

Optimal Regularization for a Data Source

Oscar Leong^{*1}, Eliza O'Reilly², Yong Sheng Soh³, and Venkat Chandrasekaran⁴

¹Department of Computing and Mathematical Sciences, Caltech

²Department of Applied Mathematics and Statistics, Johns Hopkins University

³Department of Mathematics, National University of Singapore

⁴Department of Computing and Mathematical Sciences and Electrical Engineering, Caltech

December 29, 2022, Revised: January 30, 2024

Abstract

In optimization-based approaches to inverse problems and to statistical estimation, it is common to augment criteria that enforce data fidelity with a regularizer that promotes desired structural properties in the solution. The choice of a suitable regularizer is typically driven by a combination of prior domain information and computational considerations. Convex regularizers are attractive computationally but they are limited in the types of structure they can promote. On the other hand, nonconvex regularizers are more flexible in the forms of structure they can promote and they have showcased strong empirical performance in some applications, but they come with the computational challenge of solving the associated optimization problems. In this paper, we seek a systematic understanding of the power and the limitations of convex regularization by investigating the following questions: Given a distribution, what is the optimal regularizer for data drawn from the distribution? What properties of a data source govern whether the optimal regularizer is convex? We address these questions for the class of regularizers specified by functionals that are continuous, positively homogeneous, and positive away from the origin. We say that a regularizer is optimal for a data distribution if the Gibbs density with energy given by the regularizer maximizes the population likelihood (or equivalently, minimizes cross-entropy loss) over all regularizer-induced Gibbs densities. As the regularizers we consider are in one-to-one correspondence with star bodies, we leverage dual Brunn-Minkowski theory to show that a radial function derived from a data distribution is akin to a “computational sufficient statistic” as it is the key quantity for identifying optimal regularizers and for assessing the amenability of a data source to convex regularization. Using tools such as Γ -convergence from variational analysis, we show that our results are robust in the sense that the optimal regularizers for a sample drawn from a distribution converge to their population counterparts as the sample size grows large. Finally, we give generalization guarantees for various families of star bodies that recover previous results for polyhedral regularizers (i.e., dictionary learning) and lead to new ones for a variety of classes of star bodies.

Keywords: Brunn-Minkowski theory, convex bodies, Γ -convergence, star bodies, statistical learning, variational analysis

2020 AMS Mathematics Subject Classification: 49, 52, 90

*Correspondence email: oleong@caltech.edu

1 Introduction

Balancing computational considerations and accurate modeling of various phenomena of interest is a fundamental and long-standing challenge in data science. In the fields of statistical estimation and inverse problems, this challenge manifests itself most prominently in the choice of an appropriate regularization functional for a data source. Regularizers are commonly employed to augment data fidelity criteria in optimization-based methods for inverse problems, and they are used to promote desired structural properties in the solution. Popular forms of structure for which regularization is routinely employed include smoothness, sparsity, matrices or tensors with small rank, and many others. As inverse problems are often ill-posed due to the available data being noisy or incomplete, the selection of a suitable regularizer is crucial to obtaining high-quality solutions. In deriving a regularizer that promotes a desired structure, a central goal is that of obtaining a computationally tractable optimization formulation.

The core property of an optimization problem that governs its computational tractability is whether the problem is convex or not. Indeed, as articulated by Rockafellar [66] some decades ago “... the great watershed in optimization isn’t between linearity and nonlinearity, but convexity and nonconvexity.” Convex optimization problems can be solved reliably to global optimality, and they come with the rich toolkit of convex analysis that provide convergence guarantees and bounds on sensitivity to perturbations of the underlying problem data. These attributes have significant consequences in the context of inverse problems; for example, in medical imaging applications solution reliability and robustness are of critical importance. As such, the design of convex regularizers¹ that promote a desired structure is a question that has been studied for many decades [40]. However, convex regularizers are limited in the types of structure they can promote. As a result, there has been a lot of interest in recent years in the incorporation of nonconvex regularizers in solution methods for inverse problems, and the resulting nonconvex optimization formulations have showcased strong empirical performance in some applications [7]. Nonetheless, these approaches come with few guarantees about solution quality, robustness to data perturbations, and convergence of numerical solution schemes.

Despite substantial literature devoted to the topic of regularizer selection, we still do not have a systematic understanding of the power and the limitations of convex regularization for general data sources. For example, it is unclear whether there exist data sources for which a convex regularizer is optimal and one need not sacrifice computational efficiency. In this paper, we take first steps towards addressing these larger objectives by formulating and answering the following specific questions: Given a data source specified as a probability distribution, what is the optimal regularizer that promotes the structure contained in that data? Are there data sources for which this optimal regularizer is convex? What are the “sufficient statistics” associated with a data source that govern the choice of the optimal regularizer?

1.1 A Maximum-Likelihood Criterion for Optimal Regularizer Selection

To make the preceding conceptual questions precise, we begin by fixing a family of functionals over which we identify optimal regularizers. Whether convex or not, a regularizer ought to satisfy basic regularity properties that lead to well-behaved optimization formulations for inverse problems. To this end, we consider the collection of functionals on \mathbb{R}^d that are (i) positive except at the origin; (ii) positively homogeneous; and (iii) continuous. These assumptions are commonplace in the context of inverse problems; in particular, positivity is natural as a regularizer serves as a penalty function to promote a desired structure, positive homogeneity aligns with the idea that we often wish to solve inverse problems at varying levels of signal-to-noise ratio using the same regularizer, and finally, continuity is appealing computationally as it yields well-behaved optimization formulations. Most regularizers employed in the literature in inverse problems satisfy these three conditions.

A simple but important consequence of these three assumptions is that the collection of functionals satisfying these conditions is in one-to-one correspondence with the collection of star bodies in \mathbb{R}^d ; we call a set $K \subseteq \mathbb{R}^d$ a *star body* if it is compact with non-empty interior and for each $x \in \mathbb{R}^d \setminus \{0\}$ the ray $\{\lambda x : \lambda > 0\}$ intersects the boundary of K exactly once. Given a star body $K \subseteq \mathbb{R}^d$, the regularizer associated with this

¹Nonlinear inverse problems can potentially give rise to nonconvex data fidelity terms. We assume for simplicity that such data fidelity terms are convex, and that they yield convex optimization problems when augmented with convex regularizers.

star body is given by its gauge function (also called the Minkowski functional of K):

$$\|x\|_K = \inf\{t > 0 : x \in t \cdot K\}. \quad (1)$$

This regularizer is nonconvex for general star bodies K and is convex if and only if K is convex. The 1-sublevel set of the regularizer is K . See Section 2 for background on star bodies.

With this formalization, we describe next what we mean by a regularizer being optimal for a given data source. As data sources in this paper are characterized by probability distributions, a natural approach to obtaining an optimal regularizer is to map our collection of regularizers to a family of densities and to identify the regularizer for which the associated density maximizes the population likelihood. Specifically, we associate a regularizer $\|\cdot\|_K$ given by the gauge of a star body $K \subseteq \mathbb{R}^d$ to the Gibbs density $p_K(x) \propto \exp(-\|x\|_K)$, which yields the following family of densities:

$$\mathcal{D} := \left\{ p_K : \mathbb{R}^d \rightarrow \mathbb{R} \mid p_K(x) := e^{-\|x\|_K} / Z_K, Z_K := \int_{\mathbb{R}^d} e^{-\|x\|_K} dx < \infty, K \text{ is a star body} \right\}.$$

Note that this class is quite expressive, as it includes log-concave densities with convex gauges as energy functions as well as a wide class of nonconvex energies. For a data distribution P on \mathbb{R}^d , the negative population log-likelihood for any $p_K \in \mathcal{D}$ is given by

$$\mathbb{E}_P[-\log p_K(x)] = \mathbb{E}_P[\|x\|_K] + \log Z_K.$$

The first term is the average value that the star body gauge assigns to points drawn according to the distribution P , and the second term corresponds to the normalizing constant. This normalization is typically intractable to evaluate for general Gibbs densities [34], but in our case it has an appealing geometric characterization in terms of the volume of the star body K as described in the following result.

Proposition 1. *For any star body $K \subset \mathbb{R}^d$, we have*

$$\int_{\mathbb{R}^d} e^{-\|x\|_K} dx = \text{vol}_d(K) \cdot \Gamma(d+1)$$

where $\Gamma(z) := \int_0^\infty t^{z-1} e^{-t} dt$ is the Gamma function.

Proof. The result follows from elementary integration. Let $\mathbf{1}_E$ denote the indicator function on the event E . Recall that for any star body $K \subseteq \mathbb{R}^d$, we have $\int_{\mathbb{R}^d} \mathbf{1}_{\{x \in K\}} dx = \text{vol}_d(K)$ and $\text{vol}_d(tK) = t^d \text{vol}_d(K)$ for $t \geq 0$. Then observe that

$$\begin{aligned} \int_{\mathbb{R}^d} e^{-\|x\|_K} dx &= \int_{\mathbb{R}^d} \left(\int_{\|x\|_K}^\infty e^{-t} dt \right) dx = \int_{\mathbb{R}^d} \left(\int_0^\infty \mathbf{1}_{\{\|x\|_K \leq t\}} e^{-t} dt \right) dx \\ &= \int_0^\infty e^{-t} \left(\int_{\mathbb{R}^d} \mathbf{1}_{\{\|x\|_K \leq t\}} dx \right) dt = \int_0^\infty e^{-t} \text{vol}_d(tK) dt = \text{vol}_d(K) \int_0^\infty t^d e^{-t} dt. \end{aligned}$$

□

This result shows that minimizing the negative population log-likelihood objective $\mathbb{E}_P[-\log p_K(x)]$ over $p_K \in \mathcal{D}$ is equivalent to minimizing the following functional over all star bodies $K \subseteq \mathbb{R}^d$

$$K \mapsto \mathbb{E}_P[\|x\|_K] + \log \text{vol}_d(K). \quad (2)$$

There are additional observations one can make regarding this functional to arrive at a simpler, yet equivalent, formulation to maximum likelihood. For any fixed star body K , we consider the optimal positive dilate λ that minimizes $\mathbb{E}_P[\|x\|_K] + \log \text{vol}_d(K)$, i.e., we wish to solve $\min_{\lambda > 0} \mathbb{E}_P[\|x\|_{\lambda K}] + \log \text{vol}_d(\lambda K)$. By noting that $\mathbb{E}_P[\|x\|_{\lambda K}] = \lambda^{-1} \mathbb{E}_P[\|x\|_K]$ and that $\text{vol}_d(\lambda K) = \lambda^d \text{vol}_d(K)$, we conclude from single-variable calculus that the expression $\lambda^{-1} \mathbb{E}_P[\|x\|_K] + \log(\lambda^d \text{vol}_d(K))$ is minimized when $\lambda_K := \mathbb{E}_P[\|x\|_K] / d$. One can check that $\mathbb{E}_P[\|x\|_{\lambda_K K}] + \log \text{vol}_d(\lambda_K K)$ is equal to $\log(\mathbb{E}_P[\|x\|_{K/\text{vol}_d(K)^{1/d}}])$ up to additive constants and multiplicative factors that do not depend on K . As $K/\text{vol}_d(K)^{1/d}$ has unit volume, we conclude that

in the minimization of the functional (2) over all star bodies K , it suffices to only consider unit-volume star bodies. Consequently, we arrive at the following natural optimization problem:

$$\arg \min_{\substack{K \subseteq \mathbb{R}^d, \text{vol}_d(K)=1, \\ K \text{ is a star body}}} \mathbb{E}_P[\|x\|_K]. \quad (3)$$

which is the main problem of interest in the paper.

Solving the optimization problem (3) entails many challenges as the space of star bodies is infinite-dimensional with its own intricate topology. In particular, characterizing basic properties about the problem, such as existence and uniqueness of minimizers, requires an analysis of the map $K \mapsto \mathbb{E}_P[\|x\|_K]$ on the space of star bodies. To make progress on these challenges, we leverage concepts and ideas from the field of star geometry, with dual Brunn-Minkowski theory playing a prominent role [47]. In particular, as we show in Section 3.2, the expected gauge $\mathbb{E}_P[\|x\|_K]$ can be written as a *dual mixed volume* between the star body K and a special star body induced by the data distribution P , which we call L_P . The star body L_P serves as a “computational sufficient statistic” that summarizes all the relevant information about P for the selection of a regularizer. A virtue of this geometric characterization of $\mathbb{E}_P[\|x\|_K]$ is that the dual mixed volume between K and L_P can be bounded below by the product of the volumes of K and L_P , with equality holding if and only if K is a (positive) dilate of L_P , i.e., $K = \lambda L_P$ for some $\lambda > 0$ [47].

This problem (3) constitutes the central mathematical object of interest in our paper. We show in Section 3.2 that under appropriate conditions the volume-normalized L_P is the unique solution to the above optimization problem. In Sections 4–7, we further analyze variational, statistical, and learning-theoretic aspects of this problem; see Section 1.2 for a summary of our contributions.

From an information-geometric perspective, our maximum likelihood criterion for regularizer selection is equivalent to computing a *moment projection* (or *m-projection*) [4] of the data distribution P onto the family of distributions $\overline{\mathcal{D}} := \{p_K \in \mathcal{D} : Z_K = \Gamma(d+1)\}$. Assuming P is absolutely continuous with respect to any $p_K \in \overline{\mathcal{D}}$ (i.e., $P \ll p_K$ for any $p_K \in \overline{\mathcal{D}}$), we have the following interpretation of the problem (3):

$$\min_{p_K \in \overline{\mathcal{D}}} D_{\text{KL}}(P||p_K) \Leftrightarrow \min_{p_K \in \overline{\mathcal{D}}} \int_{\mathbb{R}^d} \log \left(\frac{dP}{dP_K}(x) \right) dP(x) \Leftrightarrow \max_{p_K \in \overline{\mathcal{D}}} \int_{\mathbb{R}^d} \log p_K(x) dP(x) \Leftrightarrow \min_{p_K \in \overline{\mathcal{D}}} \mathbb{E}_P[\|x\|_K]$$

where P_K is the measure with density p_K . The minimizer $p_K^* := \arg \min_{p_K \in \overline{\mathcal{D}}} D_{\text{KL}}(P||p_K)$ is known as the *m-projection* of P onto the family $\overline{\mathcal{D}}$, and it is given by the Gibbs density induced by the star body that solves the optimization problem (3).

1.2 Our Contributions

In Section 3, we identify the solution of (3) using techniques from dual Brunn-Minkowski theory [47]. This branch of geometric functional analysis extends classical results in Brunn-Minkowski theory concerning mixed volumes of convex bodies [72] to the class of star bodies; we show that the objective in (3) may be viewed as a dual intrinsic volume, and this interpretation yields a precise characterization of the optimal regularizer. Given a data distribution that has a density with respect to the Lebesgue measure, we show that a radial function induced by the distribution’s density provides a kind of “summary statistic” for the data that provides all the information required to obtain the optimal regularizer for that data; in particular, this radial function allows us to characterize when the optimal regularizer is convex. Next, we establish the existence of minimizers of (3) in Section 4 for a broader class of distributions than those considered in Section 3. Along the way, we describe a compactness result akin to Blaschke’s Selection Theorem for “well-conditioned” star bodies.

In Section 5, we show that our results are robust in the sense that the optimal regularizers for a sample drawn from a distribution converge to their population counterparts as the sample size grows large. Our approach to establishing this result is based on tools from variational analysis, such as Γ -convergence [13], which we use to establish uniform convergence of the objective in (3) as the sample size grows large. These tools also allow us to show that star bodies learned on noisy data converge to star bodies learned on uncorrupted data as the noise level goes to zero (see Section 5.4).

Finally, in Sections 6 and 7, we describe additional results on the statistical learning of star body regularizers and how our optimal regularizer characterization has consequences for inverse problems. Section

6 builds upon results in Sections 3 – 5 for learning regularizers from certain parametrized families of star bodies. We prove a general uniform convergence result that gives a bound on generalization error as a function of the metric entropy of the family of bodies being considered. We then apply this result to several classes of interest, including star bodies given as unions of convex bodies, ellipsoids, polytopes, and linear images of Schatten- p norm balls. Several of these generalization bounds are novel, while others recover known results in the literature for problems such as dictionary learning. Section 7 discusses how our results have consequences for posterior sampling, a crucial ingredient for Bayesian inference in inverse problems.

1.3 Related Work

1.3.1 Regularizers Derived from Domain Expertise

The traditional paradigm in solving inverse problems incorporates the use of hand-crafted regularization functionals to promote structure, a practice dating as far back as Tikhonov regularization [86] and the Nyquist sampling theorem [52]. For instance, natural images have been observed to be compressible, or approximately sparse, under the wavelet or the discrete cosine transform. Subsequently, regularization functionals such as the ℓ_1 -norm in the transformed basis, have shown to be effective at inducing sparse structure and is computationally tractable to optimize over owing to its convexity [23, 17, 24, 85]. A prominent convex regularizer promoting smoothness is total variation regularization [69], which also includes generalizations incorporating higher-order derivatives [18, 14]. Other examples of hand-crafted convex regularization techniques include the nuclear norm for low-rankness [30, 16, 62] as well as entropic regularization [27]. Several of these regularizers also fall under the wider class of atomic norm regularizers [20, 11, 84, 74, 55]. We also highlight the rich use of Gaussian priors [79, 42, 22, 35, 51], which have been used extensively in infinite-dimensional inverse problems for regularization and uncertainty quantification.

While utilizing convex regularizers has computational benefits for optimization, there has been a recent surge of interest in the incorporation of nonconvex regularization functionals. For example, returning to sparsity as a model for structure, several approaches have proposed the use of ℓ_p regularizers for $p \in [0, 1)$ [31, 82, 50] or other nonconvex approximations to the ℓ_1 -norm [95, 73, 29, 93, 96, 58]. Please see [10] for a comprehensive overview of hand-crafted regularizers for inverse problems. A significant challenge in this paradigm, however, is that choosing an appropriate regularizer for an inverse problem requires precise domain knowledge or expertise. Moreover, for general data sources, it is difficult to make claims regarding the optimality of such regularizers.

1.3.2 Regularizers Learned from Data

A dominant paradigm in recent years in designing regularizers has been to take a data-driven approach by directly *learning* a regularization functional from example data. This framework has the advantage that the learned regularizer is particular to the data distribution of interest, and can help ameliorate a lack of domain knowledge or expertise. One prominent instantiation of this is given by learned convex regularizers from dictionary learning or sparse coding (see [53, 54, 77, 3, 6, 1, 2, 70, 71, 8] and the surveys [49, 28] for more). Given a training set of signals of interest, this class of methods computes a basis such that each example can be sparsely represented in this new basis. Once learned, such a basis can be incorporated in a regularization functional for subsequent tasks by enforcing that reconstructed data should also be sparsely represented in this basis. Geometrically, this regularizer constructs a polytope (whose extreme points are given by the basis elements) such that datapoints of interest should lie on its low-dimensional faces. Extensions of such an approach have also been investigated for particular types of bases, such as convolutional dictionaries [57, 94, 32], and those representable by semidefinite programming using an infinite collection of basis elements [76]. Such regularizers are convex, but there have also been many approaches incorporating more complicated nonconvex structures for data, such as learned regularizers parametrized by deep neural networks. Such approaches have showcased strong empirical performance and have been empirically observed to out-perform handcrafted regularizers in a range of tasks [46, 90, 67, 12, 43, 38, 63]. For a comprehensive survey on data-driven regularizers for inverse problems, please see the survey [7].

