$ be such that $\|A - A_0\|_{C,2} \leq \delta$. Define $\underline{g} = (|h_C(A_0^T u) - y| - \delta)_+^2$ and $\bar{g} = (|h_C(A_0^T u) - y| + \delta)^2$. It follows that $\underline{g} \leq g \leq \bar{g}$, which verifies (i). Next, we have $\mathbb{E}[\bar{g} - \underline{g}] \leq 4\delta\mathbb{E}[|h_C(A_0^T u) - y|] \leq 4\delta\mathbb{E}[r + |y|] \leq \epsilon$, which verifies (ii). \square

Proof of Theorem 3.1 To simplify notation in the following proof, we denote $B := B_{\|\cdot\|_{C,2}}(0)$. First, we recall the definition of the event $G_{r,v}$ and the function $s(v)$ from the proof of Proposition 2.3. Using a sequence of arguments identical as in the proof of Proposition 2.3, it follows that there exists \hat{r} sufficiently large so that $\hat{r}^2\mathbb{P}[G_{\hat{r},v}]/16 > \Phi_C(0, P_{K^*})$ for all $v \in S^{d-1}$. We claim that $\hat{A}_n \in \hat{r}B$ eventually a.s. We prove this assertion via contradiction. Suppose on the contrary that $\hat{A}_n \notin \hat{r}B$ i.o. For every $\hat{A}_n \notin \hat{r}B$, there exists $\hat{x}_n \in C$ such that $\|\hat{A}_n\hat{x}_n\| > \hat{r}$. The sequence of unit-norm vectors $\hat{A}_n\hat{x}_n/\|\hat{A}_n\hat{x}_n\|_2$, defined over the subset of indices n such that $\hat{A}_n \notin \hat{r}B$, contains a convergent subsequence whose limit point is $\hat{v} \in S^{d-1}$. Then

$$\begin{aligned} \limsup_n \Phi_C(\hat{A}_n, P_{n,K^*}) &= \limsup_n \mathbb{E}_{P_{n,K^*}}[(h_C(\hat{A}_n^T u) - y)^2] \\ &\geq \limsup_n \mathbb{E}_{P_{n,K^*}}[\mathbf{1}(G_{\hat{r},\hat{v}})(h_C(\hat{A}_n^T u) - y)^2] \\ &\geq \limsup_n \mathbb{E}_{P_{n,K^*}}\left[\frac{r^2\mathbf{1}(G_{\hat{r},\hat{v}})}{16}\right] \\ &\geq \lim_n \frac{r^2\mathbb{P}_{P_{n,K^*}}[G_{\hat{r},\hat{v}}]}{16} = \frac{r^2\mathbb{P}_{P_{K^*}}[G_{\hat{r},\hat{v}}]}{16} > \Phi_C(0, P_{K^*}). \end{aligned}$$

Here, the last equality follows from the SLLN. This implies $\Phi_C(\hat{A}_n, P_{n,K^*}) > \Phi_C(0, P_{n,K^*})$ i.o., which contradicts the minimality of \hat{A}_n . Hence $\hat{A}_n \in \hat{r}B$ eventually a.s.

Second, we show that $\text{dist}(\hat{A}_n, M_{K^*,C}) \rightarrow 0$ a.s. It suffices to show that $\hat{A}_n \in U$ eventually a.s., where U is any open set containing $M_{K^*,C}$. Let $\hat{A} \in M_{K^*,C}$ be arbitrary. By Proposition 2.1, the function $A \mapsto \Phi_C(A, P_{K^*})$ is continuous. By noting that the set of minimizers of $\Phi_C(\cdot, P_{K^*})$ is compact from Proposition 2.3, we can pick $\epsilon > 0$ sufficiently small so that $\{A \mid \Phi_C(A, P_{K^*}) < \Phi_C(\hat{A}, P_{K^*}) + \epsilon\} \subset U$. Next, since \hat{A}_n is defined as the minimizer of an empirical sum, we have $\Phi_C(\hat{A}_n, P_{n,K^*}) \leq \Phi_C(\hat{A}, P_{n,K^*})$ for all n . By applying Lemma 3.2 with the choice of $P = P_{K^*}$, we have $\Phi_C(\hat{A}_n, P_{n,K^*}) \rightarrow \Phi_C(\hat{A}_n, P_{K^*})$, and $\Phi_C(\hat{A}, P_{n,K^*}) \rightarrow \Phi_C(\hat{A}, P_{K^*})$, both uniformly and in the a.s. sense. Subsequently, by combining the previous two conclusions, we have $\Phi_C(\hat{A}_n, P_{K^*}) < \Phi_C(\hat{A}, P_{K^*}) + \epsilon$ eventually, for any $\epsilon > 0$. This proves that $\text{dist}(\hat{A}_n, M_{K^*,C}) \rightarrow 0$ a.s.

Third, we conclude that

$$\inf_{A \in M_{K^*,C}} d_H(\hat{A}_n \cdot \text{Aut } C, A \cdot \text{Aut } C) \rightarrow 0 \quad \text{a.s.} \tag{16}$$

Fix an integer n , and let $t_n = \text{dist}(\hat{A}_n, M_{K^*,C})$. Since $M_{K^*,C}$ is compact, we may pick $\bar{A} \in M_{K^*,C}$ such that $\|\hat{A}_n - \bar{A}\|_F = t_n$. Given $A \in \hat{A}_n \cdot \text{Aut } C$, we have $A = \hat{A}_n g$ for some $g \in \text{Aut } C$. Then $\bar{A}g \in \bar{A} \cdot \text{Aut } C$, and since g is an isometry by assumption (A4), we have $\|\hat{A}_n g - \bar{A}g\|_F = \|\hat{A}_n - \bar{A}\|_F = t_n$. This implies that $d_H(\hat{A}_n \cdot \text{Aut } C, \bar{A} \cdot \text{Aut } C) \leq t_n$. Since $t_n \rightarrow 0$ as $n \rightarrow \infty$, (16) follows. \square

3.2 Asymptotic Normality

In our second main result, we characterize the limiting distribution of a sequence of estimates $\{\hat{K}_n^C\}_{n=1}^\infty$ corresponding to minimizers of (2) by analyzing an associated sequence of minimizers of (5). Specifically, we show under suitable conditions that the estimation error in the sequence of minimizers of the empirical loss (5) is asymptotically normal. After developing this theory, we illustrate in Sect. 3.2.1 through a series of examples the asymptotic behavior of the set \hat{K}_n^C , highlighting settings in which \hat{K}_n^C converges, as well as situations in which our asymptotic normality characterization fails due to the requisite assumptions not being satisfied. Our result relies on two key ingredients, which we describe next.

The first set of requirements pertains to non-degeneracy of the function $\Phi_C(\cdot, P_{K^*})$. First, we require that the minimizers of $\Phi_C(\cdot, P_{K^*})$ constitute a unique orbit under the action of the automorphism group of the set C ; this guarantees the existence of a convergent sequence of minimizers of the empirical losses $\Phi_C(\cdot, P_{n,K^*})$, which is necessary to provide a Central Limit Theorem (CLT) type of characterization. Second, we require the function $\Phi_C(\cdot, P_{K^*})$ to be twice differentiable at a minimizer with a positive definite Hessian (modulo invariances due to $\text{Aut } C$); such a condition allows us to obtain a quadratic approximation of $\Phi_C(\cdot, P_{K^*})$ around a minimizer \hat{A} , and to subsequently compute first-order approximations of minimizers of the empirical losses $\Phi(\cdot, P_{n,K^*})$. These conditions lead to a geometric characterization of confidence regions of the extreme points of the limit of the sequence $\{\hat{K}_n^C\}_{n=1}^\infty$.

Our second main assumption centers around the collection of sets that serves as the constraint in the optimization problem (2). Informally, we require that this collection is not “overly complex,” so that we can appeal to a suitable CLT. The field of empirical processes provides the tools necessary to formalize matters. Concretely, as our estimates are obtained via minimization of an empirical process, we require that the following divided differences of the loss function are well controlled:

$$d_{C,A_1,A_2}(u, y) = \frac{(h_C(A_1^T u) - y)^2 - (h_C(A_2^T u) - y)^2}{\|A_1 - A_2\|_F}.$$

In particular, a natural assumption is that the *graph* associated to these divided differences, indexed over a collection centered at \hat{A} , is of suitably “bounded complexity”:

$$\left\{ \{(u, y, s) \mid d_{C,A,\hat{A}}(u, y) \geq s \geq 0 \text{ or } d_{C,A,\hat{A}}(u, y) \leq s \leq 0\} \mid A \in \mathcal{B}_{\|\cdot\|_F}(\hat{A}) \setminus \{\hat{A}\} \right\}. \quad (17)$$

A powerful framework to quantify the ‘richness’ of such collections is based on the notion of a *Vapnik–Chervonenkis* (VC) class [33]. VC classes describe collections of subsets with bounded complexity, and they feature prominently in the field of statistical learning theory, most notably in conditions under which generalization of a learning algorithm is possible.

Definition (Vapnik–Chervonenkis class) Let \mathcal{F} be a collection of subsets of a set F . A finite set D is said to be *shattered* by F if for every $A \subset D$ there exists $G \in \mathcal{F}$ such

that $A = G \cap D$. The collection \mathcal{F} is said to be a VC class if there is a finite k such that all sets with cardinality greater than k cannot be shattered by \mathcal{F} .

Whether or not the collection (17) is VC depends on the particular choice of C . If C is chosen to be either a simplex or a spectraplex then such a property is indeed satisfied—see Sect. 3.2.2 for further details. Our analysis relies on a result by Stengle and Yukich showing that certain collections of sets admitting semialgebraic representations are VC [31]. We are now in a position to state our main result of this section:

Theorem 3.4 *Suppose that the assumptions (A1)–(A4) hold. Suppose that there exists $\hat{A} \in L(\mathbb{R}^q, \mathbb{R}^d)$ such that the set of minimizers $M_{K^*,C} = \hat{A} \cdot \text{Aut } C$ (i.e., $M_{K^*,C}$ consists of a single orbit), the function $h_C(\cdot)$ is differentiable at $\hat{A}^T u$ for P_{K^*} -a.e. u , the function $\Phi_C(\cdot, P_{K^*})$ is twice differentiable at \hat{A} , and the associated Hessian Γ at \hat{A} is positive definite restricted to the subspace T^\perp , i.e., $\Gamma|_{T^\perp} \succ 0$; here $T := T_{\hat{A}} M_{K^*,C}$ denotes the tangent space of $M_{K^*,C}$ at \hat{A} . In addition, suppose that the collection of sets specified by (17) forms a VC class. Let $\tilde{A}_n \in \arg \min_A \Phi_C(A, P_{n,K^*})$, $n \in \mathbb{N}$, be a sequence of minimizers of the empirical loss function, and let $\hat{A}_n \in \arg \min_{A \in \hat{A}_n \cdot \text{Aut } C} \| \hat{A} - A \|_F$, $n \in \mathbb{N}$, specify an equivalent sequence defined with respect to the minimizer \hat{A} of the population loss. Setting $\nabla = \nabla_A((h_C(A^T u) - y)^2)$, we have that*

$$\sqrt{n}(\tilde{A}_n - \hat{A}) \xrightarrow{D} \mathcal{N}(0, (\Gamma|_{T^\perp})^{-1}(\mathbb{E}_{P_{K^*}}[\nabla \otimes \nabla|_{A=\hat{A}}]|_{T^\perp})(\Gamma|_{T^\perp})^{-1}).$$

The proof of this theorem relies on ideas from the literature of empirical processes, which we describe next in a general setting.

