

# OUTLIER-ROBUST ORTHOGONAL REGRESSION ON MANIFOLDS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Motivated by machine learning and computer vision applications, we formulate the problem of Outlier-Robust Orthogonal Regression to find a point in a manifold that satisfies as many linear equations as possible. Existing approaches addressing special cases of our formulation either lack theoretical support, are computationally costly, or somewhat ignore the manifold constraint; the latter two limit them from many applications. In this paper, we propose a unified approach based on solving a non-convex and non-smooth  $\ell^1$  optimization problem over the manifold. We give conditions on the geometry of the input data, the manifold, and their interplay, under which the minimizers recover the ground truth; notably the conditions can hold even when the inliers are skewed within the true hyperplane. We provide a Riemannian subgradient method and an iteratively reweighted least squares method, suiting different computational oracles, and prove their linear/sub-linear convergence to minimizers/critical points. Experiments demonstrate that respecting the manifold constraints increases robustness against outliers in robust essential matrix estimation and robust rotation search.

## 1 INTRODUCTION

Given a dataset  $\mathcal{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_L] \in \mathbb{R}^{D \times L}$  of  $L$  points and a manifold  $\mathcal{M} \subset \mathbb{R}^D$ , we propose *Outlier-Robust Orthogonal Regression* ( $\text{OR}^2$ ), formulated using the  $\ell^0$  semi-norm as

$$\min_{\mathbf{b}} \|\mathcal{Y}^\top \mathbf{b}\|_0 \quad \text{s.t.} \quad \mathbf{b} \in \mathcal{M}, \quad \mathbf{b} \neq \mathbf{0}. \quad (\text{OR}^2\text{-}\ell^0)$$

In words,  $(\text{OR}^2\text{-}\ell^0)$  is an  $\ell^0$  regression problem of finding some vector  $\mathbf{b}^*$  from  $\mathcal{M}$  that is *orthogonal* to as many points in  $\mathcal{Y}$  as possible. We assume such  $\mathbf{b}^*$  is orthogonal to  $N$  points in  $\mathcal{Y}$  ( $N \leq L$ ). Points of  $\mathcal{Y}$  orthogonal to  $\mathbf{b}^*$  are called *inliers*, or otherwise *outliers*. By solving  $(\text{OR}^2\text{-}\ell^0)$  one not only finds  $\mathbf{b}^*$ , but also identifies inliers and outliers, thus  $(\text{OR}^2\text{-}\ell^0)$  is an *outlier-robust* formulation.

$\text{OR}^2$  has massive applications in *machine learning* such as robust subspace recovery ( $\mathcal{M}$  is the unit sphere), hyperplane clustering (the Chow variety), fixed-rank matrix sensing (fixed-rank matrices), and in *geometric vision* including robust rotation search ( $\mathbb{S}\mathbb{O}(3)$ ), essential matrix estimation (the essential manifold), and trifocal tensor estimation (the calibrated trifocal tensor variety), among many others. To further motivate the paper, we begin with a tour of two important instances of  $\text{OR}^2$ .

**$\text{OR}^2$  over the Sphere  $\mathbb{S}^{D-1}$ .** In the late 1980s, [Spath & Watson \(1987\)](#) first considered the problem

$$\min_{\mathbf{b}} \|\mathcal{Y}^\top \mathbf{b}\|_1 \quad \text{s.t.} \quad \mathbf{b} \in \mathcal{M} = \mathbb{S}^{D-1}, \quad (1)$$

and called it  $\ell^1$  *orthogonal regression*, which motivates the name of this paper. (1) can be viewed as the  $\ell^1$  relaxation of  $(\text{OR}^2\text{-}\ell^0)$  with a choice of  $\mathcal{M} = \mathbb{S}^{D-1}$  to fix  $\mathbf{b}$ 's scale, since  $(\text{OR}^2\text{-}\ell^0)$  is invariant to scalings. It was not until the past decade that a fruitful line of research considered (1) (sometimes with other surrogate objectives) in robust subspace recovery ([Tsakiris & Vidal, 2015; 2017; 2018b; Zhu et al., 2018; Ding et al., 2019a; 2021](#)), dictionary learning ([Qu et al., 2014; Sun et al., 2017b;a; Bai et al., 2018; Gilboa et al., 2018](#)), and sparse blind deconvolution ([Qu et al., 2019](#)); see [Qu et al. \(2020\)](#) for a summary. Despite being non-convex non-smooth, (1) enjoys strong guarantees of successful recovery. For example, [Zhu et al. \(2018\)](#) shows that (1) is robust to as many outliers as the square of the number of inliers, and [Zhu et al. \(2018\)](#) furthermore provides linearly convergent algorithms to find  $\mathbf{b}^*$ .