Despite this empirical success of data-driven regularizers, there is a lack of an overarching understanding as to why such approaches better capture underlying data geometry than handcrafted approaches. For example, when and why is it the case that certain data distributions can be captured or modeled by sparsity in

a learned basis? On the other hand, which data sources should instead be modeled by nonconvex regularizers and how much information is lost by restricting oneself to a family of convex regularizers? These are the types of questions that we investigate in the present paper.

1.3.3 Characterizing Optimal Regularizers

To our knowledge, there are few works that consider the optimality of a regularizer for a given dataset. For example, in the area of dictionary learning, most prior work and performance guarantees analyze convergence guarantees of learning a synthetically generated, planted basis [1, 80, 81, 21], but these works do not consider the question of studying why a polyhedral regularizer is the “best choice” as a regularizer for a given dataset. A recent work that has tackled the question of identifying an optimal regularizer is [87]. Here, the authors consider signal recovery problems where the underlying signal has intrinsic low-dimensional structure. Given linear measurements of the signal, the goal is to characterize the optimal convex regularizer for the inverse problem. The authors introduce several notions of optimality for convex regularizers (dubbed compliance measures) and establish that canonical low-dimensional models, such as the ℓ_1 -norm for sparsity, are optimal for sparse recovery under such compliance measures. The results in this work, however, focus exclusively on optimality of convex regularizers for a given low-dimensional model of the underlying data source. In contrast, our results can prove optimality of both convex and nonconvex regularizers for general data sources modeled as distributions.

2 Preliminaries

In this section, we provide relevant background concerning convex bodies and star bodies. For a more detailed treatment of each subject, we recommend [72] for convex geometry and [33, 39] for star geometry.

First, we review some basic functions. Let K be a closed subset of \mathbb{R}^d . The *radial function* of K is defined by

$$\rho_K(x) := \sup\{t > 0 : t \cdot x \in K\}. \tag{4}$$

It follows that the gauge function of K satisfies $\|x\|_K = 1/\rho_K(x)$ for all $x \in \mathbb{R}^d$ such that $x \neq 0$. The *support function* of K at $x \in \mathbb{R}^d$ is given by

$$h_K(x) := \sup\{\langle x, y \rangle : y \in K\}.$$

Finally, we note that the gauge function and the support function are homogeneous with degree 1, while the radial function is homogeneous with degree -1 . As such, in what follows, we generally speak of these functions restricted to the unit sphere $\mathbb{S}^{d-1} := \{u \in \mathbb{R}^d : \|u\|_{\ell_2} = 1\}$. We also use $\text{vol}_d(K)$ to denote the usual d -dimensional volume of $K \subseteq \mathbb{R}^d$.

Recall that a closed set $K \subseteq \mathbb{R}^d$ is *convex* if $x, y \in K$ implies that $\theta x + (1 - \theta)y \in K$ for all $0 \leq \theta \leq 1$. We call a convex set K a *convex body* if it is also compact and contains the origin in its interior. We say that a closed set $K \subseteq \mathbb{R}^d$ is *star-shaped* (with respect to the origin) if for all $x \in K$, the line segment $[0, x] := \{tx : t \in [0, 1]\}$ is also contained in K . An equivalent definition of a star body as described in the introduction is a compact set $K \subseteq \mathbb{R}^d$ that is star-shaped and has a positive and continuous radial function ρ_K . We will denote the class of convex bodies and star bodies in \mathbb{R}^d by \mathcal{C}^d and \mathcal{S}^d , respectively. Note that \mathcal{C}^d is a strict subclass of \mathcal{S}^d for $d > 1$. Convex bodies are determined uniquely by their support function, and star bodies are determined uniquely by their radial function.

For $K, L \in \mathcal{S}^d$, it is easy to see that $K \subseteq L$ if and only if $\rho_K \leq \rho_L$. We say that L is a *dilate* of K if there exists $\lambda > 0$ such that $L = \lambda K$; that is, $\rho_L = \lambda \rho_K$. In addition, given a linear transformation $\phi \in GL(\mathbb{R}^d)$, one has $\rho_{\phi K}(x) = \rho_K(\phi^{-1}x)$ for all $x \in \mathbb{R}^d \setminus \{0\}$.

An additional important aspect of star bodies are their kernels. Specifically, we define the *kernel* of a star body K as the set of all points for which K is star-shaped with respect to, i.e., $\ker(K) := \{x \in K : [x, y] \subseteq K, \forall y \in K\}$. Note that $\ker(K)$ is a convex subset of K and $\ker(K) = K$ if and only if K is convex. For a parameter $r > 0$, we define the following subset of \mathcal{S}^d consisting of *well-conditioned* star bodies with nondegenerate kernels:

$$\mathcal{S}^d(r) := \{K \in \mathcal{S}^d : rB^d \subseteq \ker(K)\}.$$

Here B^d is the unit (Euclidean) ball in \mathbb{R}^d .

2.1 Metric Spaces

Let $K, L \in \mathcal{S}^d$ be star bodies. The *Minkowski sum* between K, L is given by $K + L := \{x + y : x \in K, y \in L\}$. Following the definition of the support function, we have that $h_{K+L} = h_K + h_L$. The *Hausdorff distance* is defined by

$$d_H(K, L) := \inf\{\varepsilon \geq 0 : K \subseteq L + \varepsilon B^d, L \subseteq K + \varepsilon B^d\}.$$

If K and L are also convex, then the Hausdorff distance can be defined in terms of the infinity norm of their support functions $d_H(K, L) = \|h_K - h_L\|_\infty$ – here, both support functions are defined over the unit-sphere \mathbb{S}^{d-1} . The Hausdorff distance also defines a metric over the space of nonempty compact sets. The metric space (\mathcal{C}^d, d_H) enjoys favorable compactness properties that will aid in establishing several results in our analysis. In particular, it is known that the metric space (\mathcal{C}^d, d_H) is *locally compact*, a seminal result known as Blaschke's Selection Theorem [72]. Prior to stating this result, we recall that a collection \mathfrak{C} of sets is bounded if there exists a sufficiently large ball RB^d of radius $R < \infty$ such that $K \subseteq RB^d$ for all $K \in \mathfrak{C}$.

Theorem 1 (Blaschke's Selection Theorem). *The metric space of convex bodies is locally compact. That is, given any bounded collection $\mathfrak{C} \subset \mathcal{C}^d$ in the space of convex bodies and a sequence $(K_n) \subset \mathfrak{C}$, there exists a subsequence (K_{n_m}) and a convex body $K \in \mathcal{C}^d$ such that $K_{n_m} \rightarrow K$ in the Hausdorff metric.*

While this result plays an important role in the context of our results on convex bodies, we also require an analogous result for star bodies. A generalization is known for star bodies equipped with the Hausdorff metric [41], however this is not the metric for which relevant continuity properties over the space of star bodies hold. We present here an alternative metric – the radial metric – over star bodies for which it is possible to derive a Blaschke Selection Theorem. In the convex geometry literature, the radial metric is often more natural over the space of star bodies than the Hausdorff metric. Concretely, let $K, L \subset \mathbb{R}^d$ be star bodies. We define the *radial sum* $\dot{+}$ between K and L as $K \dot{+} L := \{x + y : x \in K, y \in L, x = \lambda y\}$; that is, unlike the Minkowski sum, we restrict the pair of vectors to be parallel. The radial sum obeys the relationship $\rho_{K \dot{+} L}(u) := \rho_K(u) + \rho_L(u)$. We denote the *radial metric* between two star bodies K, L as

$$\delta(K, L) := \inf\{\varepsilon \geq 0 : K \subseteq L \dot{+} \varepsilon B^d, L \subseteq K \dot{+} \varepsilon B^d\}.$$

In a similar fashion, the radial metric satisfies $\delta(K, L) := \|\rho_K - \rho_L\|_\infty$.

2.2 Brunn-Minkowski Theory

The Brunn-Minkowski theory for convex bodies [72] combines Minkowski addition and notions of volume to understand the structure of the space of convex bodies and prove important geometric inequalities. Starting with Lutwak [47], a dual Brunn-Minkowski theory for star bodies equipped with radial addition was also developed. This theory studies functionals that measure the size of a star body K relative to another L in a particular way, and studies the extremals of such functionals. The importance of such results in the context our work is that we show our objective of interest has an equivalent interpretation as this functional of interest, and such results analyzing the extremals can aid in characterizing minimizers of our optimization problem. More concretely, the *dual mixed volume* between two star bodies $K, L \in \mathcal{S}^d$ for any $i \in \mathbb{R}$ is defined as:

$$\tilde{V}_i(K, L) := \frac{1}{d} \int_{\mathbb{S}^{d-1}} \rho_K(u)^i \rho_L(u)^{d-i} du.$$

Note that for all i , $\tilde{V}_i(K, K) = \frac{1}{d} \int_{\mathbb{S}^{d-1}} \rho_K(u)^d du = \text{vol}_d(K)$ is the usual d -dimensional volume of K . Lutwak [47] analyzed upper and lower bounds of $\tilde{V}_i(K, L)$ depending on the volumes of K and L , and characterized for which sets K and L would these bounds be attained. An important case for our purposes is the following result when $i = -1$:

Theorem 2 (Theorem 2 in [47]). *For star bodies $K, L \in \mathcal{S}^d$, we have*

$$\tilde{V}_{-1}(K, L)^d \geq \text{vol}_d(K)^{-1} \text{vol}_d(L)^{d+1},$$

and equality holds if and only if K and L are dilates.

3 Minimizers of the Population Risk

In this section, we will describe a setting where the unique optimal regularizer for a data distribution P can be derived and provide a characterization for when this optimal regularizer is convex. As described in the introduction, we aim to find the unique star body that solves the following optimization problem:

$$\arg \min_{K \in \mathcal{S}^d: \text{vol}_d(K)=1} \mathbb{E}_P[\|x\|_K]. \quad (5)$$

The optimal regularizer is convex when the solution of (5) is a convex body. We will see in the following that a unique solution can be obtained by interpreting the objective as a dual mixed volume.

In Section 3.1 we define the fundamental quantity determined by P that is needed to characterize the optimal star body. In Section 3.2 we will state the main result characterizing the solution to (5) and in Section 3.3 we give a condition on P that implies the optimal star body regularizer will be convex. Finally in Section 3.4 we provide examples of different data distributions P for which we apply this theory to obtain a description of the optimal regularizer.

3.1 Radial Function Associated to a Data Distribution

Given a data distribution P , we would like to find a way to summarize the relevant structure that determines an optimal regularizer. Since we have restricted to positively homogeneous functions, the regularizer is characterized by a function on the unit sphere. Thus, given a data distribution, we need a statistic related to the distribution in each unit direction. Recall that star bodies are uniquely determined by the distance from the origin to the boundary in each direction given by the radial function of the star body. Thus, we can hope to characterize an optimal regularizer through a function on the unit sphere obtained from P that determines the radial function of the optimal star body that solves (5).

To be precise, we first restrict to data distributions P that have a density with respect to Lebesgue measure on \mathbb{R}^d . We then define the following function on the unit sphere for a given P with density p :

$$\rho_P(u) := \left(\int_0^\infty r^d p(ru) dr \right)^{1/(d+1)}, \quad u \in \mathbb{S}^{d-1}. \quad (6)$$

Under certain conditions, the function ρ_P is precisely the radial function of a star body L_P in \mathbb{R}^d . Using this interpretation, we will see that ρ_P captures all the relevant information from the data distribution in order to determine a unique optimal regularizer.

Note that many probability distributions may correspond to the same radial function and associated star body. The following is thus a natural question: What aspects of the distribution P characterize the star body L_P ? Intuitively, $\rho_P(u)$ captures how much and how far away on average the mass of the distribution is in a given direction u . Thus the corresponding star body L_P will have larger radius in directions where there is more mass of the distribution and where the mass is farther away from the origin. More precisely, ρ_P^{d+1} is the density of the measure $\mu_P(\cdot) = \mathbb{E}_P[\|x\|_{\ell_2} \mathbf{1}_{\{x/\|x\|_{\ell_2} \in \cdot\}}]$ on the unit sphere. Additionally, ρ_P captures linear transformations of the data. For a random vector with distribution P and its transformation under a linear map ϕ , the radial function corresponding to the distribution of the transformed vector will define a dilate of the linear image $\phi(L_P)$. For example, rotations of the data distribution will be captured by a rotation of L_P . To illustrate what the radial function statistic does not capture, the following example describes a class of probability distributions that all correspond to a constant radial function.

Example 1. Let $f : \mathbb{S}^{d-1} \rightarrow [0, \infty)$ be a bounded function on the unit sphere, and define another function

$$g_f(u) := (1 + f(u)^{d+1})^{\frac{1}{d+1}}, \quad u \in \mathbb{S}^{d-1}.$$

Note that $g_f > f$. Consider the compact set

$$A_f = \{x \in \mathbb{R}^d : f(x/\|x\|_{\ell_2}) \leq \|x\|_{\ell_2} \leq g_f(x/\|x\|_{\ell_2})\},$$




Figure 1: Plots of the support of four different compact sets A_f for different functions f . The uniform distribution over each of these supports will correspond to the same constant radial function, and thus have the same optimal star body regularizer given by a scaled ℓ_2 -norm.

and let P_f be the uniform distribution over A_f . For any f , the radial function associated with P_f satisfies

$$\rho_{P_f}(u)^{d+1} = \frac{1}{\text{vol}_d(A_f)} \int_{f(u)}^{g_f(u)} r^d dr = \frac{1}{(d+1)\text{vol}_d(A_f)} (g_f(u)^{d+1} - f(u)^{d+1}) = \frac{1}{(d+1)\text{vol}_d(A_f)},$$

and the corresponding star body L_{P_f} is a dilate of B^d . Thus, the entire class of distributions $\{P_f : f : \mathbb{S}^{d-1} \rightarrow [0, \infty) \text{ is bounded}\}$ correspond to constant radial functions and under the volume normalization will all correspond to the same optimal star body given by a unit volume ℓ_2 ball. See Figure 1 for examples of the support of P_f for different functions f . Intuitively, each distribution P_f induces the same star body (up to scaling) because in each direction, the average mass of the distribution (weighted by its distance from the origin) is the same, regardless of the behavior of f .

3.2 Optimal Star Body Regularizers

We now state and prove our main result characterizing the optimal star body regularizer for (3). Given a data distribution P of interest, our result establishes that there is a unique star body K_* that achieves the minimum value of the objective over the space of star bodies of constant volume. The conditions on our distribution are fairly mild, and are satisfied for any distribution with a continuous density and finite second moments. They are also satisfied for more general cases, and we will explore examples of such distributions in the sequel. Our main result is stated as follows:

Theorem 3. *Let P be a distribution on \mathbb{R}^d with density p with respect to Lebesgue measure, and assume $\mathbb{E}_P[\|x\|_{\ell_2}] < \infty$. Suppose the function ρ_P as defined in (6) is positive and continuous over the unit sphere. Then there exists a unique star body $L_P \in \mathcal{S}^d$ whose radial function is ρ_P , and the set $K_* := \text{vol}_d(L_P)^{-1/d} L_P$ is the unique solution to the minimization problem (5).*

Proof of Theorem 3. The main idea of our proof is that the population objective can be written as a dual mixed volume with respect to the star body K we are optimizing and the distribution-dependent star body L_P . Once this is established, we can utilize the extremal inequality in Theorem 2 to characterize the minimizers of our objective subject to a fixed volume constraint.

We first construct the star body L_P . Define the function

$$f(x) := \left(\int_0^\infty r^d p(rx) dr \right)^{-1/(d+1)}, \quad x \in \mathbb{R}^d.$$

Note that f is positively homogeneous and non-negative. Define the subset L_P of \mathbb{R}^d by

$$L_P := \{x \in \mathbb{R}^d : f(x) \leq 1\}.$$

The set L_P is star-shaped, because for $x \in L_P$ we have for all $t \in [0, 1]$, $f(tx) = tf(x) \leq t \leq 1$, and thus, $[0, x] \subseteq L_P$. The gauge of L_P is then

$$\begin{aligned} \|x\|_{L_P} &= \inf\{t : x \in tL_P\} = \inf\{t : x/t \in L_P\} \\ &= \inf\{t : f(x/t) \leq 1\} = \inf\{t : f(x) \leq t\} = f(x). \end{aligned}$$

Thus, $f(x)$ is the gauge of L_P . In addition, note that $\rho_P(u) = 1/f(u)$ for $u \in \mathbb{S}^{d-1}$ where ρ_P is defined in (6). Thus ρ_P is the radial function of L_P , and L_P is a star body because ρ_P is positive and continuous.

We now establish that the objective in (5) can be written as a dual mixed volume. Observe that by positive homogeneity of the gauge, we have via integrating in spherical coordinates that

$$\begin{aligned} \mathbb{E}_P[\|x\|_K] &= \int_{\mathbb{R}^d} \|x\|_K p(x) dx = \int_{\mathbb{S}^{d-1}} \int_0^\infty r^d \|u\|_K p(ru) dr du \\ &= \int_{\mathbb{S}^{d-1}} \|u\|_K \rho_P(u)^{d+1} du \end{aligned}$$

where ρ_P is defined in (6). Moreover, since $\|u\|_K = 1/\rho_K(u)$, the definition of the dual mixed volume gives

$$\mathbb{E}_P[\|x\|_K] = \int_{\mathbb{S}^{d-1}} \|u\|_K \rho_P(u)^{d+1} du = \int_{\mathbb{S}^{d-1}} \rho_K(u)^{-1} \rho_P(u)^{d+1} du = d\tilde{V}_{-1}(K, L_P).$$

Hence, we can apply Theorem 2 to show that our objective satisfies

$$\mathbb{E}_P[\|x\|_K] = d\tilde{V}_{-1}(K, L_P) \geq d \text{vol}_d(K)^{-1/d} \text{vol}_d(L_P)^{(d+1)/d}$$

with equality if and only if K and L_P are dilates. Thus, $\mathbb{E}_P[\|x\|_K]$ is minimized over the collection $\{K \in \mathcal{S}^d : \text{vol}_d(K) = 1\}$ by $K_* := \text{vol}_d(L_P)^{-1/d} L_P$. \square

3.3 When is the Optimal Regularizer Convex?

Theorem 3 explicitly characterizes the relationship between the optimal star body K_* and the distribution P . As previously observed, if the star body K_* is in fact a convex body, then the optimal regularizer is a convex function. This connection helps to justify intuition for when a data source modeled by P is amenable to convex regularization. For example, consider the class of log-concave distributions which include many common distributions such as Gaussian, Laplace, and uniform distributions over convex bodies. Since the densities of such distributions have convex level sets, it is natural to guess that a convex regularizer would be most appropriate. If the logarithm of the density is both concave and positively homogeneous of some degree k , then the distribution P falls into the class of probability distributions considered in Example 2 below. In particular, the density is then a function, $\psi(r) = \exp(-r^k)$, of the gauge of a convex body L and in Example 2 we show that the optimal star body will be L by computing the radial function ρ_P . Thus, since L is convex, a convex regularizer is optimal for these log-concave distributions.