Proposition 3.5 *Suppose P is a probability measure over \mathbb{R}^q and P_n is the empirical measure corresponding to n i.i.d. draws from P . Suppose $f(\cdot, \cdot): \mathbb{R}^q \times \mathbb{R}^p \rightarrow \mathbb{R}$ is a Lebesgue-measurable function such that $f(\cdot, t): \mathbb{R}^q \rightarrow \mathbb{R}$ is P -integrable for all $t \in \mathbb{R}^p$. Denote the expectations $F(t) = \mathbb{E}_P[f(\cdot, t)]$ and $F_n(t) = \mathbb{E}_{P_n}[f(\cdot, t)]$. Let $\hat{t} \in \arg \min_t F(t)$ be a minimizer of the population loss, let $\{\hat{t}_n\}_{n=1}^\infty$, $\hat{t}_n \in \arg \min_t F_n(t)$, be a sequence of empirical minimizers, and let $f(\cdot, t) = f(\cdot, \hat{t}) + \langle t - \hat{t}, \Delta(\cdot) \rangle + \|t - \hat{t}\|_2 r(\cdot, t)$ denote the linearization of $f(\cdot, t)$ about \hat{t} (we assume that $f(\cdot, t)$ is differentiable with respect to t at \hat{t} P -a.e., with derivative denoted by $\Delta(\cdot)$). Let $\mathcal{D} = \{d_{t,\hat{t}}(\cdot) \mid t \in B_{\|\cdot\|_2}(\hat{t}) \setminus \{\hat{t}\}\}$, where $d_{t_1,t_2}(\cdot) = (f(\cdot, t_1) - f(\cdot, t_2))/\|t_1 - t_2\|_2$, denote the collection of divided differences. Suppose*

- (i) $\hat{t}_n \rightarrow \hat{t}$ a.s.,
- (ii) the Hessian $\nabla^2 := \nabla^2 F(t)|_{t=\hat{t}}$ about \hat{t} is positive definite,
- (iii) the collection of sets $\{\{(\cdot, s) \mid d_{t,\hat{t}}(\cdot) \geq s \geq 0 \text{ or } d_{t,\hat{t}}(\cdot) \leq s \leq 0\} \mid t \in B_{\|\cdot\|_2}(\hat{t}) \setminus \{\hat{t}\}\}$ form a VC class, and
- (iv) there is a function $\bar{d}(\cdot): \mathbb{R}^q \rightarrow \mathbb{R}$ such that $|d(\cdot)| \leq \bar{d}(\cdot)$ for all $d \in \mathcal{D}$, $|\langle \Delta(\cdot), e \rangle| \leq \bar{d}(\cdot)$ for all unit-norm vectors e , $\bar{d}(\cdot) > 0$, and $\bar{d}(\cdot) \in \mathcal{L}^2(P)$.

Then $\sqrt{n}(\hat{t}_n - \hat{t}) \xrightarrow{D} \mathcal{N}(0, (\nabla^2)^{-1}[(\mathbb{E}_P[\Delta \Delta^T]) - (\mathbb{E}_P[\Delta])(\mathbb{E}_P[\Delta])^T](\nabla^2)^{-1})$.

The proof of Proposition 3.5 is based on a series of ideas developed in [23, Ch. VII] (see in particular Example 19). These rely on computing entropy numbers for certain

function classes. More formally, given a probability measure P and a collection of functions \mathcal{F} with each member being a mapping from \mathbb{R}^q to \mathbb{R} , we define $\eta(\epsilon, P, \mathcal{F})$ to be the size of the smallest ϵ -cover of \mathcal{F} in the $\mathcal{L}^2(P)$ -distance. As these steps are built on substantial background material from [23], we state the key conceptual arguments in the following proof and refer the reader to specific pointers in the literature for further details.

Proof of Proposition 3.5 To simplify notation we denote $B_{\hat{t}} := B_{\|\cdot\|_2}(\hat{t}) \setminus \{\hat{t}\}$. The graphs of $\mathcal{D} := \{d_{t,\hat{t}}(\cdot) \mid t \in B_{\hat{t}}\}$ form a VC class by assumption. By [23, Cor. 17, p. 20], these graphs have polynomial discrimination (see [23, Ch. II]). In addition, the collection \mathcal{D} have an enveloping function \bar{d} by assumption. Hence by [23, Lem. 36, p. 34], there exist constants $\alpha_{\mathcal{D}}, \beta_{\mathcal{D}}$ such that $\eta(\delta(\mathbb{E}_{P_n}[\bar{d}^2])^{1/2}, P_n, \mathcal{D}) \leq \alpha_{\mathcal{D}}(1/\delta)^{\beta_{\mathcal{D}}}$, for all $0 < \delta \leq 1$ and all n .

We obtain a similar bound for graphs of $\mathcal{E} := \{\langle \Delta(\cdot), (t - \hat{t})/\|t - \hat{t}\|_2 \rangle \mid t \in B_{\hat{t}}\}$. First, we note that \mathcal{E} is a subset of a finite dimensional vector space. By [17, Lem. 9.6, p. 159], the collection of subgraphs $\{(t, s) \mid f(t) \leq s \mid f \in \mathcal{E}\}$ forms a VC class. Then, by noting that the singletons $\{(t, s) \mid s \leq 0\}$ form a VC class, and that the collection of sets formed by taking intersections of members of two VC classes is also a VC class (see [17, Lem. 9.7, p. 159]), we conclude that the collection of sets $\{(t, s) \mid f(t) \leq s \leq 0 \mid f \in \mathcal{E}\}$ is a VC class. A similar sequence of steps shows that the collection of sets $\{(t, s) \mid f(t) \geq s \geq 0 \mid f \in \mathcal{E}\}$ also forms a VC class. The collection of graphs of \mathcal{E} is the union of the previous two collections. Hence by [17, Lem. 9.7, p. 159] it is also a VC class. Subsequently, by applying the same sequence of steps as we did for \mathcal{D} , there exists constants $\alpha_{\mathcal{E}}, \beta_{\mathcal{E}}$ such that $\eta(\delta(\mathbb{E}_{P_n}[\bar{d}^2])^{1/2}, P_n, \mathcal{E}) \leq \alpha_{\mathcal{E}}(1/\delta)^{\beta_{\mathcal{E}}}$, for all $0 < \delta \leq 1$ and all n .

Next, consider the collection of functions $\mathcal{F} := \{r(\cdot, t) \mid t \in B_{\hat{t}}\}$. We have $r(\cdot, t) = d_{t,\hat{t}}(\cdot) + \langle \Delta(\cdot), (t - \hat{t})/\|t - \hat{t}\|_2 \rangle$, and hence every element in \mathcal{F} is expressible as a sum of functions in \mathcal{D} and \mathcal{E} respectively. Given $(\delta/\sqrt{2})$ -covers for \mathcal{D} and \mathcal{E} in any \mathcal{L}^2 -distance, one can show via the AM-GM inequality that the union of these covers forms a δ -cover for \mathcal{F} . Subsequently, we conclude that there exists constants α and β such that $\eta(\delta(\mathbb{E}_{P_n}[\bar{d}^2])^{1/2}, P_n, \mathcal{F}) \leq \alpha(1/\delta)^\beta$.

The remaining sequence of arguments is identical to those in [23, Ch. VII]. Our bound on the quantity $\eta(\delta(\mathbb{E}_{P_n}[\bar{d}^2])^{1/2}, P_n, \mathcal{F})$ implies that there exists θ such that $\int_0^\theta (\eta(t, P_n, \mathcal{F})/t)^{1/2} dt < \epsilon$ for all n and all $\epsilon > 0$. We apply [23, Lem. 15, p. 150] to obtain the following *stochastic equicontinuity* property: given $\gamma > 0$ and $\epsilon > 0$, there exists θ such that

$$\limsup_n \mathbb{P} \left[\sup_{f_1, f_2 \in \mathcal{F}: \mathbb{E}_P[(f_1 - f_2)^2] \leq \theta} |\mathbb{E}_{E_n}[f_1 - f_2]| > \gamma \right] < \epsilon.$$

Here, E_n denotes the signed measure $\sqrt{n}(P_n - P)$. One can check measurability of the inner supremum using the notion of *permissibility*—see [23, App. C] (in essence, we simply require f to be measurable and the index t to reside in an Euclidean space). Finally we apply [23, Thm. 5, p. 141] to conclude the result. \square

Proof of Theorem 3.4 The first step is to verify that the sequence $\{\tilde{A}_n\}_{n=1}^\infty$ quotients out the appropriate equivalences. Since $\tilde{A}_n \in \hat{A}_n \cdot \text{Aut } C$, we have $\tilde{A}_n = \hat{A}_n g_n$ for an isometry g_n . Subsequently we have $\|\tilde{A}_n - \hat{A}\|_F = \|\hat{A}_n g_n - \hat{A}\|_F = \|\hat{A}_n - \hat{A} g_n^{-1}\|_F \leq d_H(\hat{A}_n, M_{K^*}, C)$. Following the conclusions of Theorem 3.1, we have $\tilde{A}_n \rightarrow \hat{A}$ a.s. Furthermore, as a consequence of the optimality conditions in the definition of \tilde{A}_n , we also have $\tilde{A}_n - \hat{A} \in T^\perp$.