Figure 1: Inliers (blue) and outliers (orange) over  $\mathbb{S}^2$  (left); objective of (1) over  $\mathbb{S}^2$  (right), with the true  $\mathbf{b}^*$  (green) and the minimizer (red).

**OR<sup>2</sup> over the Essential Manifold.** Given two images of a static scene taken from a camera that undergoes a rotation  $\mathbf{R}^* \in \mathbb{SO}(3)$  and translation  $\mathbf{t}^*$ , how can we estimate  $\mathbf{R}^*, \mathbf{t}^*$ ? This is a classic textbook question (Ma et al., 2003; Hartley & Zisserman, 2004) in geometric vision and has been a building block of modern autonomous driving systems. To solve it, some matching algorithm is typically run first on the two images to produce *correspondences*  $\{(\mathbf{z}_j, \mathbf{z}'_j)\}_{j=1}^L \subset \mathbb{S}^2 \times \mathbb{S}^2$ . *Inlier correspondences*  $(\mathbf{z}_j, \mathbf{z}'_j)$  satisfy

$$\mathbf{z}_j^\top \mathbf{E}^* \mathbf{z}'_j = 0 \quad \Leftrightarrow \quad \text{tr}(\mathbf{z}'_j \mathbf{z}'_j^\top \mathbf{E}^*) = 0, \quad (2)$$

where  $\mathbf{E}^* = [\mathbf{t}^*]_{\times} \mathbf{R}^*$  is the so-called *essential matrix*<sup>1</sup>, lying in the *essential manifold*

$$\mathcal{M}_E := \{\mathbf{E} \in \mathbb{R}^{3 \times 3} : \mathbf{E} = [\mathbf{t}]_{\times} \mathbf{R} / \sqrt{2}, \quad \mathbf{t} \in \mathbb{S}^2, \quad \mathbf{R} \in \mathbb{SO}(3)\}, \quad (3)$$

in which we restricted  $\mathbf{t} \in \mathbb{S}^2$  as (2) is invariant to the scale of  $\mathbf{t}$ . Note then that (2) is equivalent to

$$\langle \mathbf{y}_j, \mathbf{b}^* \rangle = 0, \quad \text{where} \quad \mathbf{y}_j = \text{vec}(\mathbf{z}'_j \mathbf{z}'_j^\top), \quad \mathbf{b}^* = \text{vec}(\mathbf{E}^*), \quad (4)$$

which naturally leads to two different approaches for estimating  $\mathbf{E}^*$  (and therefore  $\mathbf{R}^*$  and  $\mathbf{t}^*$ ):

- (*Essential Relaxation?*) Observe that  $\mathbf{E}^*$  satisfies  $\text{tr}(\mathbf{E}^* (\mathbf{E}^*)^\top) = 1$ , so  $\mathbf{b}^* = \text{vec}(\mathbf{E}^*)$  lies in a *proper* subset of  $\mathbb{S}^8$ . Therefore a first approach is to relax the constraint  $\mathbf{E}^* \in \mathcal{M}_E$  and simply assume that  $\mathbf{b}^*$  in (4) lies in the unit sphere  $\mathbb{S}^8$ . This idea is classic in geometric vision and has received significant popularity (Hartley & Zisserman, 2004). In particular, with this relaxation, robustly estimating an essential matrix reduces to (1), which can be solved via diverse sets of methods mentioned above and notably (Ding et al., 2020b), with powerful theoretical guarantees. However, these guarantees all build on the assumption that inliers and outliers are somewhat uniformly distributed, which is violated here. Specifically, the entries of each  $\mathbf{y}_j = \text{vec}(\mathbf{z}'_j \mathbf{z}'_j^\top)$  in (4) are correlated; it has been further proved that often  $\mathbf{y}_j$ 's are *clustered* in some proper subspace (Ding et al., 2020b) under certain *degeneracies*, rather than uniformly distributed. Figure 1 presents a visual analogy in  $\mathbb{R}^3$ . On the left of Figure 1a (*resp.* Figure 1b), the blue points are uniform (*resp.* non-uniform) inliers and orange points are outliers. Then we see that, for uniform inliers, (1) has a unique up-to-sign minimizer in red coinciding with  $\mathbf{b}^*$  in green (Figure 1a, right), while non-uniform inliers, in contrast, lead to an arbitrary minimizer far from  $\mathbf{b}^*$  (Figure 1b, right).
- (*OR<sup>2</sup> on  $\mathcal{M}_E$ ?*) To address the issue, one can consider  $\ell^1$  orthogonal regression directly on  $\mathcal{M}_E$ :