We can also use Theorem 3 to characterize when a convex regularizer is optimal for general distributions. In particular, the density of P defines the radial function of K_* , which determines whether or not the set K_* is convex. This observation leads to an important contribution of this work: a condition on the data distribution P such that the optimal regularizer is convex.

Corollary 1. *Let P be a probability measure on \mathbb{R}^d as in Theorem 3. Define the radial function ρ_P as in (6). If $x \mapsto 1/\rho_P(x)$ is a convex function on \mathbb{R}^d , then the optimal K_* as defined in Theorem 3 is a convex body, and the optimal star body regularizer is convex.*

This corollary gives a characterization of data distributions where convexity is the correct structure that should be sought among a class of both convex and nonconvex regularizers, and provides a tool for determining whether a given data source is amenable to convex regularization. We will illustrate in Example 3 a setting where we can use this corollary to determine a parameter range for which the optimal regularizer for a parametric class of distributions is convex.

3.4 Examples of Optimal Regularizers for a Given Distribution

In order to provide more intuition on the above results, we discuss here several additional examples of distributions P and their associated star body L_P with radial function ρ_P defined in (6). The two general classes of distributions we consider are distributions that depend on a compact body and general mixture distributions whose mixture components satisfy the assumptions of our Theorem. When the data distribution

explicitly depends on the gauge of a star body L , then the optimal star body will be a dilate of L . For mixture distributions, the optimal star body will be a special sum of sets, akin to radial addition for star bodies. We now dive deeper into such examples here:

Example 2 (Densities induced by star bodies). *Suppose P has density $p(x) = \psi(\|x\|_L)$, i.e., the density only depends on the gauge function induced by a star body L . If ψ is such that the integral $\int_0^\infty t^d \psi(t) dt < \infty$, then we have that the corresponding radial function of P satisfies*

$$\rho_P(u) := \left(\int_0^\infty r^d \psi(r\|u\|_L) dr \right)^{1/(d+1)} = \rho_L(u) \left[\int_0^\infty t^d \psi(t) dt \right]^{1/(d+1)} = \rho_{c(\psi)L}(u),$$

where $c(\psi) := \left[\int_0^\infty t^d \psi(t) dt \right]^{1/(d+1)}$. Hence the optimal star body K_* is the unit volume dilate of L :

$$K_* = \text{vol}_d(L)^{-1/d} L.$$

There are several examples of distributions of interest that fall under this category, which we outline here.

- Suppose $L = \mathcal{B}_{\ell_q} := \{x \in \mathbb{R}^d : \|x\|_{\ell_q} \leq 1\}$ is the ℓ_q -function unit ball for $0 < q \leq \infty$ and the density of P is given by $p(x) = \psi(\|x\|_{\ell_q})$. Then our optimal set is a dilate of the ℓ_q -ball: $K_* = \text{vol}_d(\mathcal{B}_{\ell_q})^{-1/d} \mathcal{B}_{\ell_q}$. Note that this includes both convex norm balls $q \geq 1$ and nonconvex unit balls $q \in (0, 1)$.
- Another class of distributions P that satisfy our assumptions are uniform distributions over a star body L . Let $\mathbf{1}_E$ denote the indicator function on the event E . In this case, the density is given by

$$p(x) = \frac{\mathbf{1}_{\{x \in L\}}}{\text{vol}_d(L)} = \frac{\mathbf{1}_{\{\|x\|_L \leq 1\}}}{\text{vol}_d(L)},$$

and the radial function ρ_P is

$$\rho_P(u) = \frac{1}{\text{vol}_d(L)^{1/(d+1)}} \left(\int_0^\infty t^d \mathbf{1}_{\{t \leq 1\}} dt \right)^{1/(d+1)} \|u\|_L^{-1} =: c_{d,L} \rho_L(u)$$

where $c_{d,L} := ((d+1)\text{vol}_d(L))^{-1/(d+1)}$. Hence $L_P := c_{d,L}L$, and thus K_* is given by

$$K_* = \text{vol}_d(L)^{-1/d} L.$$

- Suppose $P = \mathcal{N}(0, \Sigma)$ is a multivariate Gaussian distribution with mean zero and covariance matrix Σ in \mathbb{R}^d . Then, a direct calculation shows that the radial function ρ_P is given by

$$\rho_P(u) = \det(2\pi\Sigma)^{-\frac{1}{2(d+1)}} \left(\int_0^\infty t^d e^{-t^2/2} dt \right)^{\frac{1}{d+1}} \|\Sigma^{-1/2}u\|_{\ell_2}^{-1}.$$

Note that the norm $\|\Sigma^{-1/2}u\|_{\ell_2}$ defines the gauge of an ellipsoid E_Σ induced by Σ . Hence the minimizer K_* is the ellipsoid

$$K_* = \kappa_d^{1/d} \det(\Sigma)^{-1/2d} E_\Sigma,$$

where $\|u\|_{E_\Sigma} = \|\Sigma^{-1/2}u\|_{\ell_2}$ and $\kappa_d = \text{vol}_d(B^d)$.

Example 3 (Gaussian mixture model). *Consider the following Gaussian mixture model $P_\varepsilon = \frac{1}{2}\mathcal{N}(0, \Sigma_{\varepsilon,1}) + \frac{1}{2}\mathcal{N}(0, \Sigma_{\varepsilon,2})$ where $\Sigma_{\varepsilon,1} := [1, 0; 0, \varepsilon] \in \mathbb{R}^{2 \times 2}$ and $\Sigma_{\varepsilon,2} := [\varepsilon, 0; 0, 1] \in \mathbb{R}^{2 \times 2}$ for $0 < \varepsilon < 1$ with density $p_\varepsilon(x)$ in \mathbb{R}^2 . Then, we have that the radial function of L_{P_ε} is given by*

$$\rho_{P_\varepsilon}(u)^3 = \int_0^\infty r^2 p_\varepsilon(ru) du = c_{\varepsilon,1} \|\Sigma_{\varepsilon,1}^{-1/2}u\|_{\ell_2}^{-3} + c_{\varepsilon,2} \|\Sigma_{\varepsilon,2}^{-1/2}u\|_{\ell_2}^{-3}$$

where $c_{\varepsilon,i} = \frac{1}{2} \det(2\pi\Sigma_{\varepsilon,i})^{-1/2} \left(\int_0^\infty t^2 e^{-t^2/2} dt \right)$ for $i = 1, 2$. The set L_{P_ε} with radial function ρ_{P_ε} is similar to a radial sum of two ellipsoids induced by $\Sigma_{\varepsilon,1}$ and $\Sigma_{\varepsilon,2}$, which is discussed in the next example. In Figure 2, we plot the solutions L_{P_ε} for different values of ε . We see in the plots that for small ε , L_{P_ε} will be a nonconvex star body, but as ε approaches 1, L_{P_ε} approaches a convex ℓ_2 ball. One can observe that $1/\rho_{P_\varepsilon}(x)$ will be convex for all ε such that $\varepsilon \in (\varepsilon_c^*, 1]$ for some $\varepsilon_c^* \in (0, 1)$. By Corollary 8, the optimal regularizer will be convex for this range of the parameter ε .

Optimal Regularization




Figure 2: We compute the (unnormalized) optimal star body L_{P_ϵ} based on the Gaussian mixture model in Example 3 for $\epsilon \in \{0.01, 0.1, 0.25, 0.75\}$. We note that as ϵ decreases, the individual components of the Gaussian mixture model concentrate more heavily on a lower-dimensional subspace of \mathbb{R}^2 . Similarly, the resulting star body also concentrates closer to the relevant subspaces. Finally, we see that as ϵ increases, there is a critical value between at which L_{P_ϵ} becomes convex. Numerically, this value is between 0.3 and 0.45.

Example 4 (General mixture distributions). *More generally, suppose that our distribution P is a mixture, i.e., let $P = \sum_{i=1}^m w_i P_i$ for weights $w_i > 0$ and distributions P_i in \mathbb{R}^d . Then the associated convex body K_* is a harmonic Blaschke linear combination, as defined in [48], of star bodies L_{P_i} . Each L_{P_i} is the associated star body with radial function ρ_{P_i} in (6). The harmonic Blaschke linear combination $\hat{+}$ between star bodies defines a new star body $K \hat{+} L$ by adding the radial functions of K and L together in the following way:*

$$\frac{\rho_{K \hat{+} L}(u)^{d+1}}{\text{vol}_d(K \hat{+} L)} = \frac{\rho_K(u)^{d+1}}{\text{vol}_d(K)} + \frac{\rho_L(u)^{d+1}}{\text{vol}_d(L)}.$$

4 Existence of Optimal Regularizers for General Data Distributions

In the preceding section, we characterized a unique optimal regularizer under certain conditions on the distribution P . For general P , it remains an open problem to characterize a unique star body that minimizes (5). However, we can prove existence of a solution to this variational problem for general P by restricting the hypothesis class to contain only *well-conditioned* star bodies, defined for any $r > 0$ by:

$$\mathcal{S}^d(r) := \{K \in \mathcal{S}^d : rB^d \subseteq \ker(K)\}.$$

Our results on the existence of minimizers of (5) rely on continuity properties of the objective and compactness of the constraint set contained in $\mathcal{S}^d(r)$. First, in Section 4.1 we will show that the objective is Lipschitz continuous with respect to the radial metric $\delta(\cdot, \cdot)$. Then, in Section 4.2 we show that our constraint sets of interest are compact by establishing Blaschke's Selection Theorem for bounded subsets of $\mathcal{S}^d(r)$ with respect to the radial metric. We will then state and prove the existence results in Section 4.3.

4.1 Continuity Properties of the Objective Functional

We begin by establishing continuity properties of our objective with respect to the radial metric over the collection $\mathcal{S}^d(r)$. When our sets are convex and bounded, we can also establish a Lipschitz bound with respect to the Hausdorff metric.

Proposition 2. *Suppose $K, L \in \mathcal{S}^d(r)$. Then for any $x, y \in \mathbb{S}^{d-1}$, we have*

$$\| \|x\|_K - \|y\|_L \| \leq \frac{1}{r^2} \delta(K, L) + \frac{1}{r} \|x - y\|_{\ell_2}.$$

If, additionally, K and L are convex and $K, L \subseteq RB^d$, we have the following bound:

$$\| \|x\|_K - \|y\|_L \| \leq \frac{R}{r^3} d_H(K, L) + \frac{1}{r} \|x - y\|_{\ell_2}$$

Proof. We first note by the triangle inequality that

$$\| \|x\|_K - \|y\|_L \| \leq \| \|x\|_K - \|x\|_L \| + \| \|x\|_L - \|y\|_L \|.$$

For the second term, note that since $rB^d \subseteq \ker(L)$, we have that by Proposition 6.2 in [68] that $L = \bigcup_{u \in L} L_u$ where each L_u is closed, convex and contains rB^d . Moreover, following the first part of the proof of Theorem 5.1 in [68], since $rB^d \subseteq L_u$ and each L_u is convex, $\|\cdot\|_{L_u}$ is sublinear and Lipschitz with Lipschitz constant bounded by $1/r$. Finally, $\|\cdot\|_L$ coincides with $\inf_{u \in L} \|\cdot\|_{L_u}$ and since the infimum of a family $1/r$ -Lipschitz functions is $1/r$ -Lipschitz, it follows that $\|\cdot\|_L$ is $1/r$ -Lipschitz.

We now focus on establishing a bound on $\| \|x\|_K - \|x\|_L \|$. First, note that since $rB^d \subseteq K$, we have that for any $\|x\|_{\ell_2} = 1$, $\rho_K(x) \geq \rho_{rB^d}(x) = r\rho_{B^d}(x) = r/\|x\|_{\ell_2} = r$. The same holds for L . Moreover, since $\|x\|_K = \rho_K(x)^{-1}$, we get

$$\| \|x\|_K - \|x\|_L \| \leq \frac{1}{\rho_K(x)\rho_L(x)} |\rho_K(x) - \rho_L(x)| \leq \frac{1}{r^2} |\rho_K(x) - \rho_L(x)|.$$

Taking the supremum over $x \in \mathbb{S}^{d-1}$ yields the desired result.

For the convex case, we use the shorthand notation $\gamma := d_H(K, L)$. By the definition of the Hausdorff metric, we have $K \subseteq L + \gamma B^d$ and $L \subseteq K + \gamma B^d$. Consider the ray generated by x . We let $\tilde{x} := x/\|x\|_K$ and $\bar{x} := x/\|x\|_L$ denote the intersection of the ray with the boundaries ∂K and ∂L . Our immediate goal is to show that $\|\bar{x} - \tilde{x}\|_{\ell_2} \leq \gamma R/r$. This will complete the proof.

First, we show that $1/\|x\|_L \leq 1/\|x\|_K + \gamma R/r$. Suppose that $\bar{x} \notin K$. If $\bar{x} \in K$ then $1/\|x\|_L \leq 1/\|x\|_K$, in which case the statement we wish to prove immediately holds. Let u be the projection of \bar{x} onto K . Let D be the cone generated by all lines through \tilde{x} and all points in rB^d . We claim that $u \notin \text{int}(D)$. Suppose on the contrary that $u \in \text{int}(D)$. Consider the line from u to x extended. Using the fact that $u \in \text{int}(D)$, the line intersects $\text{int}(rB^d)$. Pick such a point inside $\text{int}(rB^d)$ and consider an open ball inside $\text{int}(rB^d)$, which is also inside K . By convexity of K , we conclude that there is an open ball containing \tilde{x} also in K . However, this contradicts the original assumption that $\tilde{x} \in \partial K$. Therefore we conclude that $u \notin \text{int}(D)$.

By combining a simple trigonometric argument (such as by similarity) and using the fact that $u \notin \text{int}(D)$, we have

$$\|\bar{x} - \tilde{x}\|_{\ell_2} \leq \|\tilde{x}\|_{\ell_2} \frac{\|\bar{x} - u\|_{\ell_2}}{r} \leq \gamma R/r.$$

By repeating the same argument but with K and L switched, we have $|1/\|x\|_K - 1/\|x\|_L| \leq \gamma R/r$. Since $|1/\|x\|_K - 1/\|x\|_L| = |\rho_K(x) - \rho_L(x)|$, this completes the proof. \square

As a corollary, an application of Jensen's inequality and Proposition 2 shows that the population objective functional is Lipschitz continuous over star bodies with sufficiently large inner width and convex bodies with large inner and outer widths.

Corollary 2. *Let P be a distribution over \mathbb{R}^d such that $\mathbb{E}_P \|x\|_{\ell_2} < \infty$. Then for any $K, L \in \mathcal{S}^d(r)$, we have*

$$|\mathbb{E}_P[\|x\|_K] - \mathbb{E}_P[\|x\|_L]| \leq \frac{\mathbb{E}_P \|x\|_{\ell_2}}{r^2} \delta(K, L).$$

If, additionally, K and L are convex and $K, L \subseteq RB^d$, then

$$|\mathbb{E}_P[\|x\|_K] - \mathbb{E}_P[\|x\|_L]| \leq \frac{R \cdot \mathbb{E}_P \|x\|_{\ell_2}}{r^3} d_H(K, L).$$

4.2 Compactness and Blaschke's Selection Theorem for Star Bodies

In this section we will establish that a constraint set \mathfrak{C} that is a closed and bounded subset of $\mathcal{S}^d(r)$ is compact. We will then prove that the set of volume normalized well-conditioned star bodies

$$\mathcal{S}^d(r, 1) := \{K \in \mathcal{S}^d(r) : \text{vol}_d(K) = 1\}$$

is compact by establishing that this constraint set is a closed and bounded subset of $\mathcal{S}^d(r)$.

To establish compactness of these constraint sets, we must show that sequences in \mathfrak{C} exhibit convergent subsequences with limit inside \mathfrak{C} . In order to showcase such a result, we establish a version of Blaschke's Selection Theorem for a subclass of star bodies equipped with the radial metric. It has been shown that Blaschke's Selection Theorem holds for any bounded, infinite collection of star bodies *with respect to the Hausdorff metric* d_H [41]. However, since our Lipschitz condition only holds with respect to the radial metric, we need to establish a local compactness result that holds for this metric instead.

To prove such a result, we require the following result from [83], which establishes that on the space $\mathcal{S}^d(r)$, the radial and Hausdorff metrics are topologically equivalent to one another.

Theorem 4 (Theorem 3.6 in [83]). *For any $r > 0$, the radial metric δ and the Hausdorff metric d_H are topologically equivalent to one another on $\mathcal{S}^d(r)$. That is, convergent sequences in $(\mathcal{S}^d(r), \delta)$ are the same as convergent sequences in $(\mathcal{S}^d(r), d_H)$.*

With this result in hand, we can then prove Blaschke's Selection Theorem also holds with respect to the radial metric, specifically over the set $\mathcal{S}^d(r)$:

Theorem 5 (Blaschke's Selection Theorem for the Radial Metric). *Fix $0 < r < \infty$ and let \mathfrak{C} be a bounded and closed subset of $\mathcal{S}^d(r)$. Let (K_i) be a sequence of star bodies in \mathfrak{C} . Then (K_i) has a subsequence that converges in the radial and Hausdorff metric to a star body $K \in \mathfrak{C}$.*

Proof. From [41], since (K_i) is a sequence of star bodies contained in a ball, we know that there exists a subsequence, call it (K_{i_n}) , and a star body K such that $d_H(K_{i_n}, K) \rightarrow 0$ as $i_n \rightarrow \infty$. Since \mathfrak{C} is closed, $K \in \mathfrak{C}$. Finally, to complete the proof, we note that since d_H and δ are topologically equivalent on $\mathcal{S}^d(r)$ via Theorem 4, we have that convergence in d_H is equivalent to convergence in δ . Hence the subsequence (K_{i_n}) of (K_i) satisfies $\delta(K_{i_n}, K) \rightarrow 0$ as $i_n \rightarrow \infty$ where $K \in \mathfrak{C}$. \square

This result aids in establishing favorable compactness properties of our constraint set \mathfrak{C} . We collect such properties as a corollary here:

Corollary 3. *For $0 < r < \infty$, let \mathfrak{C} be a closed and bounded subset of $\mathcal{S}^d(r)$. Then, the metric space (\mathfrak{C}, δ) is compact and for every $\varepsilon > 0$, there exists a finite set of star bodies $\mathcal{S}_\varepsilon \subset \mathfrak{C}$ such that*

$$\sup_{K \in \mathfrak{C}} \inf_{L \in \mathcal{S}_\varepsilon} \delta(K, L) \leq \varepsilon.$$

The same conclusion holds with respect to the Hausdorff metric d_H .