The second step is to apply Proposition 3.5 to the sequence $\{\tilde{A}_n\}_{n=1}^\infty$ with $(h_C(A^T u) - y)^2$ as the choice of loss function $f(\cdot, t)$, P_{K^*} as the probability measure P , (u, y) as the argument, and A as the index t . First, the measurability of the loss as a function in A and (u, y) is straightforward to establish. Second, the differentiability of the loss function at \hat{A} for P_{K^*} -a.e. (u, y) follows from the assumption that $h_C(\cdot)$ is differentiable at $\hat{A}^T u$ for P_{K^*} -a.e. u . Third, we have shown that $\tilde{A}_n \rightarrow \hat{A}$ a.s. in the above. Fourth, the Hessian $\Gamma|_{T^\perp}$ is positive definite by assumption. Fifth, the graphs of $\{d_{C,A,\hat{A}}(u, y) \mid A \in B_{\|\cdot\|_F}(\hat{A}) \setminus \{\hat{A}\}\}$ form a VC class by assumption. Sixth, we need to show the existence of an appropriate function $d(\cdot)$ to bound the divided differences $d_{C,A,\hat{A}}(\cdot)$ and the inner products $\langle \nabla(\cdot)|_{A=\hat{A}}, E \rangle$, where $\|E\|_F = 1$. In the former case, we note that $|(h_C(A_1^T u) - y) + (h_C(A_2^T u) - y)| \leq \|A_1 - A_2\|_{C,2} \leq c_1 \|A_1 - A_2\|_F$ for some $c_1 > 0$; here, the first inequality follows from Lemma 2.2 and the second inequality follows from the equivalence of norms. Then, by noting that $A \in B_{\|\cdot\|_F}(\hat{A}) \setminus \hat{A}$ is bounded, the expression $|(h_C(A^T u) - y) + (h_C(\hat{A}^T u) - y)|$ is bounded above by a function of the form $c_2(1 + |y|)$. By expanding the divided difference expression and by combining the previous two bounds, one can show that $|d_{C,A,\hat{A}}(\cdot)| \leq c_3(1 + |y|)$ for some $c_3 > 0$. In the latter case, the derivative is given by $2(h_C(\hat{A}^T u) - y)u \otimes e_C(\hat{A}^T u)$. By noting that $h_C(\hat{A}^T u)$, u , and $e_C(\hat{A}^T u)$ are uniformly bounded over $u \in S^{d-1}$, and by performing a sequence of elementary bounds, one can show that $\langle \nabla, E \rangle$ is bounded above by $c_4(1 + |y|)$ uniformly over all unit-norm E . We pick $\bar{d}(\cdot)$ to be $c(1 + |y|)$, where $c = \max\{1, c_3, c_4\}$. Then $\bar{d}(\cdot) > 0$ by construction and furthermore, $\bar{d} \in L^2(P_{K^*})$ as $\mathbb{E}_{P_{K^*}}[e^2] < \infty$.

Finally, the result follows from an application of Proposition 3.5. □

This result gives an asymptotic normality characterization corresponding to a sequence of minimizers $\{\hat{A}_n\}_{n=1}^\infty$ of the empirical losses $\Phi_C(\cdot, P_{n,K^*})$. In the next result, we specialize this result to the setting in which the underlying set K^* is in fact expressible as a projection of C , i.e., $K^* = A^*(C)$ for some $A^* \in L(\mathbb{R}^q, \mathbb{R}^d)$. This specialization leads to a particularly simple formula for the asymptotic error covariance, and we demonstrate its utility in the examples in Sect. 3.2.1.

Corollary 3.6 *Suppose that the conditions of Theorem 3.4 and of Proposition 2.10 hold. Using the notation of Theorem 3.4, we have that $\mathbb{E}[\nabla \otimes \nabla|_{A=A^*}] = 2\sigma^2 \Gamma$ with Γ given by (11). In particular, the conclusion of Theorem 3.4 simplifies to $\sqrt{n}(\tilde{A}_n - \hat{A}) \xrightarrow{D} \mathcal{N}(0, 2\sigma^2(\Gamma|_{T^\perp})^{-1})$.*

Proof One can check that $\nabla_A((h_C(A^T u) - y)^2)|_{A=A^*} = -2\epsilon u \otimes e_C(A^* T u)$, from which we have that $\mathbb{E}[\nabla \otimes \nabla|_{A=A^*}] = 2\sigma^2 \Gamma$. This concludes the result. □

3.2.1 Examples

Here we give examples that highlight various aspects of the theoretical results described previously. In all of our examples, the noise $\{\varepsilon^{(i)}\}_{i=1}^n$ is given by i.i.d. centered Gaussian random variables with variance σ^2 . We begin with an illustration in which the assumptions of Theorem 3.4 are not satisfied and the asymptotic normality characterization fails to hold:

Example Let $K^* := \{0\} \subset \mathbb{R}$ be a singleton. As $S^0 \cong \{-1, 1\}$, the random variables $u^{(i)}$ are ± 1 u.a.r. Further, $h_{K^*}(u) = 0$ for all u and the support function measurements are simply $y^{(i)} = \varepsilon^{(i)}$ for all $i = 1, \dots, n$. For C being either a simplex or a spectraplex, the set $M_{K^*,C} = \{0\} \subset L(\mathbb{R}^q, \mathbb{R})$ is a singleton consisting only of the zero map. First, we consider fitting K^* with the choice of $C = \Delta^1 \subset \mathbb{R}^1$. Then we have $\hat{A}_n = (1/n) \sum_{i=1}^n \varepsilon^{(i)} u^{(i)}$, from which it follows that $\sqrt{n}(\hat{A}_n - 0)$ is normally distributed with mean zero and variance σ^2 —this is in agreement with Theorem 3.4. Second, we consider fitting K^* with the choice $C = \Delta^2 \subset \mathbb{R}^2$. Define $U_- = \{i \mid u_i = -1\}$ and $U_+ = \{i \mid u_i = 1\}$, and

$$\alpha_- = -\frac{1}{|U_-|} \sum_{i \in U_-} \varepsilon^{(i)} \quad \text{and} \quad \alpha_+ = \frac{1}{|U_+|} \sum_{j \in U_+} \varepsilon^{(j)}.$$

Then $\hat{K}_n^C = \{x \mid \alpha_- \leq x \leq \alpha_+\}$ if $\alpha_- \leq \alpha_+$ and $\hat{K}_n^C = \{(1/n) \sum_{i=1}^n \varepsilon^{(i)} u^{(i)}\}$ otherwise. Notice that α_- and α_+ have the same distribution, and hence \hat{K}_n^C is a closed interval with non-empty interior w.p. 1/2, and is a singleton w.p. 1/2. Thus, one can see that the linear map \hat{A}_n does not satisfy an asymptotic normality characterization. The reason for this failure is that the function $\Phi_C(\cdot, P_{K^*})$ is twice differentiable everywhere excluding the line $\{(c, c) \mid c \in \mathbb{R}\}$; in particular, it is not differentiable at the minimizer $(0, 0)$.

The above example is an instance where the function $\Phi_C(\cdot, P_{K^*})$ is not twice differentiable at \hat{A} . The manner in which an asymptotic characterization of $\{\hat{K}_n^C\}_{n=1}^\infty$ fails in instances where $M_{K^*,C}$ contains multiple orbits is also qualitatively similar. Next, we consider a series of examples in which the conditions of Theorem 3.4 hold, thus enabling an asymptotic normality characterization behavior of \hat{K}_n^C . To simplify our discussion, we describe settings in which the choices of C and A^* satisfy the conditions of Corollary 3.6, which leads to simple formulas for the second derivative Γ of the map $A \mapsto \Phi_C(A, P_{K^*})$ at A^* .

Polyhedral Examples

We present two examples in which K^* is polyhedral, and we choose $C = \Delta^q$ where q is the number of vertices of K^* . With this choice, the set $M_{K^*,C}$ comprises linear maps $A \in L(\mathbb{R}^q, \mathbb{R}^d)$ whose columns are the extreme points of K^* .

Proposition 3.7 *Let K^* be a polytope with q extreme points and let A^* be the linear map whose columns are the extreme points of K^* (in any order). The second derivative*

of $\Phi_C(\cdot, P)$ at A^* is given by $\Gamma(D) = \sum_{j=1}^q \int_{u \in H_{A^*,j}} \langle u \otimes e_j, D \rangle u \otimes e_j \, du$, where $H_{A^*,j} := \{u \in S^{d-1} \mid u^T A e_j > u^T A e_k \text{ for all } k \neq j\}$.

Proof Let $A \in L(\mathbb{R}^q, \mathbb{R}^d)$ be a linear map whose columns are pairwise distinct. Then the set $V_A := \{u \mid u^T A e_i = u^T A e_j, 1 \leq i < j \leq d\}$ is a subset of S^{d-1} with measure zero. We note that the function $h_C(\cdot)$ is differentiable at $A^T u$ whenever $u \in S^{d-1} \setminus \{V_A\}$, and hence the function $h_C(A^T u)$ is differentiable with respect to u P -a.e. By Proposition 2.9, the derivative of $\Phi_C(\cdot, P)$ at A is $2\mathbb{E}_P[(h_C(A^T u) - y)u \otimes e_C(A^T u)]$.

The subset of linear maps whose columns are pairwise distinct is open. Hence the derivative of $\Phi_C(\cdot, P)$ exists in a neighborhood of A^* . It is clear that the derivative is also continuous at A^* . Finally, we note that the collection $\{H_{A^*,j}\}_{j=1}^q$ forms a disjoint cover of S^{d-1} up to a subset of measure zero, and apply Proposition 2.10 to conclude the expression for Γ . □

We note that the operator Γ has block diagonal structure since each integral $\int_{u \in H_{A^*,j}} \langle u \otimes e_j, D \rangle u \otimes e_j \, du$ is supported on a $d \times d$ dimensional sub-block. Subsequently, by combining the conclusions of Theorem 3.1 and Proposition 3.7, we can conclude the following about \hat{K}_n^C : (i) it is a polytope with q extreme points, (ii) each vertex of \hat{K}_n^C is close to a distinct vertex of K^* , (iii) the deviations (after scaling by a factor of \sqrt{n}) between every vertex-vertex pair are asymptotically normal with inverse covariance specified by a $d \times d$ block of Γ , and further these deviations are pairwise independent.

Example Let K^* be the regular q -gon in \mathbb{R}^2 with vertices $v_k := (\cos(2k\pi/q), \sin(2k\pi/q))^T, k = 0, \dots, q - 1$. Let $\hat{v}_{n,k}$ be the vertex of \hat{K}_n^C closest to v_k . The deviation $\sqrt{n}(\hat{v}_{n,k} - v_k)$ is asymptotically normal with covariance $2\sigma^2 M_{k,k}^{-1}$, where

$$M_{k,k} = \frac{I}{q} + \frac{\sin(2\pi/q)}{2\pi} \begin{pmatrix} \cos(4k\pi/q) & \sin(4k\pi/q) \\ \sin(4k\pi/q) & -\cos(4k\pi/q) \end{pmatrix}.$$

The eigenvalues of $M_{k,k}$ are $1/q + (1/2\pi) \sin(2\pi/q)$ and $1/q - (1/2\pi) \sin(2\pi/q)$, and the corresponding eigenvectors are $(\cos(2k\pi/q), \sin(2k\pi/q))^T$ and $(\sin(2k\pi/q), -\cos(2k\pi/q))^T$ respectively. Consequently, the deviation $\hat{v}_{n,k} - v_k$ has magnitude $\approx \sigma \sqrt{q/n}$ in the direction v_k , and has magnitude $\approx \sigma \sqrt{3q^3/\pi^2 n}$ in the direction v_k^\perp . Figure 2 shows K^* as well as the confidence intervals (ellipses) of the vertices of \hat{K}_n^C for large n and $q = 5$.