$$\min_{\mathbf{b}} \|\mathcal{Y}^\top \mathbf{b}\|_1 \quad \text{s.t.} \quad \mathbf{b} \in \mathcal{M} = \mathcal{M}_E. \quad (5)$$

That being said, we are not aware of any existing theoretical guarantees for (5) regarding its minimizers or recovery algorithms. More broadly, to the best of our knowledge, such guarantees have rarely if ever been derived for many other similar outlier-robust estimation problems in geometric vision that admit an orthogonal regression formulation and are associated naturally with manifold constraints (e.g., rotations). This thus invokes a strong call for carrying out a theoretical and principled investigation into orthogonal regression in a general, manifold optimization context.

<sup>1</sup>For any  $\mathbf{a} \in \mathbb{R}^3$ ,  $[\mathbf{a}]_{\times} := \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix}$  is a matrix representing cross product, i.e.,  $[\mathbf{a}]_{\times} \mathbf{b} = \mathbf{a} \times \mathbf{b}$ .

**Contributions.** Going beyond unit spheres and essential manifolds, we propose to study

$$\min_{\mathbf{b}} \|\mathcal{Y}^\top \mathbf{b}\|_1 \quad \text{s.t.} \quad \mathbf{b} \in \mathcal{M} \quad (\text{OR}^2\text{-}\ell^1)$$

where  $\mathcal{M}$  is assumed to be a smooth and compact submanifold of the unit sphere  $\mathbb{S}^{D-1}$  (Assumption 1). This is a non-convex and non-smooth optimization problem. We consider two efficient algorithms for (OR<sup>2</sup>-ℓ<sup>1</sup>): a Riemannian subgradient method (Algorithm 1) when one can compute a retraction map of  $\mathcal{M}$  and a projection map onto  $T_{\mathbf{b}}\mathcal{M}$ , or otherwise an iteratively reweighted least squares method (Algorithm 2) when a weighted least squares problem over  $\mathcal{M}$  can be easily solved.

- *Landscape:* We provide conditions on the geometry of the inliers, outliers,  $\mathcal{M}$ , and their interplay, under which the minimizer of (OR<sup>2</sup>-ℓ<sup>1</sup>) recovers (Theorems 1 and 2) or is close to (Theorem 3) the ground truth. Notably, (OR<sup>2</sup>-ℓ<sup>1</sup>) admits a successful recovery even if the inliers are skewed within the true hyperplane, as long as  $\mathcal{M}$  is in some sense posed close to the mode of inliers (Figure 2) and does not curve too much near the ground truth (Figure 3); a novel insight to our knowledge.
- *Convergence:* We prove that the Riemannian subgradient method converges linearly to the ground truth with suitable initialization and step sizes (Theorem 4). Moreover, we show that the iteratively reweighted least squares converges globally sub-linearly to a critical point of a surrogate to (Theorem 5); this has not been shown in the literature even for particular instances of (OR<sup>2</sup>-ℓ<sup>1</sup>).
- *Experiments:* We show that respecting the manifold constraints (as opposed to relaxing them to the unit sphere ones) increases robustness against outliers in tasks of robust estimation of essential matrix (§5) and rotation (A.2) with different parameters (e.g., number of points, outlier ratios).

## 2 PRIOR ART: OUTLIER-ROBUST ESTIMATION

Here we review two types of popular approaches for outlier-robust estimation in geometric vision.