Proof. By Theorem 5, we have that the space (\mathfrak{C}, δ) is sequentially compact, as every sequence has a convergent subsequence with limit in \mathfrak{C} . On metric spaces, sequential compactness is equivalent to compactness. Since the space is compact, it is totally bounded, thus guaranteeing the existence of a finite ε -net for every $\varepsilon > 0$ as desired. Theorem 4 guarantees the extension to the Hausdorff metric holds. \square

We now consider the particular case when $\mathfrak{C} = \mathcal{S}^d(r, 1)$. In order to establish compactness using Corollary 3, we need to establish that this subset is closed and bounded. We first show that the volume constraint in $\mathcal{S}^d(r, 1)$ is sufficient to establish that the set is bounded.

Lemma 1. *For any $r > 0$, the collection $\mathcal{S}^d(r, 1)$ is a bounded subset of $\mathcal{S}^d(r)$. In particular, for $R_r := \frac{d+1}{r^{d-1}\kappa_{d-1}}$ where $\kappa_{d-1} := \text{vol}_{d-1}(B^{d-1})$,*

$$\mathcal{S}^d(r, 1) \subseteq \{K \in \mathcal{S}^d(r) : K \subseteq R_r B^d\}.$$

Proof. Let $K \in \mathcal{S}^d(r, 1)$ and define $R(K) := \min\{R \geq 0 : K \subseteq RB^d\}$. For each K , $R(K) < \infty$, however we need to obtain a uniform bound R_r such that $R(K) \leq R_r$ for all $K \in \mathcal{S}^d(r, 1)$. Now, let $x \in \partial K$ be such that $\|x\|_{\ell_2} = R(K)$, and consider the cone $C(K) := \text{conv}(\{rB^{d-1} \cap x^\perp, x\})$ where $\text{conv}(\cdot)$ denotes the convex hull. Since $rB^{d-1} \subseteq \ker(K)$, $C(K) \subseteq K$. By the volume constraint $\text{vol}_d(K) = 1$, we have that $\text{vol}_d(C(K)) \leq \text{vol}_d(K) = 1$. This then implies that

$$\text{vol}_d(C(K)) = R(K) \cdot \frac{\text{vol}_{d-1}(rB^{d-1})}{d+1} \leq 1.$$

Since this bound holds for any $K \in \mathcal{S}^d(r, 1)$, we conclude by observing the uniform bound

$$\sup_{K \in \mathcal{S}^d(r, 1)} R(K) \leq \frac{d+1}{r^{d-1} \text{vol}_{d-1}(B^{d-1})} =: R_r.$$

□

Lastly we show that the constraint set $\mathcal{S}^d(r, 1)$ is closed under the Hausdorff topology.

Lemma 2. *For any $0 < r < \infty$, the set $\mathcal{S}^d(r, 1)$ is closed under the Hausdorff topology.*

Proof. Let (L_i) be a sequence in $\mathcal{S}^d(r, 1)$ such that $L_i \rightarrow L$ in the Hausdorff metric. We claim that $L \in \mathcal{S}^d(r, 1)$. We first show $rB^d \subseteq \ker(L)$. Note that $rB^d \subseteq \ker(L_i)$ for each i so that $\ker(L_i)$ has non-empty interior. Moreover, since each L_i is closed, so is $\ker(L_i)$. Lastly, the sequence (L_i) is bounded uniformly by Lemma 1, so the sequence $(\ker(L_i))$ is as well. Thus, by Blaschke's Selection Theorem for convex bodies (Theorem 1), there exists a subsequence, call it $(\ker(L_{i_n}))$, that is convergent in the Hausdorff metric to some convex body \tilde{L} . We claim that $\tilde{L} \subseteq \ker(L)$. Fix $v \in \tilde{L}$ and let $x \in L$ be arbitrary. We will show that $[v, x] \subseteq L$. Since $\ker(L_{i_n}) \rightarrow \tilde{L}$ and $L_i \rightarrow L$ in the Hausdorff metric, there exists sequences (v_{i_n}) and (x_{i_n}) such that $v_{i_n} \in \ker(L_{i_n})$ and $x_{i_n} \in L_{i_n}$ for each i_n and $v_{i_n} \rightarrow v$ and $x_{i_n} \rightarrow x$. Then the line segments $[v_{i_n}, x_{i_n}] \rightarrow [v, x]$ in the Hausdorff metric where $[v_{i_n}, x_{i_n}] \subseteq L_{i_n}$ since each $v_{i_n} \in \ker(L_{i_n})$. Hence we must have that $[v, x] \subseteq L$. Since $x \in L$ was arbitrary, we must have that $v \in \ker(L)$ as desired so $\tilde{L} \subseteq \ker(L)$. Thus, since $rB^d \subseteq \ker(L_{i_n})$ for all i_n and $\ker(L_{i_n}) \rightarrow \tilde{L}$, we obtain $rB^d \subseteq \tilde{L} \subseteq \ker(L)$ as claimed. We finally note that since each L_i satisfies $\text{vol}_d(L_i) = 1$, continuity of the volume functional gives $\text{vol}_d(L) = 1$. We conclude that $L \in \mathcal{S}^d(r, 1)$. □

4.3 Existence of Minimizers

We now state and prove our main results on the existence of minimizers using the results above.

Theorem 6. *Suppose P is such that $\mathbb{E}_P[\|x\|_{\ell_2}] < \infty$. If $\mathfrak{C} \subseteq \mathcal{S}^d(r)$ is closed and bounded, then there exists a solution to the following optimization problem:*

$$\arg \min_{K \in \mathfrak{C}} \mathbb{E}_P[\|x\|_K]. \tag{7}$$

In particular, for fixed $0 < r < \infty$, a solution to the following optimization problem exists:

$$\arg \min_{K \in \mathcal{S}^d(r): \text{vol}_d(K)=1} \mathbb{E}_P[\|x\|_K]. \tag{8}$$

Proof. The result is a direct consequence of the fact our objective is continuous and the constraint set is compact. Indeed, using Corollary 2, we have that the objective functional $K \mapsto \mathbb{E}_P[\|x\|_K]$ is $\frac{\mathbb{E}_P[\|x\|_{\ell_2}]}{r^2}$ -Lipschitz over \mathfrak{C} . Moreover, the constraint set \mathfrak{C} is compact by Corollary 3. As we are minimizing a continuous functional over a compact domain, the existence of minimizers is immediate. Combining Lemma 1 and Lemma 2 shows that $\mathcal{S}^d(r, 1)$ is a closed and bounded subset of $\mathcal{S}^d(r)$, meaning that the existence of a minimizer holds in this case as well. □

The existence of a minimizer to (7) can be applied to the case when \mathfrak{C} is restricted to well-conditioned convex bodies defined by

$$\mathcal{C}^d(r) := \{K \in \mathcal{C}^d : rB^d \subseteq K\}.$$

We highlight this case because for computational and modeling considerations, one may wish to restrict to searching over convex regularizers. For instance, enforcing convexity can help ensure that the optimization framework to solve downstream tasks such as inverse problems will be more computationally tractable. There may also be further domain knowledge suggesting that convexity is the correct structure for the data distribution of interest. We will also consider in Sections 5 and 6 empirical and population risk minimizers over natural subclasses of convex bodies, for which the following result ensures existence.

Corollary 4. Fix $0 < r < \infty$ and suppose P is such that $\mathbb{E}_P[\|x\|_{\ell_2}] < \infty$. Suppose $\mathfrak{C} \subseteq \mathcal{C}^d(r)$ is closed and bounded. Then a minimizer to the optimization problem exists:

$$\min_{K \in \mathfrak{C}} \mathbb{E}_P[\|x\|_K].$$

Proof. Corollary 4 follows from Proposition 6 because $\mathcal{C}^d(r)$ is a closed subset of $\mathcal{S}^d(r)$, but it can also be proved directly by appealing to the corresponding results for convex bodies equipped with the Hausdorff metric. \square

At this point, it is natural to ask how Theorem 6 relates to the results in Section 3. In particular, if P satisfies the conditions of Theorem 3, when is the solution to (5) equal to the solution of (8) for some $r > 0$? The next result answers this question using properties of the radial function ρ_P .

Proposition 3. Let P satisfy the conditions of Theorem 3, and define ρ_P as in (6). If $x \mapsto 1/\rho_P(x)$ is Lipschitz, then there exists an $r > 0$ such that the unique minimizer of (5) is also the unique minimizer of (8).

Proof. It suffices to prove that there exists $r > 0$ such that $L_P \in \mathcal{S}^d(r)$. Indeed, Theorem 6.1 in [68] implies that there exists $r > 0$ such that $rB^d \subseteq \ker(L_P)$ if and only if the gauge $\|\cdot\|_{L_P} = 1/\rho_P(\cdot)$ is Lipschitz. \square

5 Convergence of Minimizers

In this section, we study the behavior of the minimizers with increasing data. Our main result is to show that the empirical minimizers converge to the minimizer of the population risk. To state our main result, we introduce the following notation. For a distribution P over \mathbb{R}^d , let P_m denote the empirical distribution of P over m i.i.d. observations drawn from P . Let $F(K; P) := \mathbb{E}_P[\|x\|_K]$. Then $F(K; P_m)$ and $F(K; P)$ denote the empirical risk functional and population risk functional, respectively. As in the previous section, we will restrict to the class of well-conditioned star bodies $\mathcal{S}^d(r)$ and the bounded subset $\mathcal{S}^d(r, 1)$ in order to obtain the required continuity and compactness properties associated to our optimization problem for the desired convergence to hold.

We first state the main result here:

Theorem 7. Let P be a distribution over \mathbb{R}^d such that $\mathbb{E}_P\|x\|_{\ell_2} < \infty$ and fix $0 < r < \infty$. Then the sequence of minimizers $(K_m^*) \subseteq \mathcal{S}^d(r, 1)$ of $F(K; P_m)$ over $\mathcal{S}^d(r, 1)$ has the property that any convergent subsequence converges in the radial and Hausdorff metric to a minimizer of the population risk almost surely:

$$\text{any convergent } (K_{m_\ell}^*) \text{ satisfies } K_{m_\ell}^* \rightarrow K_* \in \arg \min_{K \in \mathcal{S}^d(r, 1)} F(K; P).$$

Moreover, a convergent subsequence of (K_m^*) exists.

Note that our result states that any limit point of (K_m^*) is a minimizer of the population risk, which is vacuously true if (K_m^*) does not have a convergent subsequence. However, since we have established $\mathcal{S}^d(r, 1)$ is a closed and bounded subset of $\mathcal{S}^d(r)$ and $(K_m^*) \subseteq \mathcal{S}^d(r, 1)$, an application of Theorem 5 proves that a convergent subsequence of (K_m^*) must exist.

If we restrict ourselves to searching over convex bodies $\mathcal{C}^d(r, 1) := \{K \in \mathcal{C}^d : rB^d \subseteq K, \text{vol}_d(K) = 1\}$, we can obtain the following result focusing on convex bodies, whose proof is a direct consequence of Theorem 7 and the fact that the radial topology and Hausdorff topology are equivalent for convex bodies:

Theorem 8. Let P be a distribution over \mathbb{R}^d such that $\mathbb{E}_P\|x\|_{\ell_2} < \infty$. Then the sequence of minimizers $(K_m^*) \subseteq \mathcal{C}^d(r, 1)$ of $F(K; P_m)$ over $\mathcal{C}^d(r, 1)$ has the property that any convergent subsequence converges in the radial and Hausdorff metric to a minimizer of the population risk almost surely:

$$\text{any convergent } (K_{m_\ell}^*) \text{ satisfies } K_{m_\ell}^* \rightarrow K_* \in \arg \min_{K \in \mathcal{C}^d(r, 1)} F(K; P).$$

Moreover, a convergent subsequence of (K_m^*) exists.

An outline of this section is as follows. First, in Section 5.1, we describe the core mathematical tool needed to prove convergence of minimizers, namely the notion of Γ -convergence from variational analysis. Then, in Section 5.2, we establish a result that will be helpful in establishing the requirements of Γ -convergence. In particular, we prove uniform convergence of the empirical risk objective to the population objective over $\mathcal{S}^d(r, 1)$. This result requires our continuity and compactness results that were used to prove existence of minimizers in Section 4. Finally, combining results from both sections, we prove Theorem 7 in Section 5.3 by arguing that the minimizers of the empirical risk must converge to a minimizer of the population risk. Such tools will also be shown to establish robustness guarantees in Section 5.4 for our problem, in the sense that if we obtain a regularizer on data convolved with Gaussian noise, the regularizer converges to the optimal regularizer for clean data as the amount of noise decreases.

5.1 Γ -convergence

Our key technique for establishing Theorem 7 is a tool from variational analysis known as Γ -convergence [13]. In essence, Γ -convergence allows one to conclude the convergence of minimizers of a sequence of functionals to a minimizer of a particular limiting functional. In addition, because our objects of interest are *sets* and they reside in a more general metric space rather than Euclidean space, greater care is needed in understanding the limits of minimizers.

We first state the definition of Γ -convergence here and cite some useful results that will be needed in our proofs.

Definition 1. *Let (F_i) be a sequence of functions $F_i : X \rightarrow \mathbb{R}$ on some topological space X . Then we say that F_i Γ -converges to a limit F and write $F_i \xrightarrow{\Gamma} F$ if the following conditions hold:*

- *For any $x \in X$ and any sequence (x_i) such that $x_i \rightarrow x$, we have*

$$F(x) \leq \liminf_{i \rightarrow \infty} F_i(x_i).$$

- *For any $x \in X$, we can find a sequence $x_i \rightarrow x$ such that*

$$F(x) \geq \limsup_{i \rightarrow \infty} F_i(x_i).$$

In fact, if the first condition holds, then the second condition could be taken to be the following: for any $x \in X$, there exists a sequence $x_i \rightarrow x$ such that $\lim_{i \rightarrow \infty} F_i(x_i) = F(x)$.

In addition to Γ -convergence, we also require the notion of equi-coercivity, which states that minimizers of a sequence of functions are attained over a compact domain.

Definition 2. *A family of functions $F_i : X \rightarrow \mathbb{R}$ is equi-coercive if for all α , there exists a compact set $K_\alpha \subseteq X$ such that $\{x \in X : F_i(x) \leq \alpha\} \subseteq K_\alpha$.*

Finally, this notion combined with Γ -convergence guarantees convergence of minimizers, which is known as the Fundamental Theorem of Γ -convergence [13].

Proposition 4 (Fundamental Theorem of Γ -Convergence). *If $F_i \xrightarrow{\Gamma} F$ and the family (F_i) is equi-coercive, then the every limit point of the sequence of minimizers (x_i) where $x_i \in \arg \min_{x \in X} F_i(x)$ converges to some $x \in \arg \min_{x \in X} F(x)$.*

The goal of the subsequent sections will be in deriving the necessary conditions to establish Γ -convergence of the empirical risk to the population risk. Then, we will show how our compactness results in Section 4.2 in conjunction with the results in Section 5.2 can establish the requirements of Proposition 4.

5.2 Uniform Convergence of Empirical Risk Objective

In order to show that our empirical risk functional Γ -converges to the population risk, we require a result that shows the empirical risk uniformly converges to the population risk in the limit of increasing data. To obtain this, we appeal to the following Uniform Strong Law of Large Numbers (ULLN) [60], which shows that if there exists an ε -net over our hypothesis class, then we can establish uniform convergence of the empirical risk to the population risk.

Theorem 9 (Theorem 3 in [60]). *Let Q be a probability measure, and let Q_m be the corresponding empirical measure. Let \mathcal{G} be a collection of Q -integrable functions. Suppose that for every $\varepsilon > 0$ there exists a finite collection of functions \mathcal{G}_ε such that for every $g \in \mathcal{G}$ there exists $\bar{g}, \underline{g} \in \mathcal{G}_\varepsilon$ satisfying (i) $\underline{g} \leq g \leq \bar{g}$, and (ii) $\mathbb{E}_Q[\bar{g} - \underline{g}] < \varepsilon$. Then $\sup_{g \in \mathcal{G}} |\mathbb{E}_{Q_m}[g] - \mathbb{E}_Q[g]| \rightarrow 0$ almost surely.*

We now state and prove the uniform convergence result.