Example Let K^* be the ℓ_∞ -ball in \mathbb{R}^d . For any vertex v of K^* , let $\hat{w}_{n,v}$ denote the vertex of \hat{K}_n^C closest to v . The deviation $\sqrt{n}(\hat{w}_{n,v} - v)$ is asymptotically normal with covariance $2\sigma^2 M_{v,v}^{-1}$, where

$$M_{v,v} = \frac{1}{2^d d} \left(\left(1 - \frac{2}{\pi}\right) I + \frac{2}{\pi} v v^T \right).$$

Hence the deviation $\hat{w}_{n,v} - v$ has magnitude $\approx \sigma 2^{(d+1)/2} (2/\pi + (1-2/\pi)/d)^{-1/2} n^{-1/2}$ in the span of v and magnitude $\approx \sigma 2^{(d+1)/2} (1 - 2/\pi)^{-1/2} \sqrt{d/n}$ in the orthogonal complement v^\perp .




Fig. 2 Estimating a regular pentagon as the projection of Δ^5 . In the limit (large n), the estimator \hat{K}_n^C is a pentagon. The typical deviation of the vertices of \hat{K}_n^C from that of the pentagon (scaled by a factor of \sqrt{n}) is represented by the surrounding ellipses

Non-Polyhedral Examples

Next we present two examples in which K^* is non-polyhedral. Unlike the previous polyhedral examples in which columns of the linear map \tilde{A}_n map directly to vertices, our description of \hat{K}_n^C requires a different interpretation of Corollary 3.6 that is suitable for sets with infinitely many extreme points. Specifically, we characterize the deviation $\sqrt{n}(\tilde{A}_n - A^*)$ in terms of a perturbation to the set $K^* = A^*(C)$ by considering the image of C under the map $\sqrt{n}(\tilde{A}_n - A^*)$.

Example Suppose $K^* = B_{\|\cdot\|_2}(c)$ is the unit ℓ_2 -ball in \mathbb{R}^d with center c . We consider $C := \{(1, v)^T \mid \|v\|_2 \leq 1\} \subset \mathbb{R}^{d+1}$, and A^* is any linear map of the form $[c \ Q] \in L(\mathbb{R}^{d+1}, \mathbb{R}^d)$, where $Q \in O(d, \mathbb{R})$ is any orthogonal matrix. Then the restricted Hessian $\Gamma|_{T^\perp}$ is a self-adjoint operator with rank $d + \binom{d+1}{2}$. The eigenvectors of $\Gamma|_{T^\perp}$ represent ‘modes of oscillations’, which we describe in greater detail. We begin with the case $d = 2$, $c = 0$. The set resulting from the deviation $\sqrt{n}(\tilde{A}_n - A^*)$ applied to C can be decomposed into five different modes of perturbation (these exactly correspond to the eigenvectors of the operator $\Gamma|_{T^\perp}$). Parametrizing the extreme points of $A^*(C)$ by $\{(\cos \theta \ \sin \theta)^T \mid \theta \in [0, 2\pi)\}$, the contribution of each mode at the point $(\cos \theta \ \sin \theta)^T$ is a small perturbation in the directions $(1 \ 0)^T$, $(0 \ 1)^T$, $(\cos \theta \ \sin \theta)^T$, $(\cos \theta \ -\sin \theta)^T$, and $(\sin \theta \ \cos \theta)^T$ respectively—Fig. 3 provides an illustration. The first and second modes represent oscillations of \hat{K}_n^C about c , the third mode represents dilation, and the fourth and fifth modes represent flattening. The analysis for a general d is similar, with d modes representing oscillations about c and $\binom{d+1}{2}$ modes whose contributions represent higher-dimensional analogs of flattening (of which dilation is a special case).

Example Let K^* be the spectral norm ball in \mathbb{S}^2 . The extreme points of K^* consist of three connected components: $\{I\}$, $\{-I\}$, and $\{UDU^T \mid U \in \text{SO}(2, \mathbb{R})\}$ where D is a diagonal matrix with entries $(1, -1)$. To simplify our discussion, we apply a scaled isometry to K^* so that $\{I\}$, $\{-I\}$, and $\{UDU^T \mid U \in \text{SO}(2, \mathbb{R})\}$ are mapped to the points $\{(0, 0, 1)^T\}$, $\{(0, 0, -1)^T\}$, and $\{(\cos \theta, \sin \theta, 0)^T \mid \theta \in [0, 2\pi)\}$ in \mathbb{R}^3 ,




Fig. 3 Modes of oscillations for an estimate of the ℓ_2 -ball in \mathbb{R}^2

respectively. We choose

$$\begin{aligned}
 C &:= \{X \in \mathcal{O}^4 \mid X_{12} = X_{13} = X_{14} = X_{23} = X_{24} = X_{21} \\
 &= X_{31} = X_{41} = X_{32} = X_{42} = 0\} \cong \mathcal{O}^1 \times \mathcal{O}^1 \times \mathcal{O}^2,
 \end{aligned}
 \tag{18}$$

and A^* to be the map defined by $A^*(X) = (\langle A_1, X \rangle, \langle A_2, X \rangle, \langle A_3, X \rangle)^T$ where

$$A_1 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad A_3 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

For large n , \hat{K}_n^C is a convex set with extreme points P_1 and P_2 near $(0, 0, 1)^T$ and $(0, 0, -1)^T$ respectively, and a set of extreme points specified by an ellipse P_3 near $\{(\cos \theta, \sin \theta, 0)^T \mid \theta \in [0, 2\pi)\}$. The operator $\Gamma|_{\mathcal{T}}$ is block diagonal with rank 14—it comprises two 3-dimensional blocks Γ_1 and Γ_2 associated with P_1 and P_2 , and an 8-dimensional block Γ_3 associated with P_3 . One conclusion is that the distributions of P_1 , P_2 , and P_3 are asymptotically independent. Moreover, the deviations of P_1 and P_2 about $\{(0, 0, 1)^T\}$ and $\{(0, 0, -1)^T\}$ are asymptotically normal with inverse covariance specified by Γ_1 and Γ_2 , respectively. We consider the behavior of P_3 in further detail. The operator Γ_3 is the sum of an operator $\Gamma_{3,xy}$ with rank 5 describing the variation of P_3 in the xy -plane, and another operator $\Gamma_{3,z}$ with rank 3 describing the variation of P_3 in the direction of the z -axis. The operator $\Gamma_{3,xy}$, when restricted to the appropriate subspace and suitably scaled, is equal to the operator we encountered in the previous example in the setting in which K^* is the ℓ_2 -ball in \mathbb{R}^2 . The operator $\Gamma_{3,z}$ comprises a single mode representing oscillations of P_3 in the z direction (see subfigure (b) in Fig. 4), and two modes representing “wobbling” of P_3 with respect to the xy -plane (see subfigures (c) and (d) in Fig. 4). The set C we consider in this example is the intersection of the spectraplex \mathcal{O}^4 with an appropriate subspace specified by (18). A natural question is if the same analysis holds for $C = \mathcal{O}^4$. Unfortunately, the introduction of additional dimensions introduces degeneracies (in the form of zero eigenvalues into $\Gamma|_{\mathcal{T}^\perp}$) which violates the requirements of Theorem 3.4.




Fig. 4 Estimating K^* , the spectral norm ball in \mathbb{S}^2 , as the projection of the set C (18). The above figure describes the modes of oscillation corresponding to the set of extreme points of \hat{K}_n^C given by an ellipse (see the above accompanying discussion). There are eight modes altogether—five occur in the xy -plane described in Fig. 3 and the remaining three are shown in (b), (c), and (d)

3.2.2 Specialization of Theorem 3.4 to Linear Images of Spectraplices

Proposition 3.8 *Let C be the spectraplex. Then the collection specified by (17) forms a VC class.*

Proof Define the polynomial

$$p(u, y, s, D, e_1, e_2, e_3, e_4) := \langle \hat{A}^T u - yI, e_1 e_1^T \rangle^2 - \langle (\hat{A} + D)^T u - yI, e_2 e_2^T \rangle^2 - s \langle e_3, D e_4 \rangle,$$

where $e_1, e_2, e_4 \in \mathbb{R}^q$ and $e_3 \in \mathbb{R}^d$. By [31, Thm. 1], the following collection of sets forms a VC class:

$$\left\{ \left\{ (u, y, s) \mid \sup_{\|e_1\| \leq 1} \inf_{\|e_2\|, \|e_3\|, \|e_4\| \leq 1} p(u, y, s, D, e_1, e_2, e_3, e_4) \geq 0 \right\} \mid D \in L(\mathbb{R}^q, \mathbb{R}^d) \right\}.$$

Similarly, by [31, Thm. 1], the collection $\{(u, y, s) \mid s \geq 0\} \mid D \in L(\mathbb{R}^q, \mathbb{R}^d)\}$ also forms a VC class. By [17, Lem. 9.7, p. 159], the collection of sets formed by taking intersections of members of two VC classes is a VC class. Hence it follows that the collection $\{(u, y, s) \mid (h_C((\hat{A} + D)^T u) - y)^2 - (h_C(\hat{A}^T u) - y)^2 \geq s \|D\|_F \geq 0\}$ is a VC class. Subsequently, by setting $D = A - \hat{A}$ and by noting that a sub-collection of a VC class is still a VC class, the collection $\{(u, y, s) \mid d_{C, A, \hat{A}}(u, y) \geq s \geq 0\} \mid A \in B_{\|\cdot\|_F}(\hat{A}) \setminus \{\hat{A}\}\}$ forms a VC class. A similar sequence of arguments shows that the collection $\{(u, y, s) \mid d_{C, A, \hat{A}}(u, y) \leq s \leq 0\} \mid A \in B_{\|\cdot\|_F}(\hat{A}) \setminus \{\hat{A}\}\}$ also forms a VC class. Last, by [17, Lem. 9.7, p. 159], a union of VC classes is a VC class, and so our result follows. \square

3.3 Preservation of Facial Structure

Our third result describes conditions under which the constrained estimator (2) preserves the facial structure of the underlying set K^* . We begin our discussion with some stylized numerical experiments that illustrate various aspects that inform our subsequent theoretical development. First, we consider reconstruction of the ℓ_1 -ball in \mathbb{R}^3




Fig. 5 Reconstructions of the unit ℓ_1 -ball (left) in \mathbb{R}^3 from 200 noisy support function measurements using our method with $C = \Delta^6$ (second from left), and with $C = \Delta^{12}$ (third from left). The LSE is the rightmost figure




Fig. 6 Two reconstructions of the unit ℓ_∞ -ball in \mathbb{R}^3 from 75 noisy support function measurements using our method. The choice of lifting set is $C = \Delta^8$. The ℓ_∞ -ball is the leftmost figure, and the reconstructions are the second and third figures from the left

from 200 noisy support function evaluations with the choices $C = \Delta^6$ and $C = \Delta^{12}$. Figure 5 shows these reconstructions along with the LSE. When $C = \Delta^6$, our results show a one-to-one correspondence between the faces of the reconstruction obtained using our method (second subfigure from the left) with those of the ℓ_1 -ball (leftmost subfigure); in contrast, we do not observe an analogous correspondence in the other two cases.