**RANSAC.** One of the most widely applied approach for (OR<sup>2</sup>-ℓ<sup>0</sup>) is RANdom SAmple Consensus (Fischler & Bolles, 1981) that arised in the early 1980s. The most basic version operates by iteratively i) sampling say  $d < D$  points from  $\mathcal{Y}$ , ii) fitting to them a model from  $\mathcal{M}$ , iii) computing a score that measures how many points in  $\mathcal{Y}$  are consensus with this model. It terminates and returns the model with the highest score, when the probability of having sampled an outlier-free model is above a prescribed threshold.

A well-known drawback of RANSAC is, however, that the iteration complexity increases dramatically with the percent of outliers in  $\mathcal{Y}$  as well as  $d$ . Efforts mitigating the latter have unfortunately led to a dilemma: While it is possible to reduce  $d$  to the theoretical lower bound, the dimension of  $\mathcal{M}$ , by utilizing as many algebraic constraints from  $\mathcal{M}$  as possible (Nistér, 2004), yet decreasing  $d$  undermines the accuracy of step ii) against noise in  $\mathcal{Y}$ .

We conclude by remarking that numerous heuristics have been proposed to improve RANSAC in aspects of sampling (Chum & Matas, 2005; Torr et al., 2002; Barath et al., 2019a), score function (Torr & Zisserman, 2000; Stewart, 1995; Moisan et al., 2016), degeneracy handling (Chum et al., 2005; Ivashchkin et al., 2021), refitting (Chum et al., 2003; Lebeda et al., 2012; Barath & Matas, 2018; Barath et al., 2020; 2019b; Yang et al., 2021), and combinations thereof (Raguram et al., 2013; Ivashchkin et al.). Despite their empirical success on specific problem and dataset instances, little has been done on the theoretical understanding of when and why these variants offer accurate or efficient recovery. Indeed, the only theoretical result we are aware of is on the *basic* RANSAC algorithm in the problem of robust subspace recovery (Maunu & Lerman, 2019, §4).

**Continuous Robust Surrogates.** Since the ℓ<sup>0</sup> objective of (OR<sup>2</sup>-ℓ<sup>0</sup>) is discrete, one naturally considers its continuous surrogates to facilitate continuous optimization. Indeed, people have replaced the objective in (OR<sup>2</sup>-ℓ<sup>0</sup>) with ℓ<sup>1</sup>, Huber, Logcosh or truncated quadratic, albeit these works are restricted to particular instances of  $\mathcal{M}$ , to be reviewed in the sequel. Observe that the relaxed problem is still *non-convex* since the feasible set  $\mathcal{M}$  is so.

When  $\mathcal{M}$  is the set of rotations  $\mathbb{SO}(3)$ , Lajoie et al. (2018); Yang & Carlone (2019; 2022); Peng et al. (2022a) considers a surrogate of (OR<sup>2</sup>-ℓ<sup>0</sup>) using the truncated quadratic objective. Notably, they established its equivalence to a quadratically constrained quadratic programming (QCQP), which is in turn relaxed to a semi-definite programming (SDP), a *convex* problem. This allows computing and analyzing the minimizer of QCQP and SDP problems under certain conditions; it

has been shown to be rather robust to noise and outliers. However, this approach scales poorly with the number of samples  $L$ : indeed, the SDP has  $O(L^2)$  variables, hindering real-time computation for even  $L = 50$ .

A fruitful series of research consider instead optimizing directly some surrogates of  $(\text{OR}^2\text{-}\ell^0)$ , allowing much faster computation. When  $\mathcal{M}$  is  $\mathbb{S}^{D-1}$ , the minimizer of  $(\text{OR}^2\text{-}\ell^1)$  recovers the ground truth (Zhu et al., 2018), and (Projected) Riemannian sub-gradient method with suitable initialization and step sizes converges (piece-wise (Zhu et al., 2018)) linearly (Zhu et al., 2019; Li et al., 2019) to the minimizer. Nevertheless, these results requires, among other conditions, that the inliers to be uniformly distributed in the hyperplane orthogonal to the true  $\mathbf{b}^*$ , a condition that can often be violated, e.g., in geometric vision (see Example 1 and experiments). On the other hand, when  $\mathcal{M}$  is the normalize essential variety, a family of surrogates *empirically* enjoys accurate and efficient recovery (Zhao, 2022); this in principle applies to a rich family of  $\mathcal{M}$  as long as weighted least squares problems can be solved (§3.2). Be that as it may, there is no theoretical guarantee whatsoever telling when the minimizer is the ground truth or why algorithms should converge to it.