Theorem 10. *Fix $0 < r < \infty$ and let P be a distribution on \mathbb{R}^d such that $\mathbb{E}_P\|x\|_{\ell_2} < \infty$. Then, we have strong consistency in the sense that*

$$\sup_{K \in \mathcal{S}^d(r,1)} |F(K; P_m) - F(K; P)| \rightarrow 0 \text{ as } m \rightarrow \infty \text{ almost surely.}$$

Proof of Theorem 10. By Corollary 2, we have that the map $K \mapsto \mathbb{E}_P[\|x\|_K]$ is CM_P -Lipschitz over $\mathcal{S}^d(r, 1)$ where $C = C(r) := 1/r^2$ and $M_P := \mathbb{E}_P\|x\|_{\ell_2} < \infty$. Now, consider the set of functions $\mathcal{G} := \{\|\cdot\|_K : K \in \mathcal{S}^d(r, 1)\}$. We will show that we can construct a finite set of functions that approximate $\|\cdot\|_K$ for any $K \in \mathcal{S}^d(r, 1)$ via an ε -covering argument. This will allow us to apply Theorem 9. By Corollary 3, there exists an ε -net $\mathcal{S}_\varepsilon \subseteq \mathcal{S}^d(r, 1)$ of finite cardinality. For fixed $\varepsilon > 0$, construct a η -cover \mathcal{S}_η of $\mathcal{S}^d(r, 1)$ such that $\sup_{K \in \mathcal{S}^d(r,1)} \min_{L \in \mathcal{S}_\eta} \delta(K, L) \leq \eta$ where $\eta \leq \varepsilon/(2CM_P)$. Define the following sets of functions:

$$\mathcal{G}_{\eta,-} := \{(\|\cdot\|_K - CM_P\eta)_+ : K \in \mathcal{S}_\eta\} \text{ and } \mathcal{G}_{\eta,+} := \{\|\cdot\|_K + CM_P\eta : K \in \mathcal{S}_\eta\}.$$

Let $\|\cdot\|_K \in \mathcal{G}$ be arbitrary. Let $K_0 \in \mathcal{S}_\eta$ be such that $\delta(K, K_0) \leq \eta$. Define $\bar{f} = \|\cdot\|_{K_0} + CM_P\eta \in \mathcal{G}_{\eta,+}$ and $\underline{f} = (\|\cdot\|_{K_0} - CM_P\eta)_+ \in \mathcal{G}_{\eta,-}$. It follows that $\underline{f} \leq f \leq \bar{f}$. Moreover, by our choice of η , we have that

$$\mathbb{E}_P[\bar{f} - \underline{f}] \leq 2CM_P\eta \leq \varepsilon.$$

Thus, the conditions of Theorem 9 are met and we get

$$\sup_{K \in \mathcal{S}^d(r,1)} |F(K; P_m) - F(K; P)| = \sup_{K \in \mathcal{S}^d(r,1)} |\mathbb{E}_{P_m}[\|x\|_K] - \mathbb{E}_P[\|x\|_K]| \rightarrow 0 \text{ as } m \rightarrow \infty \text{ a.s.}$$

□

5.3 Proof of Theorem 7

We first establish the two requirements of Γ -convergence. For the first, fix $K \in \mathcal{S}^d(r, 1)$ and consider a sequence $K_m \rightarrow K$. Then we have that

$$\begin{aligned} F(K; P) &= F(K; P) - F(K; P_m) + F(K; P_m) - F(K_m; P_m) + F(K_m; P_m) \\ &\leq |F(K; P) - F(K; P_m)| + |F(K; P_m) - F(K_m; P_m)| + F(K_m; P_m) \\ &\leq \sup_{K \in \mathcal{S}^d(r,1)} |F(K; P) - F(K; P_m)| + |F(K; P_m) - F(K_m; P_m)| + F(K_m; P_m). \end{aligned} \quad (9)$$

By Theorem 10, we have that the first term goes to 0 as m goes to ∞ almost surely. We now show that $|F(K; P_m) - F(K_m; P_m)| \rightarrow 0$ as $m \rightarrow \infty$ almost surely. To see this, observe that by Corollary 2 we have

$$|F(K; P_m) - F(K_m; P_m)| \leq \frac{\mathbb{E}_{P_m}[\|x\|_{\ell_2}]}{r^2} \delta(K, K_m).$$

Since $\mathbb{E}_{P_m}[\|x\|_{\ell_2}] \rightarrow \mathbb{E}_P[\|x\|_{\ell_2}] < \infty$ and $\delta(K, K_m) \rightarrow 0$ as $m \rightarrow \infty$ almost surely, we attain $|F(K; P_m) - F(K_m; P_m)| \rightarrow 0$. Thus, taking the limit inferior of both sides in equation (9) yields

$$F(K; P) \leq \liminf_{m \rightarrow \infty} F(K_m; P_m).$$

For the second requirement, we exhibit a realizing sequence so let $K \in \mathcal{S}^d(r, 1)$ be arbitrary. By Corollary 3, for any $m \geq 1$, there exists a finite $\frac{1}{m}$ -net $\mathcal{S}_{1/m}$ of $\mathcal{S}^d(r, 1)$ in the radial metric δ . Construct a sequence $(K_m) \subset \mathcal{S}^d(r, 1)$ such that for each m , $K_m \in \mathcal{S}_{1/m}$ and satisfies $\delta(K_m, K) \leq 1/m$. Hence this sequence satisfies $K_m \rightarrow K$ in the radial metric and $K_m \in \mathcal{S}^d(r, 1)$ for all $m \geq 1$. Hence we can apply Theorem 10 to get

$$|F(K_m; P_m) - F(K; P)| \leq |F(K_m; P_m) - F(K; P_m)| + |F(K; P_m) - F(K; P)| \rightarrow 0 \text{ as } m \rightarrow \infty \text{ a.s.}$$

so $\lim_{m \rightarrow \infty} F(K_m; P_m) = F(K; P)$.

Now, we show that $(F(\cdot; P_m))$ is equi-coercive on $\mathcal{S}^d(r, 1)$. In fact, this follows directly from Theorem 5, the variant of Blaschke's Selection Theorem we proved for the radial metric. Thus equi-coerciveness of the family $(F(\cdot; P_m))$ trivially holds over $\mathcal{S}^d(r, 1)$. As a result, applying Proposition 4 to the family $F(\cdot; P_m) : \mathcal{S}^d(r, 1) \rightarrow \mathbb{R}$, if we define the sequence of minimizers

$$K_m^* \in \arg \min_{K \in \mathcal{S}^d(r, 1)} F(K; P_m)$$

we have that any limit point of K_m^* converges to some

$$K_* \in \arg \min_{K \in \mathcal{S}^d(r, 1)} F(K; P)$$

almost surely, as desired. The existence of a convergent subsequence of (K_m^*) follows from $\mathcal{S}^d(r, 1)$ being a closed and bounded subset of $\mathcal{S}^d(r)$ and Theorem 5.

5.4 Robustness Guarantees

Our results utilizing Γ -convergence can also be used to establish robustness guarantees on optimal regularizers for uncorrupted and noisy data. In particular, we are interested in deviations between regularizers K that are optimal over uncorrupted data $y = x$ for $x \sim P$ versus regularizers K_σ that are optimal over noisy data $y = x + z$ where $x \sim P$ and $z \sim \mathcal{N}(0, \sigma^2 I_d)$. A natural question to ask is whether K_σ converges to K as the amount of noise $\sigma \rightarrow 0$. To understand this question, we first consider characterizing the deviation in the population risk under an uncorrupted distribution and a noisy distribution.

To state our results, we first require the following metric on the space of probability distributions. Let $T(P, Q)$ denote all joint distributions γ for (X, Y) that have marginals P and Q . Define the p -Wasserstein distance between two probability distributions P and Q as

$$W_p(P, Q) := \left(\inf_{\gamma \in T(P, Q)} \mathbb{E} \|X - Y\|^p \right)^{1/p}, \quad p \geq 1.$$

A particularly useful result for our purposes is in the case $p = 1$, as the 1-Wasserstein distance has the following dual characterization [92]:

$$W_1(P, Q) := \sup_{f \in \text{Lip}(1)} |\mathbb{E}_P[f(y)] - \mathbb{E}_Q[f(x)]|.$$

Here, $\text{Lip}(R)$ denotes the set of all R -Lipschitz functions from \mathbb{R}^d to \mathbb{R} :

$$\text{Lip}(R) := \left\{ f : \mathbb{R}^d \rightarrow \mathbb{R} : \sup_{x \neq y} \frac{|f(x) - f(y)|}{\|x - y\|_{\ell_2}} \leq R \right\}.$$

Suppose we have a probability distribution P and define $P_\sigma := P * \mathcal{N}(0, \sigma^2 I_d)$ as the convolution of P and $\mathcal{N}(0, \sigma^2 I_d)$:

$$y \sim P_\sigma \iff y = x + z, \quad x \sim P, \quad z \sim \mathcal{N}(0, \sigma^2 I_d).$$

Optimal Regularization

We first characterize the deviation in population risk between P and P_σ as a function of σ and then aim to understand properties of star bodies learned on noisy data as the noise vanishes. We first establish the following result on the population risk, which shows that we can uniformly bound the deviation in population risk as a function of σ if the collection of star bodies have sufficiently large inner width:

Theorem 11. *Let P be a probability distribution in \mathbb{R}^d and define $P_\sigma := P * \mathcal{N}(0, \sigma^2 I_d)$ for some $\sigma > 0$. Suppose there exists an $0 < r < \infty$ such that $\mathfrak{C} \subseteq \mathcal{S}^d(r)$. Then we have that*

$$\sup_{K \in \mathfrak{C}} |F(K; P) - F(K; P_\sigma)| \leq \sigma \frac{\sqrt{d}}{r}.$$

To prove this, we first consider the general question of how far $F(\cdot; P)$ deviates from $F(\cdot; Q)$ for two distributions P and Q . We prove the following:

Proposition 5. *Let P and Q be two probability distributions in \mathbb{R}^d . Suppose there exists an $0 < r < \infty$ such that $\mathfrak{C} \subseteq \mathcal{S}^d(r)$. Then the following bound holds:*

$$\sup_{K \in \mathfrak{C}} |F(K; P) - F(K; Q)| \leq \frac{1}{r} W_1(P, Q).$$

Proof. Fix $K \in \mathfrak{C}$. It was proven in Proposition 2 that if $rB^d \subseteq \ker(K)$, then the mapping $x \mapsto \|x\|_K$ is $1/r$ -Lipschitz so that $\|\cdot\|_K \in \text{Lip}(1/r)$. Thus, using the dual characterization of $W_1(P, Q)$, we obtain the bound

$$\begin{aligned} |F(K; P) - F(K; Q)| &= |\mathbb{E}_P[\|y\|_K] - \mathbb{E}_Q[\|x\|_K]| \\ &\leq \sup_{f \in \text{Lip}(1/r)} |\mathbb{E}_P[f(y)] - \mathbb{E}_Q[f(x)]| \\ &\leq \frac{1}{r} W_1(P, Q). \end{aligned}$$

Taking the supremum over all $K \in \mathfrak{C}$ completes the proof. □

Now we can prove the main Theorem.

Proof of Theorem 11. By our previous proposition, we simply need to compute $W_1(P, P_\sigma)$. But it is well-known [97] that $W_1(P, P_\sigma) \leq \sigma\sqrt{d}$. Applying the uniform bound and using $W_1(P, P_\sigma) \leq \sigma\sqrt{d}$ completes the proof. □

Based on our quantitative convergence bound, we can prove the following:

Theorem 12. *Fix $0 < r < \infty$ and let P be a distribution on \mathbb{R}^d such that $\mathbb{E}_P \|x\|_{\ell_2} < \infty$. Let (σ_t) be a non-negative, monotone decreasing sequence converging to 0 and set $P_{\sigma_t} := P * \mathcal{N}(0, \sigma_t^2 I_d)$. Then we have that the sequence of minimizers on noisy data*

$$K_t \in \arg \min_{K \in \mathcal{S}^d(r, 1)} F(K; P_{\sigma_t})$$

has the property that any convergent subsequence converges to a minimizer on uncorrupted data in the radial metric and Hausdorff metric:

$$\text{any convergent } (K_{t_\ell}) \text{ satisfies } K_{t_\ell} \rightarrow K_* \in \arg \min_{K \in \mathcal{S}^d(r, 1)} F(K; P).$$

Moreover, a convergent subsequence of (K_t) exists.

Proof. The proof of this result follows similarly to the proof of Theorem 7 by using Theorem 11 to help establish Γ -convergence of the family $(F(\cdot; P_{\sigma_t}))_t$ to $F(\cdot; P)$ over $\mathcal{S}^d(r, 1)$. □

6 Statistical Learning Guarantees

As a final application of our results, we showcase sample complexity guarantees to control the deviation from the empirical risk to the population risk. In many data science contexts, it is imperative to understand the tradeoff between the expressiveness of the hypothesis class and the number of samples needed to efficiently solve the problem. Specifically, one might wish to ask, how many samples does it take to construct a regularizer (from a structured family) that approximates the best possible choice from a hypothesis class to some target accuracy – such questions are often phrased as *sample complexity* problems. In our framework, this requires an analysis of the particular collection of star bodies $\mathfrak{C} \subseteq \mathcal{S}^d$ being optimized over and the behavior of the objective $K \mapsto \|\cdot\|_K$ over a finite number of samples drawn from a data distribution P . We showcase uniform convergence guarantees over the collection \mathfrak{C} , where the difference between the empirical risk and the population risk is controlled with overwhelming probability by a quantity $\eta(m, \mathfrak{C}, P)$ that depends on the number of samples m , the distribution P , and the complexity of the collection of bodies \mathfrak{C} :

$$\sup_{K \in \mathfrak{C}} |F(K; P_m) - F(K; P)| \leq \eta(m, \mathfrak{C}, P) \tag{10}$$

As the number of samples increases, this approximation improves where the rate depends on the data distribution and structure of the bodies. As a consequence, we can establish that empirical minimizers $K_m^* \in \arg \min_{K \in \mathfrak{C}} F(K; P_m)$ obtain population risk not far from the population minimizer $K_* \in \arg \min_{K \in \mathfrak{C}} F(K; P)$ in the sense that with high probability

$$F(K_m^*; P) \leq F(K_*; P) + 2\eta(m, \mathfrak{C}, P).$$

At a high-level, our generalization results require the following key ingredients to specify a uniform convergence bound:

- First, identify a family of star/convex bodies \mathfrak{C} of interest.
- Show that the functional $K \mapsto \|\cdot\|_K$ is well-behaved in the sense that it is $C_{\mathfrak{C}}$ -Lipschitz over \mathfrak{C} with respect to a metric $D(\cdot, \cdot)$. This will allow us to exploit some of our previous continuity results for our objective functional over the space of well-conditioned star bodies.
- Characterize the “complexity” of \mathfrak{C} . While there are many notions of complexity that have been studied in statistical learning theory, one notion that will be particularly useful for our purposes is the notion of a covering number of a subset of a metric space. For a metric space (X, D) , we define the ε -covering number of a subset $T \subset X$ with respect to the metric D as the smallest number of balls of radius ε that covers T :

$$N(T, D, \varepsilon) := \min \left| \left\{ \mathcal{T} \subset T : T \subset \bigcup_{q \in \mathcal{T}} \{x \in T : D(x, q) \leq \varepsilon\} \right\} \right|.$$

In Section 6.1, we describe a general uniform convergence bound of the form (10) that quantifies the rate of convergence of the empirical risk to the population risk over a general class of bodies \mathfrak{C} . We will then showcase the utility of this bound by showing how it can be applied to several classes of bodies of interest in applications. First, we establish a bound on the class of star bodies given as unions of convex bodies with large inner width in Section 6.2. To our knowledge, such metric entropy estimates for this class of star bodies is new. Then, we apply this result to parametrized families of sets. The first class we consider is ellipsoids in Section 6.3. Afterwards, in Section 6.4, we discuss polytopes and how such a bound can be applied to dictionary learning. Finally, we showcase an application to semidefinite regularizers in Section 6.5, which are regularizers induced by linear images of Schatten p -norm balls.

6.1 General Uniform Convergence Result

Given the ingredients described previously, we can prove a general uniform convergence result that shows when the functional $K \mapsto \|\cdot\|_K$ is Lipschitz over a collection of star bodies \mathfrak{C} , we can obtain generalization

bounds based on the complexity of the collection \mathfrak{C} and the data distribution P . We will show that this general theorem can then be applied to many classes of bodies of interest to showcase the utility of our result.

Theorem 13. *Let P be a distribution over \mathbb{R}^d such that for $x \sim P$, $\|x\|_{\ell_2} \leq 1$ almost surely. Consider a subset $\mathfrak{C} \subset \mathcal{S}^d$ such that $r_{\mathfrak{C}}B^d \subseteq K$ for all $K \in \mathfrak{C}$ for some $0 < r_{\mathfrak{C}} < \infty$. Suppose \mathfrak{C} has a finite ε -covering number in a metric D . That is, there exists a range $(0, \tilde{\varepsilon}]$ for some $\tilde{\varepsilon} > 0$ such that $\log N(\mathfrak{C}, D, \varepsilon) < \infty$ for all $0 < \varepsilon \leq \tilde{\varepsilon}$. Suppose the map $K \mapsto \|\cdot\|_K$ is $C_{\mathfrak{C}}$ -Lipschitz over \mathfrak{C} in the maximum norm with respect to D , i.e., for all $K, L \in \mathfrak{C}$*

$$\| \|\cdot\|_K - \|\cdot\|_L \|_{\infty} := \max_{\|x\|_{\ell_2}=1} |\|x\|_K - \|x\|_L| \leq C_{\mathfrak{C}}D(K, L).$$

Then for any $\gamma \in (0, 1)$, we have that with probability $1 - \gamma$ over m samples drawn from P , the following bound holds for all $K \in \mathfrak{C}$,

$$F(K; P) \leq F(K; P_m) + 2 \inf_{\varepsilon \in (0, \tilde{\varepsilon}]} \left\{ \varepsilon + c_{P,m} \sqrt{\frac{2}{m} \log N\left(\mathfrak{C}, D, \frac{\varepsilon}{C_{\mathfrak{C}}}\right)} \right\} + c_{\mathfrak{C}} \sqrt{\frac{2 \ln 4/\gamma}{m}}.$$

Here $c_{P,m}$ is a constant that depends on \mathfrak{C} , P , and m and $c_{\mathfrak{C}} := 4/r_{\mathfrak{C}}$.

Proof. Let $X_m := \{x_1, \dots, x_m\}$ denote the m samples drawn from P and define the function class

$$\mathcal{F}_{\mathfrak{C}} := \{\|\cdot\|_K : \mathbb{R}^d \rightarrow \mathbb{R} : K \in \mathfrak{C}\}.$$

Note that our class of functions $\mathcal{F}_{\mathfrak{C}}$ is uniformly bounded. In particular, by our assumption on the class \mathfrak{C} , there exists an inner radius $0 < r_{\mathfrak{C}} < \infty$ such that $r_{\mathfrak{C}}B^d \subseteq K$ for all $K \in \mathfrak{C}$. Hence for any $\|x\|_{\ell_2} \leq 1$, we have $\|x\|_K \leq 1/r_{\mathfrak{C}}$. By standard statistical learning results [75, Theorem 26.5], the deviation between the empirical and population risk is controlled by the (empirical) Rademacher complexity $\mathcal{R}(\mathcal{F}_{\mathfrak{C}}; X_m)$ [9] of the function class $\mathcal{F}_{\mathfrak{C}}$. That is, it holds that with probability $1 - \gamma$, we have that for all $K \in \mathfrak{C}$,

$$\mathbb{E}_P[\|x\|_K] \leq \mathbb{E}_{P_m}[\|x\|_K] + 2\mathcal{R}(\mathcal{F}_{\mathfrak{C}}; X_m) + c_{\mathfrak{C}} \sqrt{\frac{2 \ln 4/\gamma}{m}}.$$

Here, $\mathcal{R}(\mathcal{F}; X_m)$ is defined by

$$\mathcal{R}(\mathcal{F}; X_m) := \mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \sigma_i f(x_i) \right]$$

where σ_i are Rademacher random variables for $i \in [m]$ and $c_{\mathfrak{C}} := 4/r_{\mathfrak{C}}$. Moreover, one can show [61, Proposition 5.2] that the Rademacher complexity is bounded by the covering number of the function class $\mathcal{F}_{\mathfrak{C}}$ with respect to the pseudo-norm $\|\cdot\|_{1, X_m}$:

$$\mathcal{R}(\mathcal{F}_{\mathfrak{C}}; X_m) \leq \inf_{\varepsilon \in (0, \tilde{\varepsilon}]} \left\{ \varepsilon + c_{P,m} \sqrt{\frac{2}{m} \log N(\mathcal{F}_{\mathfrak{C}}, \|\cdot\|_{1, X_m}, \varepsilon)} \right\}$$

where

$$c_{P,m} := \inf \left\{ c > 0 : \sup_{K \in \mathfrak{C}} \|\|\cdot\|_K\|_{1, X_m} \leq c \right\} \leq 1/r_{\mathfrak{C}}.$$

Here, $\|f\|_{1, X_m} := \frac{1}{m} \sum_{i=1}^m |f(x_i)|$.