Second, we consider reconstruction of the ℓ_∞ -ball in \mathbb{R}^3 from 75 noisy support function measurements with $C = \Delta^8$. From Fig. 6 we see that both reconstructions (obtained from two different sets of 75 measurements) break most of the faces of the ℓ_∞ -ball. In these examples the association between the faces of the underlying set and those of the reconstruction is somewhat transparent as the sets are polyhedral. The situation becomes more delicate with non-polyhedral sets. We describe next a numerical experiment in which we estimate the *Race Track* in \mathbb{R}^2 from 200 noisy support function measurements with $C = \mathcal{O}^4$:

$$\text{Race Track} := \text{conv} \{ (x, y)^T \mid \| (x, y)^T - (-1, 0)^T \|_2 \leq 1 \text{ or } \| (x, y)^T - (1, 0)^T \|_2 \leq 1 \}.$$

From this experiment, it appears that the exposed extreme points of the *Race Track* are recovered, although the two one-dimensional edges are not recovered and seem to be distorted into curves. However, the exact correspondence between faces of the *Race Track* and those of the reconstruction seems less clear (Fig. 7).

Our first technical contribution in this subsection is a formal notion of ‘preservation of face structure’. Motivated by the preceding example, the precise manner in which we do so is via the existence of an invertible affine transformation between the faces




Fig. 7 Reconstructions of the *Race Track* from 200 noisy support function measurements using (2) with $C = \mathcal{O}^4$

of the underlying set and the faces of the reconstruction. For analytical tractability, our result focuses on exposed faces (faces that are expressible as the intersection of the underlying set with a hyperplane).

Definition Let $\{K_n\}_{n=1}^\infty \subset \mathbb{R}^d$ be a sequence of compact convex sets converging to some $K \subset \mathbb{R}^d$. Let $F \subset K$ be an exposed face. We say that F is *preserved* by the sequence $\{K_n\}_{n=1}^\infty$ if there is a sequence $\{F_n\}_{n=n_0}^\infty$, $F_n \subseteq K_n$, such that:

- (i) $F_n \rightarrow F$.
- (ii) F_n are exposed faces of K_n .
- (iii) There is an invertible affine transformation b_n such that $F = b_n F_n$ and $F_n = b_n^{-1} F$.

As our next contribution, we consider conditions under which exposed faces of K^* are preserved. To gain intuition for the types of assumptions that may be required, we review the results of the numerical experiments presented above. In the setting with K^* being the ℓ_1 -ball in \mathbb{R}^3 , all the faces are simplicial and the reconstruction with $C = \Delta^6$ preserves all the faces. In contrast, in the experiment with K^* being the ℓ_∞ -ball in \mathbb{R}^3 and $C = \Delta^8$, some of the faces of the ℓ_∞ -ball are broken into smaller simplicial faces in the reconstruction. These observations suggest that we should only expect preservation of simplicial faces, at least in the polyhedral context. However, in attempting to reconstruct the ℓ_1 -ball with $C = \Delta^{12}$, some of the simplicial faces of the ℓ_1 -ball are broken into smaller simplicial faces in the reconstruction. This is due to the overparametrization of the class of polytopes over which the regression (2) is performed (polytopes with at most twelve vertices) relative to the complexity of the ℓ_1 -ball (a polytope with six vertices). We address this point in our theorem via the single-orbit condition discussed in Sect. 2.2, which also plays a role in Theorem 3.4. Finally, in the non-polyhedral setting with K^* being the *Race Track* and $C = \mathcal{O}^4$, the two one-dimensional faces deform to curves in the reconstruction. This is due to the fact that (generic) small perturbations to the linear image of \mathcal{O}^4 that gives K^* lead to a deformation of the edges of K^* to curves. To ensure that faces remain robust to perturbations of the linear images, we require that the normal cones associated to faces that are to be preserved must be sufficiently large.

Theorem 3.9 *Suppose that $K^* \subseteq \mathbb{R}^d$ is a compact convex set with non-empty interior. Let $C \subset \mathbb{R}^q$ be a compact convex set such that $\text{Span } C \cong \mathbb{R}^q$. Suppose that there is a linear map $A^* \in L(\mathbb{R}^q, \mathbb{R}^d)$ such that $K^* = A^*(C)$, and $M_{K^*, C} = A^* \cdot \text{Aut } C$. Let*

$\{\hat{A}_n\}_{n=1}^\infty$, $\hat{A}_n \in \arg \min_A \Phi_C(A, P_{n, K^*})$, be a sequence of minimizers of the empirical loss function, and let $\{\hat{K}_n^C\}_{n=1}^\infty$, $\hat{K}_n^C = \hat{A}_n(C)$, be the corresponding sequence of estimators of K^* . Given an exposed face $F^* \subset K^*$, let $G = \{x \mid A^*x \in F^*\} \cap C$ be its pre-image, and let $N_C(G) := \{v \mid \langle v, x \rangle \geq \langle v, y \rangle \text{ for all } x \in G, y \in C\}$ be the normal cone of G w.r.t. C . If

- (i) the linear map A^* is injective when restricted to $\text{aff } G$, and
- (ii) $\dim \text{Span } N_C(G) > q - \text{rank } A^*$,

then F^* is preserved by the sequence $\{\hat{K}_n^C\}_{n=1}^\infty$.

Before giving a proof of this result, we remark next on some of the consequences.

Remark Suppose K^* is a full-dimensional polytope with q extreme points and we choose $C = \Delta^q$. It is easy to see that there is a linear map A^* such that $K^* = A^*(\Delta^q)$ and that $M_{K^*, \Delta^q} = A^* \cdot \text{Aut } \Delta^q$. Let $F^* \subseteq K^*$ be any face (note that all faces of a polytope are exposed), and let G be its pre-image in Δ^q . Note that G is an exposed face, and hence G is of the form $\{\Pi x \mid x \geq 0, \langle x, 1 \rangle = 1, x_{s+1} = \dots = x_q = 0\}$ for some $\Pi \in \text{Aut } \Delta^q$ and some $s \leq q$. The map A^* being injective on $\text{aff } G$ implies that the image of G under A^* is isomorphic to G ; i.e., F^* is simplicial. The normal cone $N_{\Delta^q}(G)$ is given by $\{\Pi z \mid z \leq 0, z_1 = \dots = z_s = 0\}$, and the requirement $\dim \text{aff } N_{\Delta^q}(G) > q - \text{rank } A^*$ holds precisely when $s < d$; i.e., the face F^* is *proper*. Thus, Theorem 3.9 implies that all *proper simplicial* faces of K^* are preserved in the reconstruction.

Remark Suppose K^* is the image under A^* of the spectraplex $C = \mathcal{O}^p$ and that $M_{K^*, C} = A^* \cdot \text{Aut } C$. Let F^* be an exposed face and let G be its pre-image in \mathcal{O}^p . Then G is a face of \mathcal{O}^p , and is of the form

$$G = \left\{ UDU^T \mid D = \begin{pmatrix} D_{11} & 0 \\ 0 & 0 \end{pmatrix}, D_{11} \in \mathcal{O}^r \right\}$$

for some $U \in O(p, \mathbb{R})$ and some $r \leq p$. Note that

$$N_{\mathcal{O}^p}(G) = \left\{ UDU^T \mid D = \begin{pmatrix} 0 & 0 \\ 0 & -D_{22} \end{pmatrix}, D_{22} \in \mathbb{S}^{p-r}, D_{22} \geq 0 \right\},$$

Thus, the requirement that $\dim \text{aff } N_{\mathcal{O}^p}(G) > \binom{p+1}{2} - \text{rank } A^*$ holds precisely when $d > pr - \binom{r-1}{2}$. We consider this result in the context of our earlier example involving the *Race Track*. Specifically, one can represent the *Race Track* as a linear image of \mathcal{O}^4 given by the following linear map:

$$A^*(X) = \begin{pmatrix} \langle A_1, X \rangle \\ \langle A_2, X \rangle \end{pmatrix}, \quad A_1 = \begin{pmatrix} -1 & 1 & & \\ 1 & -1 & & \\ & & 1 & 1 \\ & & 1 & 1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 1 & & & \\ & -1 & & \\ & & 1 & \\ & & & -1 \end{pmatrix}.$$

It is clear that $\text{rank } A^* = 2$. Let F^* be the face connecting $(-1, 0)^T$ and $(1, 0)^T$, and let $G_{\mathcal{O}^4}$ be the pre-image of F^* in \mathcal{O}^4 . One can check that

$$G_{\mathcal{O}^4} = \left\{ \begin{pmatrix} x & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & y & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \mid x, y \geq 0, x + y \leq 1 \right\},$$

$$N_{\mathcal{O}^4}(G_{\mathcal{O}^4}) = \left\{ Z = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & * & 0 & * \\ 0 & 0 & 0 & 0 \\ 0 & * & 0 & * \end{pmatrix} \mid Z \leq 0 \right\}.$$

It follows that $\dim \text{aff } N_{\mathcal{O}^4}(G_{\mathcal{O}^4}) = 3$. As the dimension of \mathcal{O}^4 is 10, our requirement on $\dim \text{aff } N_{\mathcal{O}^4}(G_{\mathcal{O}^4})$ is not satisfied.

Proof of Theorem 3.9 As we noted in the above, let $\tilde{A}_n \in \arg \min_{A \in \hat{A}_n \cdot \text{Aut } C} \|\hat{A}_n - A\|_F$, and denote $F_n = \tilde{A}_n(G)$.

[$F_n \rightarrow F$]: Since $M_{K^*, C} = A^* \cdot \text{Aut } C$, it follows from Theorem 3.1 that $\tilde{A}_n \rightarrow A^*$, from which we have $F_n \rightarrow F^*$.

[F_n are faces of K_n]: Since F^* is an exposed face of K^* , there exists $y \in \mathbb{R}^d$ and $c \in \mathbb{R}$ such that $\langle y, x \rangle = c$ for all $x \in F^*$, and $\langle y, x \rangle > c$ for all $x \in K^* \setminus F^*$. This implies that $\langle A^{*T}y, \tilde{x} \rangle = c$ for all $\tilde{x} \in G$, and $\langle A^{*T}y, \tilde{x} \rangle > c$ for all $\tilde{x} \in C \setminus G$. In particular, it implies that the row space of A^* intersects the relative interior of $N_C(G)$ in the direction A^*y .