### 3 FORMULATION

In this section, we first formulate the problem in a precise manner (§3.1), and then discuss two algorithms useful when different mappings can be efficiently computed with respect to the manifold  $\mathcal{M}$  (§3.2). The theoretical analysis of the objective and algorithms are deferred to §4.

#### 3.1 SET UP: DATA MODEL, ASSUMPTIONS, NOTATIONS

To begin with, we define the manifold and the data model, making clear the associated assumptions. The first assumption we make is the following.

**Assumption 1** (Regularity on manifold).  $\mathcal{M}$  is a smooth and compact submanifold of  $\mathbb{S}^{D-1}$ .

Let  $\mathcal{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_L] \in \mathbb{R}^{D \times L}$  be a given *normalized* dataset, i.e.,  $\mathbf{y}_i \in \mathbb{S}^{D-1}$ . When the context is clear we also treat  $\mathcal{Y}$  as the set  $\{\mathbf{y}_i\}_{i=1}^L$  with some abuse of notation. Let  $\mathcal{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]$  be the subset of  $\mathcal{Y}$  containing points from  $\mathcal{Y}$  that are orthogonal to a ground truth vector  $\mathbf{b}^* \in \mathcal{M}$ , which we refer to as *inliers*. It turns out to be handy to define the ground truth hyperplane  $\mathcal{H} := \text{span}(\mathbf{b}^*)^\perp$ . On the other hand, denote the points from  $\mathcal{Y} \setminus \mathcal{X}$  by  $\mathcal{O} = [\mathbf{o}_1, \dots, \mathbf{o}_M]$ , termed *outliers*; observe that  $L = N + M$ . Finally, we make a mild assumption on the outliers.

**Assumption 2** (Outliers are in general position). Any size- $D$  subset of  $\mathcal{O}$  are linearly independent.

*Remark 1.* The requirements that points  $\mathbf{y}_i$ 's from  $\mathcal{Y}$  lie on the unit sphere  $\mathbb{S}^{D-1}$  and  $\mathcal{M}$  is a submanifold of  $\mathbb{S}^{D-1}$  do not compromise generality. Note that the objective of  $(\text{OR}^2\text{-}\ell^0)$  is invariant to the scale of  $\mathbf{b}$  and points from  $\mathcal{Y}$ , i.e.,  $(\text{OR}^2\text{-}\ell^0)$  concerns  $\mathbf{b}$  and  $\mathbf{y}_i$ 's only in *directions* but not *scaling*. As such, for an arbitrary dataset we can project<sup>2</sup> its points onto  $\mathbb{S}^{D-1}$  to form a qualified  $\mathcal{Y}$ ; for an *arbitrary* smooth manifold  $\overline{\mathcal{M}}$ , taking  $\Pi_{\mathbb{S}^{D-1}}(\overline{\mathcal{M}} \setminus \{\mathbf{0}\})$  gives a  $\mathcal{M}$  obeying Assumption 1.

*Remark 2.* Assumption 2 is not a strong assumption, since for outliers from a uniform distribution over  $\mathbb{S}^{D-1}$  it holds with probability 1.

**Notations.** We denote the objective function of  $(\text{OR}^2\text{-}\ell^1)$  as  $f(\cdot)$ , i.e.,  $f(\mathbf{b}) = \|\mathcal{Y}^\top \mathbf{b}\|_1$ . For a set  $\mathcal{S} \subset \mathbb{R}^D$  and a vector  $\mathbf{x} \in \mathbb{R}^D$ ,  $\Pi_{\mathcal{S}}(\mathbf{x}) := \operatorname{argmin}_{\mathbf{y}} \{\|\mathbf{x} - \mathbf{y}\|_2^2 : \mathbf{y} \in \mathcal{S}\}$  is the projection of  $\mathbf{x}$  onto  $\mathcal{S}$  provided it exists. For any  $\mathbf{v} \in \mathbb{R}^D$ ,  $\langle \mathbf{v} \rangle$  denotes the linear subspace spanned by  $\mathbf{v}$ . For any linear subspace  $\mathcal{S}$  of  $\mathbb{R}^D$ ,  $\mathcal{S}^\perp$  is its orthogonal complement. Let  $\operatorname{sign}(\cdot)$  be such that  $\operatorname{sign}(x)$  is 1 if  $x > 0$ ,  $-1$  if  $x < 0$  and 0 otherwise. Denote by  $T_{\mathbf{b}}\mathcal{M}$  the tangent space of  $\mathcal{M}$  at  $\mathbf{b} \in \mathcal{M}$ . Given  $\mathbf{b} \in \mathcal{M}$  we denote by  $\operatorname{retr}_{\mathbf{b}} : T_{\mathbf{b}}\mathcal{M} \rightarrow \mathcal{M}$  the retraction map that sends a vector  $\mathbf{v}$  in the tangent space  $T_{\mathbf{b}}\mathcal{M}$  back to the manifold  $\mathcal{M}$ . We use  $\operatorname{rgrad}$  for the Riemannian gradient.