Now, we can use the fact that $K \mapsto \|\cdot\|_K$ is Lipschitz over $K \in \mathfrak{C}$ to obtain a bound on $N(\mathcal{F}_{\mathfrak{C}}, \|\cdot\|_{1, X_m}, \varepsilon)$ based on $N(\mathfrak{C}, D, \eta)$ for some $\eta := \eta(\varepsilon)$ to be determined. First, we note that it suffices to obtain a bound on the covering number of $\mathcal{F}_{\mathfrak{C}}$ with respect to the maximum norm $\|\cdot\|_{\infty}$ as opposed to $\|\cdot\|_{1, X_m}$ since by our assumption on the data, $\|x_i\|_{\ell_2} \leq 1$ almost surely so that

$$\|\|\cdot\|_K\|_{1, X_m} = \frac{1}{m} \sum_{i=1}^m \|x_i\|_K \leq \frac{1}{m} \sum_{i=1}^m \|\|\cdot\|_K\|_{\infty} = \|\|\cdot\|_K\|_{\infty}.$$

Hence $N(\mathcal{F}_{\mathfrak{C}}, \|\cdot\|_{1, X_m}, \varepsilon) \leq N(\mathcal{F}_{\mathfrak{C}}, \|\cdot\|_{\infty}, \varepsilon)$.

Fix $\varepsilon \in (0, \bar{\varepsilon}]$ and set $\eta := \varepsilon/C_{\mathfrak{C}}$. Select an η -covering of \mathfrak{C} , denoted by \mathfrak{C}_{η} , of cardinality $\log |\mathfrak{C}_{\eta}| = \log N(\mathfrak{C}, D, \varepsilon/C_{\mathfrak{C}})$. Then define the finite set of functions $\mathcal{F}_{\mathfrak{C}_{\eta}} := \{\|\cdot\|_K : K \in \mathfrak{C}_{\eta}\}$. Then $\log |\mathfrak{C}_{\eta}| = \log |\mathcal{F}_{\mathfrak{C}_{\eta}}|$ and for any $K \in \mathfrak{C}$, there exists a $K_{\eta} \in \mathfrak{C}_{\eta}$ such that $D(K, K_{\eta}) \leq \eta$. By the assumption that $K \mapsto \|\cdot\|_K$ is $C_{\mathfrak{C}}$ -Lipschitz in the D metric, we have

$$\|\|\cdot\|_K - \|\cdot\|_{K_{\eta}}\|_{\infty} \leq C_{\mathfrak{C}}D(K, K_{\eta}) \leq C_{\mathfrak{C}}\eta = \varepsilon.$$

Hence, we have obtained an ε -covering of $\mathcal{F}_{\mathfrak{C}}$ with respect to $\|\cdot\|_{\infty}$ of cardinality $N(\mathfrak{C}, D, \varepsilon/C_{\mathfrak{C}})$. \square

The above result establishes a uniform generalization bound for any collection of star bodies \mathfrak{C} that satisfies two requirements: 1) the functional $K \mapsto \|\cdot\|_K$ is Lipschitz continuous over \mathfrak{C} with respect to a metric D and 2) \mathfrak{C} enjoys finite entropy bounds. In this work, we established that our objective is Lipschitz continuous over the space of well-conditioned star bodies $\mathcal{S}^d(r, 1)$ with respect to the radial metric δ . Moreover, by Corollary 3, this space $(\mathcal{S}^d(r, 1), \delta)$ enjoys finite covering numbers. While there are several prior works that have analyzed entropy bounds for collections of convex bodies with respect to the Hausdorff metric, we are unaware of prior work analyzing such entropy bounds for star bodies in either the radial or Hausdorff metric. Our general bound can be applied to $\mathcal{S}^d(r, 1)$ under the radial metric, but we do not have quantitative upper bounds on the covering number $\log N(\mathcal{S}^d(r, 1), \delta, \varepsilon)$. Nevertheless, we show that our generalization bound can be applied to several interesting classes of bodies, including star bodies given as unions of convex bodies along with other classes of convex bodies that satisfy the conditions of Theorem 13 with respect to the Hausdorff metric d_H . This will recover some known results in the literature on matrix factorization while generating novel statistical learning guarantees for several classes of bodies of interest, including unions of convex bodies, ellipsoids, polytopes, and linear images of Schatten p -norm balls.

We note that in our general bound, we assume that the data distribution is supported in the unit Euclidean ball. This assumption could be relaxed to allow for a larger radius, which would then be absorbed into the constants $c_{P,m}$ and $c_{\mathfrak{C}}$ that appear in our upper bound. It would be interesting to derive a result that would allow for a broader range of distributions. We focus on this class of almost surely bounded distributions for simplicity here.

In our results, we use \lesssim to denote inequality up to absolute constants. A result that we will repeatedly use in the sequel is the following approximate monotonicity property of covering numbers:

Proposition 6 (Exercise 4.2.10 in [91]). *For two subsets T_1, T_2 of a metric space X with metric D , $T_1 \subseteq T_2$ implies $N(T_1, D, \varepsilon) \leq N(T_2, D, \varepsilon/2)$.*

6.2 Application: Star bodies as Unions of Well-Conditioned Convex Bodies

We first consider a special class of star bodies that are given by unions of convex bodies. In particular, consider the following class of convex bodies with unit volume and inner width at least r : $\mathcal{C}^d(r, 1) := \{K \in \mathcal{C}^d : rB^d \subseteq K, \text{vol}_d(K) = 1\}$, which is a subset of the space of well-conditioned star bodies we considered in Section 4. As shown in Lemma 1, we also have $K \subseteq R_r B^d$ for all $K \in \mathcal{C}^d(r, 1)$, where $R_r := \frac{d+1}{r^{d-1}\kappa_{d-1}}$ and $\kappa_{d-1} := \text{vol}_{d-1}(B^{d-1})$. Now, for any $L \in \mathbb{N}$, define the following subset of star bodies $\mathcal{S}_L^d(r) \subset \mathcal{S}^d$ by

$$\mathcal{S}_L^d(r) := \left\{ K = \bigcup_{i=1}^L K_i : K_i \in \mathcal{C}^d(r, 1) \forall i \in [L] \right\}.$$

This class of star bodies contains sets that can be obtained as unions of convex bodies with sufficiently large inner width and unit volume. By exploiting our previous continuity results and using known metric entropy estimates for bounded convex bodies, we can apply Theorem 13 to obtain a generalization bound as follows:

Corollary 5. *Let P be a distribution over \mathbb{R}^d such that for $x \sim P$, $\|x\|_{\ell_2} \leq 1$ almost surely. Fix $0 < r < \infty$. Then for any $\gamma \in (0, 1)$, we have that with probability $1 - \gamma$ over m samples drawn from P , the following bound holds for all $K \in \mathcal{S}_L^d(r)$,*

$$F(K; P) \leq F(K; P_m) + 2 \inf_{\varepsilon > 0} \left\{ \varepsilon + c_{P,m} \sqrt{\frac{L}{m} \log \left(\frac{\alpha(\varepsilon, r, d)}{L} \right)} \right\} + c_r \sqrt{\frac{2 \ln 4/\gamma}{m}}$$

Optimal Regularization

where $\log \alpha(\varepsilon, r, d) \lesssim d \log (R_r^2/\varepsilon r^5)$ and $c_r \lesssim 1/r$.

Proof. First, recall that by Proposition 2, we established that $K \mapsto \|\cdot\|_K$ is $C_{\mathfrak{C}} := 1/r^2$ -Lipschitz continuous in the radial Hausdorff metric over $\mathcal{S}_L^d(r)$. Note that this part of the Proposition only requires $rB^d \subseteq K$, which is satisfied for $K \in \mathcal{S}_L^d(r)$. Now, we must establish that the metric entropy $\log N(\mathcal{S}_L^d(r), \delta, \varepsilon)$ is finite under the radial metric. To prove this, we will build an ε -cover of $\mathcal{S}_L^d(r)$ using known covers of bounded convex bodies in the Hausdorff metric and show how the radial distance between sets in $\mathcal{S}_L^d(r)$ can be upper bounded by their distance in the Hausdorff metric.

We record some useful known results. It has been shown [37] that the space $\mathfrak{C}(r) := \{K \in \mathcal{C}^d : K \subseteq R_r B^d\}$ has an ε^2 -covering number in the Hausdorff metric of size $\log N(\mathfrak{C}(r), d_H, \varepsilon) \leq c \cdot \frac{d-1}{2} \log \left(\frac{R_r}{\varepsilon}\right) < \infty$ for all $\varepsilon \leq R_r \varepsilon_0$ for some c and ε_0 depending on d . Call such a covering $\tilde{\mathcal{C}}_{\varepsilon^2}^d$ and define the relevant subset $\mathcal{C}_{\varepsilon^2}^d := \tilde{\mathcal{C}}_{\varepsilon^2}^d \cap \mathcal{C}^d(r, 1)$. Then for any $K \in \mathcal{C}^d(r, 1)$, there exists a $\tilde{K} \in \mathcal{C}_{\varepsilon^2}^d$ such that $d_H(K, \tilde{K}) \leq \varepsilon^2$. Using this covering, we now construct a cover of $\mathcal{S}_L^d(r)$. Define the following finite class of star bodies:

$$\mathcal{S}_{L, \varepsilon^2}^d := \left\{ K = \bigcup_{i=1}^L K_i : K_i \in \mathcal{C}_{\varepsilon^2}^d \forall i \in [L] \right\} \subset \mathcal{S}_L^d(r).$$

We claim that this is an $R_r \varepsilon^2 / r^3$ -covering of $\mathcal{S}_L^d(r)$ in the radial Hausdorff metric. Fix $K \in \mathcal{S}_L^d(r)$. Then $K = \bigcup_{i=1}^L K_i$ where $K_i \in \mathcal{C}^d(r, 1)$ for each $i \in [L]$. By the definition of $\mathcal{C}_{\varepsilon^2}^d$, for each K_i , there exists a corresponding $\tilde{K}_i \in \mathcal{C}_{\varepsilon^2}^d$ such that $d_H(K_i, \tilde{K}_i) \leq \varepsilon^2$. We claim that $\delta(K, \tilde{K}) \leq R_r \varepsilon^2 / r^3$ where $\tilde{K} := \bigcup_{i=1}^L \tilde{K}_i \in \mathcal{S}_{L, \varepsilon^2}^d$. First, note that each K_i, \tilde{K}_i satisfies $rB^d \subseteq K_i, \tilde{K}_i \subseteq R_r B^d$. For a star body K that is the union of star bodies K_1, \dots, K_L , its radial function satisfies $\rho_K(x) = \min_{i \in [L]} \rho_{K_i}(x)$. Moreover, we have the relation $\rho_K(x) = 1/\|x\|_K$. Thus, for any $\|x\|_{\ell_2} = 1$, we have that

$$\begin{aligned} |\rho_K(x) - \rho_{\tilde{K}}(x)| &= \left| \min_{i \in [L]} \rho_{K_i}(x) - \min_{i \in [L]} \rho_{\tilde{K}_i}(x) \right| = \left| \min_{i \in [L]} \frac{1}{\|x\|_{K_i}} - \min_{i \in [L]} \frac{1}{\|x\|_{\tilde{K}_i}} \right| \\ &= \left| \max_{i \in [L]} \|x\|_{K_i} - \max_{i \in [L]} \|x\|_{\tilde{K}_i} \right| \\ &\leq \max_{i \in [L]} \left| \|x\|_{K_i} - \|x\|_{\tilde{K}_i} \right| \\ &\leq \frac{R_r}{r^3} d_H(K_{i_*}, \tilde{K}_{i_*}) \\ &\leq \frac{R_r}{r^3} \cdot \varepsilon^2. \end{aligned}$$

where we set $i_* \in \arg \max_{i \in [L]} \left| \|x\|_{K_i} - \|x\|_{\tilde{K}_i} \right|$, used Proposition 2 in the second to last inequality, and the definition of $\mathcal{C}_{\varepsilon^2}^d$ in the final inequality. Since this holds for any $\|x\|_{\ell_2} = 1$, we can take the maximum over the sphere to obtain

$$\delta(K, \tilde{K}) = \max_{\|x\|_{\ell_2}=1} |\rho_K(x) - \rho_{\tilde{K}}(x)| \leq \frac{R_r}{r^3} \cdot \varepsilon^2.$$

Since $K \in \mathcal{S}_L^d(r)$ was arbitrary, this shows that $\mathcal{S}_{L, \varepsilon^2}^d$ is an $R_r \varepsilon^2 / r^3$ -covering of $\mathcal{S}_L^d(r)$ in the radial metric. Note that this covering is of cardinality $|\mathcal{S}_{L, \varepsilon^2}^d| = \binom{|\mathcal{C}_{\varepsilon^2}^d|}{L}$. Using the upper bound $\binom{n}{k} \leq (n \cdot e/k)^k$, we can estimate the metric entropy as

$$\log N(\mathcal{S}_L^d(r), \delta, \varepsilon) \leq \log \binom{|\mathcal{C}_{\varepsilon^2}^d|}{L} \lesssim L \log \left(\frac{|\mathcal{C}_{\varepsilon^2}^d|}{L} \right)$$

where $\log |\mathcal{C}_{\varepsilon^2}^d| \lesssim d \log(R_r/\varepsilon)$. □

Remark on the choice of parameters. This result focuses on the class of star bodies that can be expressed as unions of convex bodies with fixed volume and inner width at least $r > 0$, but a similar bound can also be obtained when the convex bodies are restricted to lie in RB^d for a user-specified outer radius

$R < \infty$. Note that the outer radius R_r grows with the dimension if r is kept constant. In high-dimensions, however, the volume of the unit ball concentrates on the boundary, so it is natural to instead choose a value of r that also increases with the dimension, e.g., setting $r = \alpha\sqrt{d}$ for some constant $0 < \alpha < 1$. This would ensure that R_r and r are on the same order.

Simplification for well-conditioned convex regularizers. Note that this above result also gives a generalization guarantee for well-conditioned convex regularizers in the case $L = 1$. One could improve this result in terms of powers of r for the convex case by directly using relevant results solely in the Hausdorff topology. Note that this cannot be done for the case of unions of convex bodies, since the radial topology is more natural over the space of star bodies.

6.3 Application: Ellipsoidal Regularizers

We additionally consider another class of convex bodies that are linear images of a norm ball. Specifically, we showcase a generalization bound for learning regularizers induced by ellipsoids.

Corollary 6 (Generalization Bound for Ellipsoids). *Let P be a distribution over \mathbb{R}^d such that for $x \sim P$, $\|x\|_{\ell_2} \leq 1$ almost surely. Fix $0 < r < R < \infty$. Then for any $\gamma \in (0, 1)$, we have that with probability $1 - \gamma$ over m samples drawn from P , the following bound holds for all $K \in \{A(B^d) : A \in \mathbb{R}^{d \times d}, A = A^T, RI \succeq A \succeq rI\}$:*

$$F(K; P) \leq F(K; P_m) + 2 \inf_{\varepsilon > 0} \left\{ \varepsilon + c_{P,m} \sqrt{\frac{d(d+1)}{2m} \log \left(\frac{3R^2}{\varepsilon r^3} \right)} \right\} + c_r \sqrt{\frac{2 \ln 4 / \gamma}{m}}.$$

Proof of Corollary 6. Let $\mathfrak{A} := \{A \in \mathbb{R}^{d \times d} : A = A^T, rI \preceq A \preceq RI\}$. Observe that for any $A \in \mathfrak{A}$ and $x \in A(B^d)$, we have $\|x\|_{\ell_2} \leq \|A\| \|z\|_{\ell_2} \leq R$ where $x = Az$ and $\|z\|_{\ell_2} \leq 1$, meaning $A(B^d) \subseteq RB^d$. Moreover, $\|Az\|_{\ell_2} \geq r\|z\|_{\ell_2}$ for any $z \in \mathbb{R}^d$ so $A(B^d) \supseteq rB^d$. Thus, $K \mapsto \|\cdot\|_K$ is Lipschitz over $\mathfrak{C} = \{A(B^d) : A \in \mathfrak{A}\}$ with respect to the Hausdorff metric.

To complete the proof, we show that the covering number of \mathfrak{C} in the Hausdorff metric can be bounded by the covering number of \mathfrak{A} with respect to the spectral norm $\|\cdot\|$. This follows from noting that for any $A, D \in \mathfrak{A}$, we have by definition of the Hausdorff metric that

$$\begin{aligned} d_H(A(B^d), D(B^d)) &= \max_{\|x\|_{\ell_2}=1} |h_{A(B^d)}(x) - h_{D(B^d)}(x)| \\ &= \max_{\|x\|_{\ell_2}=1} |h_{B^d}(A^T x) - h_{B^d}(D^T x)| \\ &= \max_{\|x\|_{\ell_2}=1} |\|A^T x\|_{\ell_2} - \|D^T x\|_{\ell_2}| \\ &\leq \max_{\|x\|_{\ell_2}=1} \|(A - D)^T x\|_{\ell_2} =: \|A - D\|. \end{aligned}$$

But note that the covering number of \mathfrak{A} with respect to the spectral norm $\|\cdot\|$ is bounded by

$$\log N(\mathfrak{A}, \|\cdot\|, \varepsilon) \leq \frac{d(d+1)}{2} \cdot \log \left(\frac{3R}{\varepsilon} \right)$$

since the set of PSD matrices is a Riemannian manifold of dimension $d(d+1)/2$ [36]. This gives the bound

$$\log N \left(\mathfrak{C}, d_H, \frac{\varepsilon r^3}{R} \right) \leq \log N \left(\mathfrak{A}, \|\cdot\|, \frac{\varepsilon r^3}{R} \right) \leq \frac{d(d+1)}{2} \log \left(\frac{3R^2}{\varepsilon r^3} \right).$$

□

6.4 Application: Polyhedral Regularizers and Dictionary Learning

We now apply our result to another important class of convex bodies: polytopes. This bound also characterizes a generalization bound for dictionary learning. Let $\mathcal{P}_T \subset \mathcal{C}^d$ denote the space of polytopes with T vertices:

$$\mathcal{P}_T := \{\text{conv}(S) : S \subset \mathbb{R}^d \text{ has cardinality at most } T\}.$$

For $0 < r < R < \infty$, let $\mathcal{P}_T(r, R) := \{K \in \mathcal{P}_T : rB^d \subseteq K \subseteq RB^d\}$. Note that by Lemma 1, $\{K \in \mathcal{P}_T : rB^d \subseteq K, \text{vol}_d(K) = 1\} \subseteq \mathcal{P}_T(r, R_r)$, so the following bound can be applied to this class of volume-normalized polytopes as well.

Corollary 7. *Let P be a distribution over \mathbb{R}^d such that for $x \sim P$, $\|x\|_{\ell_2} \leq 1$ almost surely. Fix $0 < r < R < \infty$. Then for any $\gamma \in (0, 1)$, we have that with probability $1 - \gamma$ over m samples drawn from P , the following bound holds for all $K \in \mathcal{P}_T(r, R)$,*

$$F(K; P) \leq F(K; P_m) + 2 \inf_{\varepsilon > 0} \left\{ \varepsilon + c_{P,m} \sqrt{\frac{Td}{m} \log \left(\frac{MR^2}{\varepsilon r^3} \right)} \right\} + c_r \sqrt{\frac{2 \ln 4/\gamma}{m}}.$$

Here, M is a positive absolute constant.