By combining the earlier conclusion that $\tilde{A}_n \rightarrow A^*$ a.s., and that $\dim \text{Span } N_C(G) + \text{rank } A^* > q$, we conclude that the row spaces of the maps \tilde{A}_n eventually intersect the relative interior of $N_C(G)$ a.s. That is to say, there are an integer n_0 and sequences $\{y_n\}_{n=n_0}^\infty \subset \mathbb{R}^d$, $\{c_n\}_{n=n_0}^\infty \subset \mathbb{R}$ such that $\langle y_n, x \rangle = c_n$ for all $x \in F_n$, and $\langle y_n, x \rangle > c_n$ for all $x \in \hat{K}_n^C \setminus F_n$, $n \geq n_0$, a.s. In other words, the sets F_n are exposed faces of \hat{K}_n^C eventually a.s.

[One-to-one affine correspondence]: To establish a one-to-one affine correspondence between F_n and F we need to treat the case where $0 \in \text{aff } G$ and the case where $0 \notin \text{aff } G$ separately.

First suppose that $0 \in \text{aff } G$. Let $H_F = \text{aff } F$ and $H_G = \text{aff } G$. Since $0 \in H_G$, it follows that H_F and H_G are subspaces. Moreover given that A^* is injective restricted to $H_G = \text{aff } G$, it follows that H_F and H_G have equal dimensions. Hence the map t defined as the restriction of A^* onto $L(H_G, H_F)$ is square and invertible. Next let $H_{F_n} = \text{aff } F_n$, and let t_n denote the restriction of \tilde{A}_n to $L(H_G, H_{F_n})$. Given that $\tilde{A}_n \rightarrow A^*$, the maps $\{t_n\}_{n=1}^\infty$ are also square and invertible eventually a.s. It follows that one can define a linear map $b_n \in L(\mathbb{R}^d, \mathbb{R}^d)$ that coincides with $t \circ t_n^{-1}$ restricted to $L(H_{F_n}, H_F)$, is permitted to be any square invertible map on $L(H_{F_n}^\perp, H_F^\perp)$, and is zero everywhere else. Notice that b_n is invertible by construction. It straightforward to check that $F = b_n F_n$ and $F_n = b_n^{-1} F$.

Next suppose that $0 \notin \text{aff } G$. The treatment in this case is largely similar as in the previous case. Let H_F be the smallest subspace containing $\{(x, 1) \mid x \in F\} \subseteq \mathbb{R}^{d+1}$, where the set F is embedded in the first d coordinates. Let H_{F_n} be similarly defined. Let $H_G = \text{aff}(G \cup \{0\})$ —note that this defines a subspace. Since $0 \notin \text{aff } G$, there is a non-zero $z \in \mathbb{R}^q$ such that $\langle z, x \rangle = 1$ for all $x \in G$ (i.e., there exists a hyperplane containing G). Define the linear map $t \in L(H_G, H_F)$ as

$$t = \Pi_{H_F} \left[\begin{array}{c} (A^*) \\ (z^T) \end{array} \Big|_{H_G} \right],$$

where Π_{H_F} is the projection map onto the subspace H_F . Since A^* is injective on G , it follows that H_F and H_G have the same dimensions, and that t is square and invertible. One can define a square invertible map t_n analogously. The remainder of the proof proceeds in a similar fashion to the previous case, and we omit the details. Here, note that a linear invertible map operating on the lifted space \mathbb{R}^{d+1} defines an affine linear invertible map in the embedded space \mathbb{R}^d . □

4 Algorithm

We describe a procedure based on alternating updates for solving the optimization problem (2). In terms of the linear map A , the task of solving (2) can be reformulated as follows:

$$\arg \min_{A \in L(\mathbb{R}^q, \mathbb{R}^d)} \Phi(A, P_n) := \frac{1}{n} \sum_{i=1}^n (y^{(i)} - h_C(A^T u^{(i)}))^2. \tag{19}$$

As described previously, the problem (19) is non-convex as formulated; consequently, our approach is not guaranteed to return a globally optimal solution. However, we demonstrate the effectiveness of these methods with random initialization in numerical experiments in Sect. 5.

We describe our method as follows. For a fixed A , we compute $e^{(i)} \in C, i = 1, \dots, n$, so that $\langle e^{(i)}, A^T u^{(i)} \rangle = h_C(A^T u^{(i)})$, i.e., $e^{(i)} = e_C(A^T u^{(i)})$. With these $e^{(i)}$'s fixed, we update A by solving the following least squares problem:

$$\arg \min_{A \in L(\mathbb{R}^q, \mathbb{R}^d)} \frac{1}{n} \sum_{i=1}^n (y^{(i)} - \langle e^{(i)}, A^T u^{(i)} \rangle)^2. \tag{20}$$

This least squares problem can sometimes be ill-conditioned, in which case we employ Tikhonov regularization (with debiasing); see Algorithm 1.

When specialized to the choice $C = \Delta^q$, our procedure reduces to the algorithm proposed by Magnani and Boyd [20] for max-affine regression. More broadly, as noted in [20], for $C = \Delta^q$ our method is akin to Lloyd’s algorithm for K -means clustering [19]. Specifically, Lloyd’s algorithm begins with an initialization of q centers, and it alternates between (i) assigning data-points to centers based on proximity (keeping the centers fixed), and (ii) updating the location of cluster centers to minimize the squared-

loss error. In our context, suppose we express the linear map $A = [a_1 | \dots | a_q] \in \mathbb{R}^{d \times q}$ in terms of its columns. The algorithm begins with an initialization of the q columns, and it alternates between (i) assigning measurement pairs $(u^{(i)}, y^{(i)})$, $1 \leq i \leq n$, to the respective columns $\{a_j\}_{1 \leq j \leq q}$ so that the inner product $\langle u^{(i)}, a_j \rangle$ is maximized (keeping the columns fixed), and (ii) updating the columns $\{a_j\}_{1 \leq j \leq q}$ to minimize the squared-loss error.

Algorithm 1 Convex Regression via Alternating Minimization

Input: A collection $\{(u^{(i)}, y^{(i)})\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$ of support function evaluations; a compact convex set $C \subset \mathbb{R}^q$; an initialization $A \in L(\mathbb{R}^q, \mathbb{R}^d)$; a choice of regularization parameter $\gamma > 0$

Algorithm: Repeat until convergence

1. [Update optimizers of support function] $e^{(i)} \leftarrow e_C(A^T u^{(i)})$

2. [Update A by solving (20) via Tikhonov regularization (with debiasing)]

$A \leftarrow (V \otimes V + \gamma I)^{-1}(VY + \gamma A)$ where $V \leftarrow (u^{(1)} \otimes e^{(1)} | \dots | u^{(n)} \otimes e^{(n)})$, $Y \leftarrow (y^{(1)}, \dots, y^{(n)})^T$

Output: Final iterate A

5 Numerical Experiments

In this section we describe the results of numerical experiments on fitting convex sets to support function evaluations in which we contrast our framework based on solving (2) to previous methods based on solving (1). The first few experiments are on synthetically generated data, while the final experiment is on a reconstruction problem with real data obtained from the Computed Tomography (CT) scan of a human lung. For each experiment, we apply Algorithm 1 described in Sect. 4 with multiple random initializations, and we select the solution that minimizes the least squared error. The (polyhedral) LSE reconstructions in our experiments are based on the algorithm proposed in [10, Sect. 4].

5.1 Reconstructing the ℓ_1 -ball and the ℓ_2 -ball

We consider reconstructing the ℓ_1 -ball $\{g \mid \|g\|_1 \leq 1\} \subset \mathbb{R}^3$ and the ℓ_2 -ball $\{g \mid \|g\|_2 \leq 1\} \subset \mathbb{R}^3$ from noiseless and noisy support function evaluations based on the model (3). In particular, we evaluate the performance of our framework relative to the reconstructions provided by the LSE for $n = 20, 50, 200$ measurements. For both the ℓ_1 -ball and the ℓ_2 -ball in the respective noisy cases, the measurements are corrupted with additive Gaussian noise of variance $\sigma^2 = 0.1$. The reconstructions based on our framework (2) of the ℓ_1 -ball employ the choice $C = \Delta^6$, while those of the ℓ_2 -ball use $C = \mathcal{O}^3$. Figures 8 and 9 give the results corresponding to the ℓ_1 -ball and the ℓ_2 -ball, respectively.

Considering first a setting with noiseless measurements, we observe that our approach gives an exact reconstruction for both the ℓ_1 -ball and the ℓ_2 -ball. For the ℓ_1 -ball this occurs when we have $n = 200$ measurements, while the LSE provides a

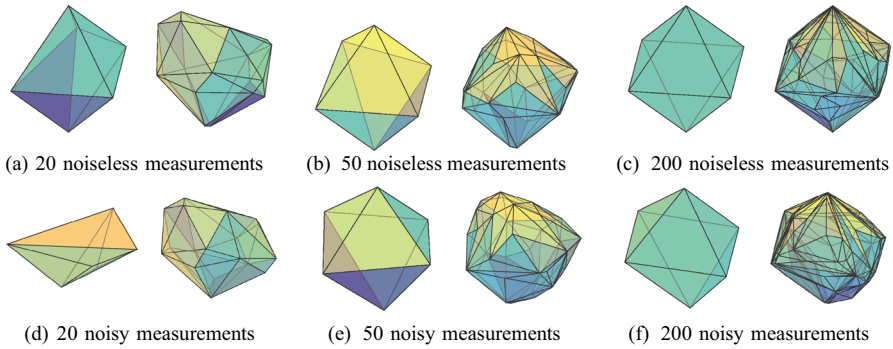


Fig. 8 Reconstruction of the unit ℓ_1 -ball in \mathbb{R}^3 from noiseless (first row) and noisy (second row) support function measurements. The reconstructions obtained using our method (with $C = \Delta^6$ in (2)) are on the left of every subfigure, while the LSE reconstructions are on the right of every subfigure

reconstruction with substantially more complicated facial structure that does not reflect that of the ℓ_1 -ball. Indeed, the LSE only approaches the ℓ_1 -ball with respect to the Hausdorff metric, but despite being the best solution in terms of minimizing the least-squares criterion, the reconstruction offered by this method provides little information about the facial geometry of the ℓ_1 -ball. Further, even with $n = 20, 50$ measurements, our reconstructions bear far closer resemblance to the ℓ_1 -ball, while the LSE in these cases looks very different from the ℓ_1 -ball. For the ℓ_2 -ball, our approach provides an exact reconstruction with just $n = 20$ measurements, while the LSE only begins to resemble the ℓ_2 -ball with $n = 200$ measurements (and even then, the reconstruction is a polyhedral approximation).

Turning our attention next to the noisy case, the contrast between the results obtained using our framework and those of the LSE approach is even more stark. For both the ℓ_1 -ball and the ℓ_2 -ball, the LSE reconstructions bear little resemblance to the underlying convex set, unlike the estimates produced using our method. Notice that the reconstructions of the ℓ_2 -ball using our algorithm are not even ellipsoidal when the number of measurements is small (e.g., when $n = 20$), as linear images of the spectraplex \mathcal{O}^3 may be non-ellipsoidal in general and need not even consist of smooth boundaries. Nonetheless, as the number of measurements available to our algorithm increases, the estimates improve in quality and offer improved reconstructions—with smooth boundaries—of the ℓ_2 -ball.