<sup>2</sup>A minor detail here is that there is not a unique projection of the zero vector  $\mathbf{0}$  onto  $\mathbb{S}^{D-1}$ , a point that we can remove from  $\mathcal{Y}$  regardless without affecting minimizers of  $(\text{OR}^2\text{-}\ell^0)$ .

**Algorithm 1:** RSGM for  $(\text{OR}^2\text{-}\ell^1)$ 

- 1 Initialization  $\mathbf{b}_0 \in \mathcal{M}$ , step sizes:  $\{\alpha_k > 0\}_{k=1,2,\dots}$
- 2 For  $k \leftarrow 1, 2, \dots, K$ :

$$\mathbf{g}_k \in \partial f(\mathbf{b}_k) \quad (6)$$

$$\mathbf{r}_k \leftarrow \Pi_{T_{\mathbf{b}_k} \mathcal{M}} \mathbf{g}_k \quad (7)$$

$$\mathbf{b}_{k+1} \leftarrow \text{retr}_{\mathbf{b}_k}(-\alpha_k \mathbf{r}_k) \quad (8)$$

**Algorithm 2:** IRLS for  $(\text{OR}^2\text{-}\ell^1)$ 

- 1 Initialization  $\mathbf{b}_0 \in \mathcal{M}$ , parameter  $\delta > 0$
- 2 For  $k \leftarrow 0, \dots, K$ :

$$w_j^{k+1} \leftarrow \frac{1}{\max\{|\mathbf{y}_j^\top \mathbf{b}_k|, \delta\}}, \quad \forall j = 1, \dots, L \quad (9)$$

$$\mathbf{b}_{k+1} \in \operatorname{argmin}_{\mathbf{b} \in \mathcal{M}} \sum_{j=1}^L w_j^{k+1} \cdot (\mathbf{y}_j^\top \mathbf{b})^2 \quad (10)$$

## 3.2 ALGORITHMS

To solve  $(\text{OR}^2\text{-}\ell^1)$ , we describe two algorithms: a Riemannian subgradient method when a retraction map of  $\mathcal{M}$  and a projection map onto  $T_{\mathbf{b}} \mathcal{M}$  are accessible, or otherwise an iteratively reweighted least squares method when a weighted least squares problem over  $\mathcal{M}$  can be easily solved.

**Riemannian Sub-gradient Method (RSGM).** Algorithm 1 gives the classic RSGM algorithm for  $(\text{OR}^2\text{-}\ell^1)$ , which iteratively computes a subgradient (6), projects it onto the tangent space (7), and takes a retraction step (8). To give a concrete example, suppose for the moment that  $\mathcal{M}$  is the unit sphere. In this case the computation of the subgradient in step (6) can be contextualized as  $\mathbf{g}_k \leftarrow \sum_{j=1}^L \text{sign}(\mathbf{y}_j^\top \mathbf{b}_k) \mathbf{y}_j$ , the projection step (7) is given as  $\mathbf{r}_k \leftarrow \mathbf{g}_k - \mathbf{b}_k \mathbf{b}_k^\top \mathbf{g}_k$ , and the retraction step (8) can be implemented simply as a projection onto the sphere,  $\mathbf{b}_{k+1} \leftarrow (\mathbf{b}_k - \alpha_{k+1} \mathbf{r}_k) / \|\mathbf{b}_k - \alpha_{k+1} \mathbf{r}_k\|_2$ . In general, the subgradient step (6) can be computed for typical objective functions and the two Riemannian optimization steps (7) and (8) can be computed efficiently for a wide range of manifolds  $\mathcal{M}$ 's. Therefore, we assume that Algorithm 1 can be implemented efficiently, and will analyze its convergence rate in §4.3 and performance in applications in §5.