Proof. First, note that due to our assumptions on $\mathcal{P}_T(r, R)$, we have again by Proposition 2 that $K \mapsto \|x\|_K$ is R/r^3 -Lipschitz with respect to the Hausdorff metric over this class of convex bodies. It has additionally been shown [37] (proof of Lemma C.1) that $\mathcal{P}_T(R) := \{K \in \mathcal{P}_T : K \subseteq RB^d\}$ enjoys a finite ε -covering number with respect to the Hausdorff metric, which scales as

$$\log N(\mathcal{P}_T(R), d_H, \varepsilon) \lesssim Td \cdot \log \left(\frac{MR}{\varepsilon} \right)$$

for some positive absolute constant M . Using the monotonicity property (Proposition 6) to bound the $\varepsilon r^3/R$ -covering number of $\mathcal{P}_T(r, R)$ by that of $\mathcal{P}_T(R)$ achieves the desired result. \square

The class of bounded polytopes of sufficiently large inner width covers the setting of dictionary learning under some conditions on the dictionaries of interest. In particular, recall that in dictionary learning, the goal is to learn a basis $A \in \mathbb{R}^{d \times p}$ such that a given dataset $\{x_i\}_{i=1}^N$ can be sparsely represented as $x_i \approx Az_i$ where z_i is sparse. In most cases, sparsity is induced by solving for latent codes z_i with small ℓ_1 -norm. In this case, the learned regularizer can be shown to be a gauge function induced by a linear image of the ℓ_1 -ball:

$$\|x\|_{A(\mathcal{B}_{\ell_1})} = \inf\{\|z\|_{\ell_1} : z = Ax\}.$$

This follows by definition, since $\|x\|_{A(\mathcal{B}_{\ell_1})} = \inf\{t > 0 : x \in t \cdot A(\mathcal{B}_{\ell_1})\} = \inf\{t > 0 : x = Az, z \in t \cdot \mathcal{B}_{\ell_1}\} = \inf\{\|z\|_{\ell_1} : x = Az\}$. Thus, the regularizers of interest in dictionary learning are gauges induced by linear images of the ℓ_1 -ball, where the linear map A is the object to be learned. That is, for $A \in \mathbb{R}^{d \times p}$, the convex body of interest is given by

$$A(\mathcal{B}_{\ell_1}) := \{Az \in \mathbb{R}^d : \|z\|_{\ell_1} \leq 1\} = \text{conv}(\{A_1, \dots, A_p\}).$$

To apply our result for polytopes, we need to show that a class of suitable dictionaries $\mathfrak{A} \subset \mathbb{R}^{d \times p}$ induces a collection of convex bodies $\{A(\mathcal{B}_{\ell_1}) : A \in \mathfrak{A}\}$ that satisfies the geometric conditions $rB^d \subseteq A(\mathcal{B}_{\ell_1}) \subseteq RB^d$ for all $A \in \mathfrak{A}$. For a parameter $\eta > 0$, consider the collection of dictionaries

$$\mathfrak{A} := \{A \in \mathbb{R}^{d \times p} : \|A_i\|_{\ell_2} = 1, \text{rank}(A) = d, \sqrt{\sigma_{\min}(AA^T)} \geq \eta > 0\}$$

and fix $A \in \mathfrak{A}$. We first show $A(\mathcal{B}_{\ell_1}) \subseteq RB^d$ for some $R > 0$. Consider $x \in A(\mathcal{B}_{\ell_1})$. Then $x = Az$ for $\|z\|_{\ell_1} \leq 1$. Since $\|A\|^2 \leq \|A\|_{\ell_2}^2 = p$ where $\|\cdot\|$ and $\|\cdot\|_{\ell_2}$ denote the spectral and Frobenius norm, respectively, we see that

$$\|x\|_{\ell_2} = \|Az\|_{\ell_2} \leq \|A\| \|z\|_{\ell_2} \leq \sqrt{p} \|z\|_{\ell_1} \leq \sqrt{p}.$$

Thus $A(\mathcal{B}_{\ell_1}) \subseteq \sqrt{p}B^d$. We now show $A(\mathcal{B}_{\ell_1}) \supseteq rB^d$ for some $0 < r < R$. Note that this inclusion holds if and only if the support functions satisfy $h_{A(\mathcal{B}_{\ell_1})} \geq h_{rB^d}$. Note that for any $\|u\|_{\ell_2} = 1$, we have that the support function of $A(\mathcal{B}_{\ell_1})$ is given by

$$h_{A(\mathcal{B}_{\ell_1})}(u) = h_{\mathcal{B}_{\ell_1}}(A^T u) = \|A^T u\|_{\ell_\infty}.$$

But since $A \in \mathfrak{A}$, note that

$$\min_{\|u\|_{\ell_2}=1} \|A^T u\|_{\ell_\infty} \geq \frac{1}{\sqrt{p}} \min_{\|u\|_{\ell_2}=1} \|A^T u\|_{\ell_2} = \frac{1}{\sqrt{p}} \cdot \sqrt{\sigma_{\min}(AA^T)} \geq \frac{\eta}{\sqrt{p}}.$$

This guarantees $\eta/\sqrt{p}B^d \subseteq A(\mathcal{B}_{\ell_1})$.

Thus each $K \in \{A(\mathcal{B}_{\ell_1}) : A \in \mathfrak{A}\}$ satisfies $\eta/\sqrt{p}B^d \subseteq K \subseteq \sqrt{p}B^d$. Hence the functional $\|x\|_K$ is Lipschitz over the collection $\mathfrak{C} := \{A(\mathcal{B}_{\ell_1}) : A \in \mathfrak{A}\}$. Moreover, the set

$$\{A(\mathcal{B}_{\ell_1}) : A \in \mathfrak{A}\} \subseteq \mathcal{P}_p(\sqrt{p})$$

since each $A(\mathcal{B}_{\ell_1})$ is a polytope with at most p extreme points. Thus we achieve the same bound in Corollary 7 with $T = p$, $r = \eta/\sqrt{p}$, and $R = \sqrt{p}$. These bounds match previously achieved generalization bounds for dictionary learning [88, 36] up to logarithmic factors.

6.5 Application: Semidefinite Regularizers

Our final class of convex bodies consists of linear images of Schatten p -norm balls. Such classes have found wide interest in data science. For example, if one wants to learn an infinite collection of dictionary elements, a natural generalization of dictionary learning would require learning a linear image of the nuclear norm ball [76]. The set of dictionary elements in this case corresponds to the convex hull of the set of rank-1 matrices.

For a matrix $A \in \mathbb{R}^{q \times q}$, let $\sigma_i(A)$ denote the i -th singular value of A . For $p \in [1, \infty]$, we define the Schatten p -norm of A as

$$\|A\|_{S_p} := \begin{cases} (\sum_{i=1}^q \sigma_i(A)^p)^{1/p} & \text{if } p < \infty, \\ \sigma_{\max}(A) & \text{else.} \end{cases}$$

Let $\mathcal{B}_{S_p} := \{A \in \mathbb{R}^{q \times q} : \|A\|_{S_p} \leq 1\}$ denote the Schatten p -norm ball and define the space of linear maps from $\mathbb{R}^{q \times q}$ to \mathbb{R}^d as $L(\mathbb{R}^{q \times q}, \mathbb{R}^d)$. For a linear map $\mathcal{L} \in L(\mathbb{R}^{q \times q}, \mathbb{R}^d)$, define the induced operator norm $\|\cdot\|_{S_p, 2}$ as

$$\|\mathcal{L}\|_{S_p, 2} := \max_{Z \in \mathcal{B}_{S_p}} \|\mathcal{L}(Z)\|_{\ell_2}.$$

For a collection of linear maps $\mathfrak{L}_p \subset L(\mathbb{R}^{q \times q}, \mathbb{R}^d)$, let \mathfrak{C}_p denote the induced collection of convex bodies which are linear images of Schatten p -norm balls

$$\mathfrak{C}_p := \{\mathcal{L}(\mathcal{B}_{S_p}) : \mathcal{L} \in \mathfrak{L}_p \subset L(\mathbb{R}^{q \times q}, \mathbb{R}^d)\}.$$

We will prove the following result, which prescribes a set of conditions on the collection \mathfrak{L}_p to obtain a generalization bound.

Corollary 8. *Let P be a distribution over \mathbb{R}^d such that for $x \sim P$, $\|x\|_{\ell_2} \leq 1$ almost surely. For $p \geq 1$, let \tilde{p} denote the conjugate² of p and consider the set of bounded linear maps*

$$\mathfrak{L}_p(r, R) := \left\{ \mathcal{L} \in L(\mathbb{R}^{q \times q}, \mathbb{R}^d) : r \leq \min_{\|x\|_{\ell_2}=1} \|\mathcal{L}^T x\|_{S_{\tilde{p}}} \leq \|\mathcal{L}\|_{S_p, 2} \leq R \right\}.$$

Then for any $\gamma \in (0, 1)$, we have that with probability $1 - \gamma$ over m samples drawn from P , the following bound holds for all $K \in \{\mathcal{L}(\mathcal{B}_{S_p}) : \mathcal{L} \in \mathfrak{L}_p(r, R)\}$,

$$F(K; P) \leq F(K; P_m) + 2 \inf_{\varepsilon > 0} \left\{ \varepsilon + c_{P, m} \sqrt{\frac{q^2 d}{2m}} \cdot \log \left(\frac{R^2}{\varepsilon r^3} \right) \right\} + c_r \sqrt{\frac{2 \ln 4 / \gamma}{m}}.$$

²Recall that the conjugate \tilde{p} of p is $\tilde{p} = \infty$ if $p = 1$ and $1/p + 1/\tilde{p} = 1$ otherwise.

Proof. We first show that the collection $\mathfrak{C}_p := \{\mathcal{L}(\mathcal{B}_{S_p}) : \mathcal{L} \in \mathfrak{L}_p(r, R)\}$ is bounded and has a sufficiently large inner width. Clearly for $K \in \mathfrak{C}_p$, we have $K \subseteq RB^d$. Moreover, note that the support function $h_{\mathcal{L}(\mathcal{B}_{S_p})}(u) = \|\mathcal{L}^T u\|_{S_{\bar{p}}}$. Thus, for $K \in \mathfrak{C}_p$, the condition $rB^d \subseteq K$ is equivalent to

$$h_{rB^d}(u) \leq h_{\mathcal{L}(\mathcal{B}_{S_p})}(u) \quad \forall \|u\|_{\ell_2} = 1 \iff r \leq \min_{\|u\|_{\ell_2}=1} \|\mathcal{L}^T u\|_{S_{\bar{p}}}.$$

By assumption, this inequality holds.

Now we need to bound the covering number of \mathfrak{C}_p in the Hausdorff metric. We claim such a bound can be obtained by covering the space \mathfrak{L}_p in the $\|\cdot\|_{S_p, 2}$ norm. Arguing similarly to proof of Corollary 6, we can show that for $\mathcal{L}, \mathcal{A} \in \mathfrak{L}_p$, we have

$$d_H(\mathcal{L}(\mathcal{B}_{S_p}), \mathcal{A}(\mathcal{B}_{S_p})) \leq \max_{\|x\|_{\ell_2}=1} \|(\mathcal{L}^T - \mathcal{A}^T)x\|_{S_{\bar{p}}} =: \|\mathcal{L}^T - \mathcal{A}^T\|_{2, S_{\bar{p}}}.$$

But by duality, note that $\|X\|_{S_p} = \max\{|\langle X, Z \rangle| : \|Z\|_{S_{\bar{p}}} \leq 1\}$ so that for any linear map $\mathcal{L} \in L(\mathbb{R}^{q \times q}, \mathbb{R}^d)$,

$$\begin{aligned} \|\mathcal{L}^T\|_{2, S_{\bar{p}}} &= \max_{\|x\|_{\ell_2}=1} \|\mathcal{L}^T x\|_{S_{\bar{p}}} = \max_{\|x\|_{\ell_2}=1} \max_{\|Z\|_{S_{\bar{p}}} \leq 1} |\langle Z, \mathcal{L}^T x \rangle| = \max_{\|x\|_{\ell_2}=1} \max_{\|Z\|_{S_{\bar{p}}} \leq 1} |\langle \mathcal{L}(Z), x \rangle| \\ &= \max_{\|Z\|_{S_{\bar{p}}} \leq 1} \max_{\|x\|_{\ell_2}=1} |\langle \mathcal{L}(Z), x \rangle| \\ &= \max_{\|Z\|_{S_{\bar{p}}} \leq 1} \|\mathcal{L}(Z)\|_{\ell_2} \\ &= \|\mathcal{L}\|_{S_p, 2} \end{aligned}$$

which gives $d_H(\mathcal{L}(\mathcal{B}_{S_p}), \mathcal{A}(\mathcal{B}_{S_p})) \leq \|\mathcal{L} - \mathcal{A}\|_{S_p, 2}$. It suffices to then bound the covering number of $\mathfrak{L}_p(R) := \{\mathcal{L} \in L(\mathbb{R}^{q \times q}, \mathbb{R}^d) : \|\mathcal{L}\|_{S_p, 2} \leq R\}$. But note that $\mathfrak{L}_p(R)$ is isomorphic to the space of $d \times q^2$ matrices of bounded $\|\cdot\|_{S_p, 2}$ norm. This gives the estimate

$$\log N(\mathfrak{C}_p, d_H, \varepsilon) \leq \log N(\mathfrak{L}_p, \|\cdot\|_{S_p, 2}, \varepsilon) \leq \log N(\mathfrak{L}_p(R), \|\cdot\|_{S_p, 2}, \varepsilon/2) \lesssim q^2 d \log \left(1 + \frac{2R}{\varepsilon}\right).$$

Bounding the $\varepsilon r^3/R$ -covering number of \mathfrak{C}_p in the Hausdorff metric by that of $\mathfrak{L}_p(R)$ in the metric induced by the $\|\cdot\|_{S_p, 2}$ -norm achieves the claimed bound. \square

7 Consequences Beyond Optimization

In this section, we briefly illustrate the consequences of our approach to regularizer selection for solution methods of inverse problems that do not rely on optimization but instead on the closely related topic of sampling. In the preceding sections of this paper, we advocate for summarizing a data distribution P via a star body regularizer $\|\cdot\|_K$. As discussed in Section 1.1, a regularizer $\|\cdot\|_K$ defined by the gauge of a star body $K \subseteq \mathbb{R}^d$ is optimal for a data distribution P if the Gibbs density with energy $\|\cdot\|_K$ is the projection of P onto the set of regularizer-induced Gibbs densities $\bar{\mathcal{D}}$. With this notion of optimality, our results in Section 3.3 on conditions for a data distribution under which the optimal regularizer is convex have computational implications beyond optimization.

Concretely, in the Bayesian paradigm for solving inverse problems, one aims to reconstruct an object from corrupted measurements by sampling from a suitable posterior distribution. When the prior is log-concave (i.e., when it is proportional to $\exp(-V)$ for a convex function V), the induced posterior distribution is also log-concave for many common likelihood models for the observed data (e.g., in linear inverse problems with likelihood functions given by exponential family models). In such situations, we are faced with the task of sampling from a log-concave posterior. The area of log-concave sampling is rich with algorithms and results guaranteeing convergence, for a wide range of smooth and non-smooth convex energies V [89, 64, 65, 45, 44, 15, 25, 26]. Consequently, the identification of log-concave Gibbs densities that can approximate a given prior distribution P by obtaining convex regularizers corresponding to P can yield a computationally tractable sampling-based approach for solving inverse problems.

We now describe this Bayesian perspective in more detail. Suppose we collect measurements $y \in \mathbb{R}^m$ of an underlying signal $x_* \in \mathbb{R}^d$, which are modeled via likelihood function of y given x of the form $e^{-\Phi(x; y)}$

for a known map Φ . (For example, if y and x_* are related via $y = f(x_*) + \eta$ for some known forward model $f : \mathbb{R}^d \rightarrow \mathbb{R}^m$ and $\eta \sim \mathcal{N}(0, \sigma^2 I_m)$, then $\Phi(x; y) := (2\sigma^2)^{-1} \|y - f(x)\|_{\ell_2}^2$.) From the Bayesian perspective [79], one views x_* as being drawn from a prior distribution P and the quantity of interest for estimating x_* is the posterior distribution of x_* given y , which we denote by μ_P :

$$\mu_P(dx) := \frac{1}{Z_P(y)} e^{-\Phi(x;y)} dP(x)$$

where $Z_P(y) := \int_{\mathbb{R}^d} e^{-\Phi(x;y)} dP(x)$ is the normalizing constant.

If instead of using the posterior μ_P we use the posterior $\mu_{p_K^*}$ induced by the moment projection p_K^* of P onto $\overline{\mathcal{D}}$, it is natural to ask what the implications might be in terms of reconstruction. Using recent results in the Bayesian inference literature, one can show that the distance from the posterior μ_P induced by the true prior P to the induced posterior distribution $\mu_{p_K^*}$ is upper bounded by the distance between P and the moment projection p_K^* . To state the result, recall that the total variation distance between two measures P_1, P_2 is given by $d_{TV}(P_1, P_2) := \sup_A |P_1(A) - P_2(A)|$ where the supremum is over all measurable subsets.

Proposition 7. *Fix $y \in \mathbb{R}^m$. Consider two prior measures P_1, P_2 on \mathbb{R}^d with $P_1 \ll P_2$ and log-likelihood function $\Phi(\cdot; y) : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $Z_{P_i}(y) := \int_{\mathbb{R}^d} e^{-\Phi(x;y)} dP_i(x) < \infty$ for each $i = 1, 2$. Defining the posterior measures $\mu_{P_i}(dx) := Z_{P_i}(y)^{-1} e^{-\Phi(x;y)} dP_i(x)$, we have that the following bounds hold:*

$$d_{TV}^2(\mu_{P_1}, \mu_{P_2}) \leq \frac{2}{Z_{P_1}^2(y)} D_{KL}(P_1 || P_2).$$

Proof. This result follows from Theorem 8 in [78] along with an application of Pinsker's inequality [59], which states that $d_{TV}(P_1, P_2) \leq \sqrt{D_{KL}(P_1 || P_2)}/2$ for any two measures P_1, P_2 . \square

Hence, if the normalizing constants $Z_P(y), Z_{p_K^*}(y)$ are finite, the m -projection p_K^* of P onto $\overline{\mathcal{D}}$ satisfies:

$$\min_{p_K^* \in \overline{\mathcal{D}}} d_{TV}^2(\mu_P, \mu_{p_K^*}) \leq \frac{2}{Z_P(y)^2} D_{KL}(P || p_K^*).$$

Thus, identifying an element p_K^* of the class $\overline{\mathcal{D}}$ to best approximate a prior P with a view to minimizing the total variation between the corresponding posteriors $\mu_P, \mu_{p_K^*}$ may be achieved by setting p_K^* equal to the m -projection of P onto $\overline{\mathcal{D}}$; in turn, this m -projection can be obtained by deriving the optimal regularizer for P . If this optimal regularizer is convex so that p_K^* is log-concave, then under appropriate conditions on the likelihood $e^{-\Phi(x;y)}$ the posterior $\mu_{p_K^*}$ is log-concave. To summarize this line of reasoning, identifying an optimal regularizer for a data distribution P and conditions under which it is convex yield methods for obtaining samples from a log-concave posterior whose distance to the true posterior is controlled in terms of the quality of the approximation of P .