In summary, these synthetic examples demonstrate that our framework is much more effective than the LSE in terms of robustness to noise, accuracy of reconstruction given a small number of measurements, and in settings in which the underlying set is non-polyhedral.

5.2 Reconstruction via Linear Images of the Spectraplex

In the next series of synthetic experiments, we consider reconstructions of convex sets with both smooth and non-smooth features on the boundary via linear images of the spectraplex. In these illustrations, we consider sets in \mathbb{R}^2 and in \mathbb{R}^3 for which noiseless

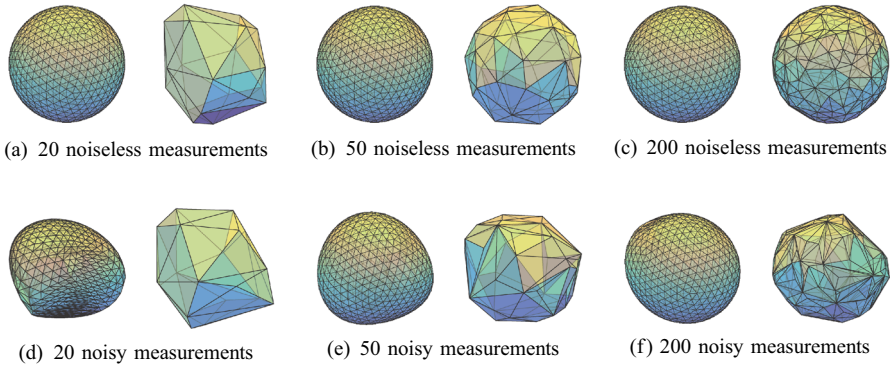


Fig. 9 Reconstruction of the unit ℓ_2 -ball in \mathbb{R}^3 from noiseless (first row) and noisy (second row) support function measurements. The reconstructions obtained using our method (with $C = \mathcal{O}^3$ in (2)) are on the left of every subfigure, while the LSE reconstructions are on the right of every subfigure




Fig. 10 Approximating the ℓ_1 -ball in \mathbb{R}^2 as a projection of the spectraplices \mathcal{O}^2 (left), \mathcal{O}^3 (center), and \mathcal{O}^4 (right)




Fig. 11 Approximating the ℓ_1 -ball in \mathbb{R}^3 as a projection of the spectraplices \mathcal{O}^3 , \mathcal{O}^4 , \mathcal{O}^5 , and \mathcal{O}^6 (from left to right)

support function evaluations are obtained and supplied as input to the problem (2), with C equal to a spectraplex \mathcal{O}^p for different choices of p . For the examples in \mathbb{R}^2 , the support function evaluations are obtained at 1000 equally spaced points on the unit circle S^1 . For the examples in \mathbb{R}^3 , the support function evaluations are obtained at 2562 regularly spaced points on the unit sphere S^2 based on an icosphere discretization.

We consider reconstruction of the ℓ_1 -ball in \mathbb{R}^2 and in \mathbb{R}^3 . Figure 10 shows the output from our algorithm when $d = 2$ for $p \in \{2, 3, 4\}$, and the reconstruction is exact for $p = 4$. Figure 11 shows the output from our algorithm when $d = 3$ for $p \in \{3, 4, 5, 6\}$. Interestingly, when $d = 3$ the computed solution for $p = 5$ does not contain any isolated extreme points (i.e., vertices) even though such features are expressible as projections of the spectraplex \mathcal{O}^5 .




Fig. 12 Reconstructions of K^* (defined in (21)) as the projection of \mathcal{O}^3 (top row) and \mathcal{O}^4 (bottom row). The figures in each row are different views of a single reconstruction, and are orientated in the $(0, 0, 1)$, $(0, 1, 0)$, $(1, 0, 1)$, and $(1, 1, 0)$ directions (from left to right) respectively

As our next illustration, we consider the following projection of \mathcal{O}^4 :

$$\text{UPillow} = \left\{ (x, y, z)^T \mid X \in \mathcal{O}^4, X_{12} = X_{21} = x, X_{23} = X_{32} = y, X_{34} = X_{43} = z \right\} \subset \mathbb{R}^3. \tag{21}$$

We term this convex set the ‘uncomfortable pillow’ and it contains both smooth and non-smooth features on its boundary. Figure 12 shows the reconstruction of UPillow as linear images of \mathcal{O}^3 and \mathcal{O}^4 computed using our algorithm. The reconstruction based on \mathcal{O}^4 is exact, while the reconstruction based on \mathcal{O}^3 smoothens out some of the ‘pointy’ features of UPillow; see for example the reconstructions based on \mathcal{O}^3 and on \mathcal{O}^4 viewed in the $(0, 1, 0)$ direction in Fig. 12).

5.3 Polyhedral Approximations of the ℓ_2 -Ball and the Tammes Problem

In the third set of synthetic experiments, we consider polyhedral approximations of the ℓ_2 -ball in \mathbb{R}^3 . This problem has been studied in many contexts under different guises. For instance, the Tammes problem seeks the optimal placement of q points on S^2 so as to maximize the minimum pairwise distance, and it is inspired by pattern formation in pollens [32].² Another body of work studies the asymptotics of polyhedral approximations of general compact convex bodies (see, for example [5]). In the optimization literature, polyhedral approximations of the second-order cone have been investigated in [4]—in particular, the approach in [4] leads to an approximation that is based on expressing the ℓ_2 -ball via a nested hierarchy of planar spherical constraints, and to subsequently approximate these constraints with regular polygons.

Our focus in the present series of experiments is to investigate polyhedral approximations of the Euclidean sphere from a *computational* perspective by employing the

² The Tammes problem is a special case of Thompson’s problem as well as Smale’s 7th problem [29].



Fig. 13 Approximating the ℓ_2 -ball in \mathbb{R}^3 as the projection of Δ^q for $q \in \{4, 5, \dots, 12\}$ (from left to right)

algorithmic tools developed in this paper. The experimental setup is similar to that of the previous subsection: we supply 2562 regularly-spaced points in S^2 (with corresponding support function values equal to one) based on an icosphere discretization as input to (2), and we select C to be the simplex Δ^q for a range of values of q . Figure 13 shows the optimal solutions computed using our method for $q \in \{4, 5, \dots, 12\}$. It turns out that the results obtained using our approach are closely related for certain values of q to optimal configurations of the Tammes problem [7,28]:

$$\arg \max_{\{a_j\}_{j=1}^q \subset S^{d-1}} \min_{1 \leq k < l \leq q} \text{dist}(a_k, a_l) = \arg \min_{\{a_j\}_{j=1}^q \subset S^{d-1}} \max_{1 \leq k < l \leq q} \langle a_k, a_l \rangle. \tag{22}$$

Specifically, the face lattice of our solutions is isomorphic to that of the Tammes problem for $q \in \{4, 5, 6, 7, 12\}$, which suggests that these configurations are stable and optimal for a broader class of objectives. We are currently not aware if the distinction between the solutions to the two sets of problems for $q \in \{8, 9, 10, 11\}$ is a result of our method recovering a locally optimal solution (in generating these results, we apply 500 initializations for each instance of q), or if it is inherently due to the different objectives that the two problems seek to optimize. For the case of $q = 8$, the difference appears to be due to the latter reason as an initialization supplied to our algorithm based on a configuration that is isomorphic to the Tammes solution led to a suboptimal local minimum.

5.4 Reconstruction of a Human Lung

In the final set of experiments we apply our algorithm to reconstruct a convex mesh of a human lung. The purpose of this experiment is to demonstrate the utility of our algorithm in a setting in which the underlying object is not convex. Indeed, in many applications in practice of reconstruction from support function evaluations, the underlying set of interest is not convex; however, due to the nature of the measurements available, one seeks a reconstruction of the convex hull of the underlying set. In the present example, the set of interest is obtained from the CT scan of the left lung of a healthy individual [9]. We note that a priori it is unclear whether the convex hull of the lung is well approximated as a linear image of either a low-dimensional simplex or a low-dimensional spectraplex.

We first obtain $n = 50$ noiseless support function evaluations of the lung (note that this object lies in \mathbb{R}^3) in directions that are generated uniformly at random over the sphere S^2 . In the top row of Fig. 14 we show the reconstructions as projections of \mathcal{O}^q for $q \in \{3, 4, 5, 6\}$, and we contrast these with the LSE. We repeat the same

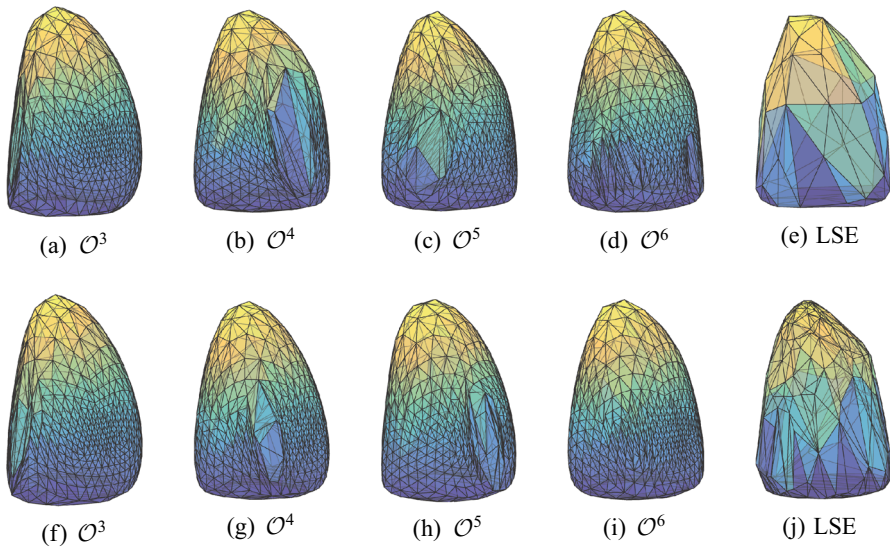


Fig. 14 Reconstructions of the left lung from 50 support function measurements (top row) and 300 support function measurements (bottom row). Subfigures (a)–(d) and (f)–(i) are projections of spectraplices with dimensions as indicated, and subfigures (e) and (j) are LSEs

experiment with $n = 300$ measurements, with the reconstructions shown in the bottom row of Fig. 14.