**Iteratively Reweighted Least Squares (IRLS).** Algorithm 2 presents the classic IRLS framework to solve  $(\text{OR}^2\text{-}\ell^1)$ . IRLS alternates between updating the weights (9) and solving a weighted least-squares problem (10). The choice of weight  $w_j^{k+1}$  in (9) is a classic one and has been used in prior works (Tsakiris & Vidal, 2018a; Ding et al., 2019b; 2020a; Iwata et al., 2020; Kümmerle et al., 2021; Peng et al., 2022b). While we have not seen the appearance of (10) in its exact form in the literature, several papers have considered (10) with particular  $\mathcal{M}$ 's for the problems of robust subspace recovery (Tsakiris & Vidal, 2018a; Ding et al., 2019b; 2020a) and essential matrix estimation (Zhao et al., 2020; Zhao, 2022). In the former case, (10) is a weighted eigenvalue problem, easily solvable; in the latter case, (10) is a QCQP, which can be solved via its SDP relaxations. Since (10) does not involve binary decision variables (different from Lajoie et al. (2018); Yang & Carlone (2019; 2022); Peng et al. (2022a)), the SDP solver of Zhao et al. (2020); Zhao (2022) is actually efficient. Given these, we implicitly assume that (10) is efficiently solvable, and later we will analyze the convergence rate of IRLS in generality and evaluate its performance in action.

## 4 THEORETICAL ANALYSIS

In this section, we first characterize the geometry of inliers, outliers, and their interplay with  $\mathcal{M}$  via some quantities (§4.1). They turn out to be useful in the analysis of the minimizers (§4.2) of and algorithms (§4.3) for  $(\text{OR}^2\text{-}\ell^1)$ .

## 4.1 GEOMETRIC QUANTITIES

We will define quantities that depend on the inliers  $\mathcal{X}$ , outliers  $\mathcal{O}$ , the true vector  $\mathbf{b}^*$  (equivalently, the true hyperplane  $\mathcal{H}$ ), and the manifold  $\mathcal{M}$ . It should be clear that these quantities are for theoretical understanding and not for running the algorithms. To begin with, define

$$c_{\text{in}} := \inf \left\{ \frac{1}{N} \sum_{j=1}^N |\mathbf{x}_j^\top \mathbf{s}| : \mathbf{b} \in \mathcal{M}, \mathbf{b} \neq \mathbf{b}^*, \mathbf{s} = \Pi_{\mathcal{S}^{D-1}}(\Pi_{\mathcal{H}}(\mathbf{b})) \right\}. \quad (11)$$




Figure 2: The non-uniform data from Figure 1 (left), and  $f(\cdot)$  over  $\mathbb{S}^2$  using the data overlaid by two  $\mathcal{M}$ 's (dark solid lines) yielding large  $c_{\text{in}}$  (middle) and small  $c_{\text{in}}$  (right) respectively. Dotted lines are locations of  $s$ 's in  $c_{\text{in}}$  (11). For clarity, the viewpoints are set at the true  $\mathbf{b}^*$  facing towards the origin.

*Remark 3.* To understand  $c_{\text{in}}$ , note that one can expect it to be large when for any  $\mathbf{b} \in \mathcal{M}$ ,  $\mathbf{b} \neq \mathbf{b}^*$ , its projection  $s$  onto  $\mathcal{H}$  is *not* orthogonal to most inliers  $\mathbf{x}_j$ 's. This is visualized in Figure 2, where  $c_{\text{in}}$  is large as the dotted lines are far from the equator  $z = 0$  (middle), and vice versa (right). To better appreciate the role of  $\mathcal{M}$  in  $c_{\text{in}}$ , we make a contrast with a similar quantity in [Zhu et al. \(2018\)](#) for guaranteeing (1):

$$\bar{c}_{\text{in}} = \inf \left\{ \frac{1}{N} \sum_{j=1}^N |\mathbf{x}_j^\top \mathbf{s}| : \mathbf{s} \in \mathcal{H} \cap \mathbb{S}^{D-1} \right\}. \quad (12)$$

Observe that  $\bar{c}_{\text{in}}$  is large only when *all directions* in  $\mathcal{H}$  are not orthogonal to most inliers, essentially requiring inliers to be uniformly distributed in  $\mathcal{H} \cap \mathbb{S}^{D-1}$ . Such an assumption can fail in applications as illustrated in Example 1. In contrast, one can have large  $c_{\text{in}}$  albeit some vectors in  $\mathcal{H}$  are orthogonal to a vast number of inliers, as long as these directions are ruled out by  $\mathcal{M}$ .