8 Conclusion and Discussion

In this paper, we analyzed the question of determining the optimal regularizer for a given data distribution. We showed that when considering the class of continuous, positively homogeneous regularizers, this question is equivalent to solving a variational optimization problem over the space of star bodies. We showed that given a distribution of interest, there exists a unique star body of fixed volume that achieves the optimal risk by exploiting dual Brunn-Minkowski theory. The density of the distribution helps characterize the radial function of the optimal star body, which can have a wide array of geometries depending on the distribution under consideration. We further analyzed the convergence of empirical risk minimizers in the limit of infinite data and robustness of solutions to noisy perturbations. We then obtained generalization guarantees for various classes of star bodies and convex bodies, recovering known results while developing novel guarantees. There are many interesting directions for future work, and we describe a few here:

Performance of regularizers in subsequent tasks. An important next step in the theory of optimal regularizers is to understand their performance in downstream tasks such as inverse problems. These types of questions are interesting from a generalization perspective, and would bring unique mathematical challenges. For example, previous works [55, 56] have established worst-case bounds on denoising performance using a convex regularizer, which depends on the statistical complexity of the regularizer’s subdifferential at the point of interest. In our setting and the convex case, this would amount to analyzing the subdifferential of $\|\cdot\|_{K_*}$, the gauge induced by the optimal star body K_* . Additionally, this would also likely involve quantifying the Gaussian width or statistical dimension [20, 19, 5] of the tangent cone of convex bodies at datapoints of interest, a quantity that has been well-studied in the context of inverse problems and statistical estimation tasks. Understanding what the analogous quantities would be in the general nonconvex star case could lead to new geometric questions regarding the boundary structure of such bodies.

Approximation power of convex bodies. The focus of our paper mainly lied in characterizing the optimal star body regularizer for a given distribution, but did not investigate questions related to modeling mismatch. It is of interest to investigate situations such as those discussed in Section 4, where the collection of sets we are searching over \mathfrak{C} does not contain the optimal regularizer K_* . For example, if the optimal regularizer K_* for a given distribution is nonconvex, what is the best convex approximation to this regularizer? Along these lines, are there benefits to considering the best convex approximation to the optimal star body regularizer from a computational or analytic perspective? The answers to such questions will further advance our understanding of the powers and limitations of convexity.

Computing the optimal regularizer. One interesting direction that stems from our work is to provide a practical method for learning optimal regularizers that are specified as the gauge of a star body for a given data set. The dictionary learning problem provides a roadmap for describing structured families of convex regularizers that are extremely expressive while being amenable to training (via alternating minimization methods) – these are regularizers specified as the gauge of some polytope parametrized as some linear image of the ℓ_1 -ball. A natural family of *nonconvex* regularizers that extends the ideas of dictionary learning are those specified as the gauge of unions of a *few* convex sets, each of which could be, for instance, specified as some projection of the ℓ_1 -ball. It would be interesting to develop numerical algorithms for learning regularizers with such structure.

Incorporating computational considerations. An outstanding challenge in the field of inverse problems is to better understand how contemporary data-driven methods for learning regularizers perform, and explain why they seem to work well in practice. An interesting direction based on our work is to view stylized instances of these practical methods as instances of our framework, but with the additional constraint that the star body arises from some computational model. For instance, we may view the dictionary learning problem as an instance of our framework with the constraint that the star body is specified as the linear image of an ℓ_1 -ball. It would be interesting to investigate the relationship between the optimal regularizers corresponding to the constrained (to being specified by some computational model) and unconstrained instances, and in the process, provide insight into the performance of practical methods for learning regularizers.

Acknowledgements

We thank Shuhan Yang for notifying us of a sign error in a previous version of this manuscript. VC was supported in part by AFOSR grants FA9550-23-1-0204, FA9550-23-1-0070 and NSF grant DMS 2113724. YS acknowledges support from the Ministry of Education (Singapore) Academic Research Fund (Tier 1) R-146-000-329-133.

References

- [1] Alekh Agarwal, Animashree Anandkumar, Prateek Jain, and Praneeth Netrapalli. Learning sparsely used overcomplete dictionaries via alternating minimization. *SIAM Journal on Optimization*, 26(4):2775–2799, 2016.
- [2] Alekh Agarwal, Animashree Anandkumar, and Praneeth Netrapalli. A clustering approach to learn sparsely-used overcomplete dictionaries. *IEEE Transactions on Information Theory*, 63(1):575–592, 2017.
- [3] Michal Aharon, Michael Elad, and Alfred Bruckstein. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006.
- [4] S. Amari and H. Nagaoka. *Methods of Information Geometry*. Translations of mathematical monographs. American Mathematical Society, 2000.
- [5] Dennis Amelunxen, Martin Lotz, Michael B. McCoy, and Joel A. Tropp. Living on the edge: Phase transitions in convex programs with random data. *Information and Inference: A Journal of the IMA*, 224–294:3(3), 2014.
- [6] Sanjeev Arora, Rong Ge, Tengyu Ma, and Ankur Moitra. Simple, efficient, and neural algorithms for sparse coding. *Conference on Learning Theory*, 2015.
- [7] Simon Arridge, Peter Maass, Ozan Öktem, and Carola-Bibiane Schönlieb. Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174, 2019.
- [8] Boaz Barak, Jonathan A. Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*, page 143–151, 2015.
- [9] Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. *Annals of Statistics*, 33:1497–1537, 2005.
- [10] Martin Benning and Martin Burger. Modern regularization methods for inverse problems. *Acta Numerica*, 27:1–111, 2018.
- [11] Badri Bhaskar and Benjamin Recht. Atomic norm denoising with applications to line spectral estimation. *IEEE Transactions on Signal Processing*, 61(23):5987–5999, 2013.
- [12] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros Dimakis. Compressed sensing using generative models. *International Conference on Machine Learning*, 2017.
- [13] Andrea Braides. A handbook of Γ -convergence. *Handbook of Differential Equations: Stationary Partial Differential Equations*, Volume 3, 2007.
- [14] Kristian Bredies, Karl Kunisch, and Thomas Pock. Total generalized variation. *SIAM Journal on Imaging Sciences*, 3:492 – 526, 2011.
- [15] Sébastien Bubeck, Ronen Eldan, and Joseph Lehec. Sampling from a log-concave distribution with projected langevin monte carlo. *Discrete and Computational Geometry*, 59:757–783, 2018.
- [16] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.
- [17] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [18] Antonin Chambolle and Pierre-Louis Lions. Image recovery via total variation minimization and related problems. *Numer. Math.*, 76:167–188, 1997.

- [19] Venkat Chandrasekaran and Michael I. Jordan. Computational and statistical tradeoffs via convex relaxation. *Proceedings of the National Academy of Sciences*, 110(13):E1181–E1190, 2013.
- [20] Venkat Chandrasekaran, Benjamin Recht, Pablo A. Parrilo, and Alan S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12:805–849, 2012.
- [21] Niladri S. Chatterji and Peter L. Bartlett. Alternating minimization for dictionary learning: Local convergence guarantees. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [22] Masoumeh Dashti and Andrew M. Stuart. Uncertainty quantification and weak approximation of an elliptic inverse problem. *SIAM Journal on Numerical Analysis*, 49(6):2524–2542, 2011.
- [23] Ingrid Daubechies, Michel Defrise, and Christine de Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.
- [24] David Donoho. For most large underdetermined systems of linear equations the minimal l1-norm solution is also the sparsest solution. *Communications on Pure and Applied Mathematics*, 59(6):797–829, 2006.
- [25] Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of langevin monte carlo via convex optimization. *Journal of Machine Learning Research*, 20:1–46, 2019.
- [26] Raaz Dwivedi, Yuansi Chen, Martin J. Wainwright, and Bin Yu. Log-concave sampling: Metropolis-hastings algorithms are fast. *Journal of Machine Learning Research*, 20:42, 2019.
- [27] P. P. B. Eggermont. Maximum entropy regularization of fredholm integral equations of the first kind. *SIAM Journal on Mathematical Analysis*, 24(6):1557–1576, 1993.
- [28] Michael Elad. Sparse and redundant representations: From theory to applications in signal and image processing. *Springer*, 2010.
- [29] Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96:1348 – 1360, 2001.
- [30] Maryam Fazel. Matrix rank minimization with applications. *Ph.D. Thesis, Department of Electrical Engineering, Stanford University*, 2002.
- [31] Simon Foucart and Ming-Jun Lai. Sparsest solutions of underdetermined linear systems via ℓ_q - minimization for $0 < q \leq 1$. *Applied and Computational Harmonic Analysis*, 26:395–407, 2009.
- [32] Cristina Garcia-Cardona and Brendt Wohlberg. Convolutional dictionary learning: A comparative review and new algorithms. *IEEE Transactions on Computational Imaging*, 4(3):366–381, 2018.
- [33] Richard J. Gardner. Geometric tomography. *Cambridge: Cambridge University Press*, 2006.
- [34] Andrew Gelman and Xiao-Li Meng. Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163–185, 1998.
- [35] Matteo Giordano and Richard Nickl. Consistency of bayesian inference with gaussian process priors in an elliptic inverse problem. *Inverse Problems*, 36(8):085001, 2020.
- [36] Rémi Gribonval, Rodolphe Jenatton, Francis Bach, Martin Kleinstueber, and Matthias Seibert. Sample complexity of dictionary learning and other matrix factorizations. *IEEE Transactions on Information Theory*, 61(6):3469–3486, 2015.
- [37] Adityanand Guntuboyina. Optimal rates of convergence for convex set estimation from support functions. *Annals of Statistics*, 40(1):385–411, 2012.
- [38] Andreas Habring and Martin Holler. A generative variational model for inverse problems in imaging. *SIAM Journal on Mathematics of Data Science*, 4(1):306–335, 2022.

- [39] G. Hansen, I. Herbut, H. Martini, and M. Moszynska. Starshaped sets. *Aequationes Mathematicae*, 94:1001–1092, 2020.
- [40] Trevor Hastie, Robert Tibshirani, and Martin Wainwright. Statistical learning with sparsity: The lasso and generalizations. *Chapman & Hall/CRC*, 2015.
- [41] Tanehisa Hirose. On the convergence theorem for star-shaped sets in E^n . *Proc. Japan Acad.*, 41(3):209–211, 1965.
- [42] B. T. Knapik, A. W. van der Vaart, and J. H. van Zanten. Bayesian inverse problems with gaussian priors. *The Annals of Statistics*, 39(5):2626–2657, 2011.
- [43] Erich Kobler, Alexander Effland, Karl Kunisch, and Thomas Pock. Total deep variation: A stable regularizer for inverse problems. *arXiv preprint arXiv:2006.08789*, 2020.
- [44] László Lovász and Santosh Vempala. Hit-and-run from a corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006.
- [45] László Lovász and Santosh Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
- [46] Sebastian Lunz, Ozan Öktem, and Carola-Bibiane Schönlieb. Adversarial regularizers in inverse problems. *Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018.
- [47] Erwin Lutwak. Dual mixed volumes. *Pacific Journal of Mathematics*, 58(2):531–538, 1975.
- [48] Erwin Lutwak. Centroid bodies and dual mixed volumes. *Proceedings of the London Mathematical Society*, s3-60(2):365–391, 1990.
- [49] Julien Mairal, Francis Bach, and Jean Ponce. Sparse modeling for image and vision processing. *Foundations and Trends in Computer Graphics and Vision*, 8(2–3):85–283, 2014.
- [50] Hosein Mohimani, Massoud Babaie-Zadeh, and Christian Jutten. A fast approach for overcomplete sparse decomposition based on smoothed ℓ_0 norm. *IEEE Transactions on Signal Processing*, 57(1):289–301, 2008.
- [51] François Monard, Richard Nickl, and Gabriel P. Paternain. Consistent inversion of noisy non-abelian x-ray transforms. *Communications on Pure and Applied Mathematics*, 74(5):1045–1099, 2021.
- [52] Harry Nyquist. Certain topics in telegraph transmission theory. *Trans. AIEE.*, 47(2):617–644, 1928.
- [53] Bruno A. Olshausen and David J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.
- [54] Bruno A. Olshausen and David J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision in Research*, 37(23):3311–3325, 1997.
- [55] Samet Oymak and Babak Hassibi. Sharp mse bounds for proximal denoising. *Foundations of Computational Mathematics*, 16:965–1029, 2016.
- [56] Samet Oymak, Christos Thrampoulidis, and Babak Hassibi. The squared-error of generalized lasso: A precise analysis. *51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2013.
- [57] Vardan Pappyan, Yaniv Romano, Jeremias Sulam, and Michael Elad. Convolutional dictionary learning via local processing. *International Conference on Computer Vision (ICCV)*, page 5296–5304, 2017.
- [58] Konstantin Pieper and Armenak Petrosyan. Nonconvex regularization for sparse neural networks. *Applied and Computational Harmonic Analysis*, 61:25 – 56, 2022.

- [59] Mark S. Pinsker. Information and information stability of random variables and processes. *Holden-Day, Inc., San Francisco, Calif.-London, Amsterdam*, 1964.
- [60] David Pollard. Convergence of stochastic processes. *Springer-Verlag*, 1984.
- [61] Patrick Rebeschini. Lecture 5: Covering numbers bounds for rademacher complexity. chaining. 2021.
- [62] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- [63] Edward T. Reehorst and Philip Schniter. Regularization by denoising: clarifications and new interpretations. *IEEE Transactions on Computational Imaging*, 5(1):52 – 67, 2019.
- [64] Gareth O. Roberts and Richard L. Tweedie. Exponential convergence of langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363, 1996.
- [65] Gareth O. Roberts and Richard L. Tweedie. Geometric convergence and central limit theorems for multidimensional hastings and metropolis algorithms. *Biometrika*, 83(1):95–110, 1996.
- [66] R. Tyrell Rockafellar. Lagrange multipliers and optimality. *SIAM Review*, 35(2):183–238, 1993.
- [67] Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by denoising (red). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017.
- [68] A.M. Rubinov. Radiant sets and their gauges. In: *Quasidifferentiability and Related Topics*, 43:235–261, 2000.
- [69] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.
- [70] Karin Schnass. On the identifiability of overcomplete dictionaries via the minimisation principle underlying k-svd. *Applied and Computational Harmonic Analysis*, 37(3):464–491, 2014.
- [71] Karin Schnass. Convergence radius and sample complexity of itkm algorithms for dictionary learning. *Applied and Computational Harmonic Analysis*, 45(1):22–58, 2016.
- [72] Rolf Schneider. Convex bodies: The brunn–minkowski theory. *Cambridge: Cambridge University Press.*, 2013.
- [73] Ivan Selesnick. Sparse regularization via convex analysis. *IEEE Transactions on Signal Processing*, 65(17):4481–4494, 2017.
- [74] Parikshit Shah, Badri Narayan Bhaskar, Gongguo Tang, and Benjamin Recht. Linear system identification via atomic norm regularization. *Proceedings of the 51st Annual Conference on Decision and Control*, 2012.
- [75] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, USA, 2014.
- [76] Yong Sheng Soh and Venkat Chandrasekaran. Learning semidefinite regularizers. *Foundations of Computational Mathematics*, 19:375–434, 2019.
- [77] Daniel A. Spielman, Huan Wang, and John Wright. Exact recovery of sparsely-used dictionaries. *Conference on Learning Theory*, 23(37):1–18, 2012.
- [78] Björn Sprungk. On the local lipschitz stability of bayesian inverse problems. *Inverse Problems*, 36:055015, 2020.
- [79] Andrew M. Stuart. Inverse problems: a bayesian perspective. *Acta Numerica*, 19:451–559, 2010.

- [80] Ju Sun, Qing Qu, and John Wright. Complete dictionary recovery over the sphere i: Overview and the geometric picture. *IEEE Transactions on Information Theory*, 63(2):853–884, 2016.
- [81] Ju Sun, Qing Qu, and John Wright. Complete dictionary recovery over the sphere ii: Recovery by riemannian trust-region method. *IEEE Transactions on Information Theory*, 63(2):885–914, 2016.
- [82] Qiyu Sun. Recovery of sparsest signals via ℓ^q -minimization. *Applied and Computational Harmonic Analysis*, 32(3):329–341, 2012.
- [83] Grzegorz Sójka. Metrics in the family of star bodies. *Advances in Geometry*, 13:117–144, 2013.
- [84] Gongguo Tang, Badri Narayan Bhaskar, Parikshit Shah, and Benjamin Recht. Compressed sensing off the grid. *IEEE Transactions on Information Theory*, 59(11):7465–7490, 2013.
- [85] Ryan Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1994.
- [86] Andrey Nikolayevich Tikhonov. On stability of inverse problems. *Dokl. Akad. Nauk SSSR*, 39(5):176–179, 1943.
- [87] Yann Traonmilin, Rémi Gribonval, and Samuel Vaiter. A theory of optimal convex regularization for low-dimensional recovery. *arXiv preprint arXiv:2112.03540*, 2021.
- [88] Daniel Vainsencher, Shie Mannor, and Alfred M. Bruckstein. The sample complexity of dictionary learning. *Journal of Machine Learning Research (JMLR)*, 12:3259–3281, 2011.
- [89] Santosh Vempala. Geometric random walks: A survey. *Combinatorial and Computational Geometry*, 52(2):573–612, 2005.
- [90] Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 945–948. IEEE, 2013.
- [91] Roman Vershynin. High-dimensional probability: An introduction with applications in data science. *Cambridge University Press*, 2020.
- [92] Cédric Villani. Topics in optimal transport. *Providence, RI: American Mathematical Society*, 2003.
- [93] Zhaoran Wang, Han Liu, and Tong Zhang. Optimal computational and statistical rates of convergence for sparse nonconvex learning problems. *Annals of Statistics*, 42(6):2164 – 2201, 2014.
- [94] Brendt Wohlberg. Efficient algorithms for convolutional sparse representations. *IEEE Transactions on Image Processing*, 25(1):301–315, 2015.
- [95] Quanming Yao and James T. Kwok. Efficient learning with a family of nonconvex regularizers by redistributing nonconvexity. *Journal of Machine Learning Research (JMLR)*, 18:1–52, 2018.
- [96] Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *Annals of Statistics*, 38(2):894 – 942, 2010.
- [97] Yixing Zhang, Xiuyuan Cheng, and Galen Reeves. Convergence of gaussian-smoothed optimal transport distance with sub-gamma distributions and dependent samples. *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, 2021.