To concretely compare the results obtained using our framework and those based on the LSE, we contrast the description complexity—the number of parameters used to specify the reconstruction—of the estimates obtained from both frameworks. An estimator computed using our approach is specified by a projection map $A \in L(\mathbb{R}^q, \mathbb{R}^d)$, and hence it requires dq parameters, while the LSE proposed by the algorithm in [10] assigns a vertex to every measurement, and hence it requires dn parameters. The LSE using $n = 300$ measurements requires 3×300 parameters to specify whereas the estimates obtained using our framework that are specified as projections of \mathcal{O}^5 and \mathcal{O}^6 —these estimates offer comparable quality to those of the LSE—require 3×15 and 3×21 parameters, respectively. This substantial discrepancy highlights the drawback of using polyhedral sets of growing complexity to approximate non-polyhedral objects in higher dimensions.

6 Conclusions and Future Directions

In this paper we describe a framework for fitting tractable convex sets to noisy support function evaluations. Our approach provides many advantages in comparison to the previous LSE-based methods, most notably in settings in which the measurements available are noisy or small in number as well as those in which the underlying set to be reconstructed is non-polyhedral. We discuss here some potential future directions.




Fig. 15 Choosing the lifting dimension in a data-driven manner. The left sub-plot shows the cross-validation error of reconstructing the ℓ_1 -ball in \mathbb{R}^3 as the projection of Δ^q over different choices of q , and the right sub-plot shows the same quantity for $K_{S_3} \subset \mathbb{R}^3$ (see accompanying text) as the projection of \mathcal{O}^p over different choices of p

Informed selection of model complexity In practice, a suitable choice of the dimension of the simplex or the spectraplex to employ in (2) may not be available in advance. Lower-dimensional choices for C provide more concisely-described reconstructions but may not fit the data well, while higher-dimensional choices provide better fidelity to the data at the risk of overfitting. Consequently, it is of practical relevance to develop methods to select C in a data-driven manner. We describe next a stylized experiment to choose C via cross-validation.

In the first illustration, we are given 100 support function measurements of the ℓ_1 -ball in \mathbb{R}^3 corrupted by Gaussian noise with standard deviation $\sigma = 0.1$. We obtain 50 random partitions of this data into two subsets of equal size. For each partition, we solve (2) with $C = \Delta^q$ (with different choices of q) on the first subset and evaluate the mean-squared error on the second subset. The left subplot of Fig. 15 shows the average mean-squared error over the 50 partitions. We observe that initially the error decreases as q increases as a more expressive model allows us to better fit the data, and subsequently, the error plateaus out. Consequently, in this experiment an appropriate choice of C would be Δ^6 . In our second illustration, we are given 200 support function measurements corrupted by Gaussian noise with standard deviation $\sigma = 0.05$ of a set $K_{S_3} = \text{conv}(S_1 \cup S_2 \cup S_3) \subset \mathbb{R}^3$, where S_1, S_2, S_3 are defined as follows. For $j = 1, 2, 3$:

$$S_j = Q_j \left(\left\{ \begin{pmatrix} \cos \theta \\ 1 \\ \sin \theta \end{pmatrix} \mid \theta \in \mathbb{R} \right\} \right), \quad Q_j = \begin{pmatrix} \cos(2\pi j/3) & -\sin(2\pi j/3) & 0 \\ \sin(2\pi j/3) & \cos(2\pi j/3) & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

In words, the sets S_1, S_2, S_3 are disjoint planar discs. One can check that K_{S_3} is representable as a linear image of \mathcal{O}^6 . The other aspects remain the same as in the first illustration, and as we observe from the right subplot of Fig. 15, an appropriate choice for C would be \mathcal{O}^6 .

Approximation power of semidefinite descriptions Theorem 3.1 demonstrates that our estimator converges almost surely to the linear image of C that best approximates the underlying set K^* . Consequently, a natural question is to understand the quality of the approximation of general convex bodies provided by linear images of the simplex (i.e., polytopes) or the spectraplex. There is a substantial body of prior work that studies approximations of convex bodies via polytopes (see, for instance [5]), but an analogous theory for approximations that are specified as linear images of the spectraplex as well as the more general collection of feasible regions of semidefinite programs is far more limited—the closest piece of work of which we are aware is a result by Barvinok [3]. Progress on these fronts would be useful for understanding the full expressive power of our framework. More generally, and as noted in [3], obtaining even a basic understanding about the approximation power of such descriptions has broad algorithmic implications concerning the use of semidefinite programs to approximate general convex programs.

Richer families of tractable convex sets A restriction in the development in this paper is that we only consider reconstructions specified as linear images of a *fixed* convex set C ; we typically choose C to be a simplex or a spectraplex, which are given by particular slices of the non-negative orthant or the cone of positive semidefinite matrices. As described in the introduction, optimizing over more general affine sections of these cones is likely to be intractable due to the lack of a compact description of the sensitivity of the optimal value of conic optimization problems with respect to perturbations of the affine section. Consequently, it would be useful to identify broader yet structured families of sets than the ones we have considered in this paper for which such a sensitivity analysis is efficiently characterized.

References

1. Averkov, G.: Optimal size of linear matrix inequalities in semidefinite approaches to polynomial optimization. *SIAM J. Appl. Algebra Geom.* **3**(1), 128–151 (2019)
2. Balázs, G.: Convex Regression: Theory, Practice, and Applications. PhD thesis, University of Alberta (2016)
3. Barvinok, A.: Approximations of convex bodies by polytopes and by projections of spectrahedra (2012). [arXiv:1204.0471](https://arxiv.org/abs/1204.0471)
4. Ben-Tal, A., Nemirovski, A.: On polyhedral approximations of the second-order cone. *Math. Oper. Res.* **26**(2), 193–205 (2001)
5. Bronstein, E.M.: Approximation of convex sets by polytopes. *J. Math. Sci. (N.Y.)* **153**(6), 727–762 (2008)
6. Cai, T.T., Guntuboyina, A., Wei, Y.: Adaptive estimation of planar convex sets. *Ann. Stat.* **46**(3), 1018–1049 (2018)
7. Danzer, L.: Finite point-sets on S^2 with minimum distance as large as possible. *Discrete Math.* **60**, 3–66 (1986)
8. Fisher, N.I., Hall, P., Turlach, B.A., Watson, G.S.: On the estimation of a convex set from noisy data on its support function. *J. Am. Stat. Assoc.* **92**(437), 84–91 (1997)
9. Gaillard, F.: Normal chest CT-lung window. *Radiopaedia*. <https://radiopaedia.org/cases/normal-chest-ct-lung-window>
10. Gardner, R.J., Kiderlen, M.: A new algorithm for 3D reconstruction from support functions. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(3), 556–562 (2009)
11. Gardner, R.J., Kiderlen, M., Milanfar, P.: Convergence of algorithms for reconstructing convex bodies and directional measures. *Ann. Stat.* **34**(3), 1331–1374 (2006)

12. Gouveia, J., Parrilo, P.A., Thomas, R.R.: Lifts of convex sets and cone factorizations. *Math. Oper. Res.* **38**(2), 248–264 (2013)
13. Gregor, J., Rannou, F.R.: Three-dimensional support function estimation and application for projection magnetic resonance imaging. *Int. J. Imaging Syst. Technol.* **12**(1), 43–50 (2002)
14. Guntuboyina, A.: Optimal rates of convergence for convex set estimation from support functions. *Ann. Stat.* **40**(1), 385–411 (2012)
15. Hall, P., Turlach, B.A.: On the estimation of a convex set with corners. *IEEE Trans. Pattern Anal. Mach. Intell.* **21**(3), 225–234 (1999)
16. Hannah, L.A., Dunson, D.B.: Multivariate convex regression with adaptive partitioning. *J. Mach. Learn. Res.* **14**, 3261–3294 (2013)
17. Kosorok, M.R.: *Introduction to Empirical Processes and Semiparametric Inference*. Springer Series in Statistics. Springer, New York (2008)
18. Lele, A.S., Kulkarni, S.R., Willsky, A.S.: Convex-polygon estimation from support-line measurements and applications to target reconstruction from laser-radar data. *J. Optical Soc. Amer. A* **9**(10), 1693–1714 (1992)
19. Lloyd, S.P.: Least squares quantization in PCM. *IEEE Trans. Inf. Theory* **28**(2), 129–137 (1982)
20. Magnani, A., Boyd, S.P.: Convex piecewise-linear fitting. *Optim. Eng.* **10**(1), 1–17 (2009)
21. Nesterov, Yu., Nemirovskii, A.: *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM Studies in Applied Mathematics, vol. 13. Society for Industrial and Applied Mathematics, Philadelphia (1994)
22. Pollard, D.: Strong consistency of k -means clustering. *Ann. Stat.* **9**(1), 135–140 (1981)
23. Pollard, D.: *Convergence of Stochastic Processes*. Springer Series in Statistics. Springer, New York (1984)
24. Prince, J.L., Willsky, A.S.: Reconstructing convex sets from support line measurements. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**(4), 377–389 (1990)
25. Prince, J.L., Willsky, A.S.: Convex set reconstruction using prior shape information. *CVGIP: Graph. Models Image Process.* **53**(5), 413–427 (1991)
26. Saunderson, J.: Limitations on the expressive power of convex cones without long chains of faces. *SIAM J. Optim.* **30**(1), 1033–1047 (2020)
27. Schneider, R.: Convex Bodies: The Brunn–Minkowski Theory. *Encyclopedia of Mathematics and Its Applications*, vol. 151. Cambridge University Press, Cambridge (2014)
28. Schütte, K., van der Waerden, B.L.: Auf welcher Kugel haben 5, 6, 7, 8 oder 9 Punkte mit Mindestabstand Eins Platz? *Math. Ann.* **123**, 96–124 (1951)
29. Smale, S.: Mathematical problems for the next century. *Math. Intell.* **20**(2), 7–15 (1998)
30. Stark, H., Peng, H.: Shape estimation in computer tomography from minimal data. *J. Opt. Soc. Am. A* **5**(3), 331–343 (1988)
31. Stengle, G., Yükich, J.E.: Some new Vapnik–Chervonenkis classes. *Ann. Stat.* **17**(4), 1441–1446 (1989)
32. Tammes, P.M.L.: On the Origin of Number and Arrangement of the Places of Exit on the Surface of Pollen-Grains. *Recueil des Travaux Botaniques Néerlandais*, vol. 27. Koninklijke Nederlandse Botanische Vereniging (1930)
33. Vapnik, V.N., Chervonenkis, A.Ya.: On the uniform convergence of relative frequencies of events to their probabilities. *Proc. USSR Acad. Sci.* **181**(4), 781–783 (1968). (in Russian)
34. Yannakakis, M.: Expressing combinatorial optimization problems by linear programs. *J. Comput. Syst. Sci.* **43**(3), 441–466 (1991)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Yong Sheng Soh¹  · Venkat Chandrasekaran²

Yong Sheng Soh
soh_yong_sheng@ihpc.a-star.edu.sg

Venkat Chandrasekaran
venkatc@caltech.edu

- ¹ Institute of High Performance Computing, 1 Fusionopolis Way, # 16-16 Connexis, Singapore 138632, Singapore
- ² Department of Computing and Mathematical Sciences and Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125, USA