*Example 1.* In geometric vision applications, inliers tend to lie close to a *proper* subspace of the true hyperplane. For instance, in *essential matrix* or *trifocal tensor estimation*, data resulting from 3D points on an affine plane provably lies in a subspace of lower dimension ([Chum et al., 2005](#); [Ding et al., 2020b](#)); this is more generally known as critical surfaces ([Hartley & Zisserman, 2003](#), §22).

With the above said, we can further define similar quantities for outliers:

$$c_{\text{out},\min} := \min_{\mathbf{b} \in \mathbb{S}^{D-1}} \frac{1}{M} \sum_{j=1}^M |\mathbf{o}_j^\top \mathbf{b}|, \quad c_{\text{out},\max} := \max_{\mathbf{b} \in \mathbb{S}^{D-1}} \frac{1}{M} \sum_{j=1}^M |\mathbf{o}_j^\top \mathbf{b}|. \quad (13)$$

*Remark 4.* It is now not hard to see that  $c_{\text{out},\min}$  and  $c_{\text{out},\max}$  are close if the outliers are uniformly distributed in  $\mathbb{S}^{D-1}$ . While in principle one can consider the  $\mathcal{M}$ -constrained version of these outlier-related quantities, we find doing so complicates the picture without shedding useful insights.

Finally, we define

$$\gamma_{\text{in},\theta} := \max \left\{ \frac{1}{N} \left\| \mathbf{\Pi}_{T_{\mathbf{b}}\mathcal{M}^\perp \cap \langle \mathbf{b} \rangle^\perp} \sum_{j=1}^N \text{sgn}(\mathbf{x}_j^\top \mathbf{b}) \mathbf{x}_j \right\|_2 : \mathbf{b} \in \mathcal{B}_{\mathcal{M}}(\mathbf{b}^*, \theta) \right\}, \quad (14)$$

$$\bar{\eta}_{\text{out}} := \max_{\mathbf{b} \in \mathbb{S}^{D-1}} \frac{1}{M} \left( \left\| \mathbf{\Pi}_{T_{\mathbf{b}}\mathcal{M}} \sum_{j=1}^M \text{sign}(\mathbf{o}_j^\top \mathbf{b}) \mathbf{o}_j \right\|_2 + D \right), \quad (15)$$

where  $\mathcal{B}_{\mathcal{M}}(\mathbf{b}^*, \theta)$  is a ball of angular radius  $\theta$  centered at  $\mathbf{b}^*$  over  $\mathcal{M}$

$$\mathcal{B}_{\mathcal{M}}(\mathbf{b}^*, \theta) := \{ \mathbf{b} \in \mathcal{M} : \angle(\mathbf{b}, \mathbf{b}^*) \leq \theta \}. \quad (16)$$

*Remark 5.* Roughly speaking,  $T_{\mathbf{b}}\mathcal{M}^\perp \cap \langle \mathbf{b} \rangle^\perp$  contains directions *tangent* to  $\mathbb{S}^{D-1}$  and *normal* to  $\mathcal{M}$  at  $\mathbf{b}$ , as seen by the arrows in Figure 3. As such,  $\gamma_{\text{in},\theta}$  is small if within  $\mathcal{B}_{\mathcal{M}}(\mathbf{b}^*, \theta)$ , all such directions are orthogonal to most inliers  $\mathbf{x}_j$ 's; this is, e.g., the case for the left and middle plots in Figure 3 with  $\theta = \frac{\pi}{2}$  and  $\theta = \frac{\pi}{3}$  respectively, but *not* the case for the right plot with  $\theta = \frac{\pi}{3}$ . Observe also that  $\gamma_{\text{in},\theta}$  is increasing in  $\theta$ . Last but not least,  $\bar{\eta}_{\text{out}}$  is small if outliers are uniformly distributed in  $\mathbb{S}^{D-1}$ .






Figure 3:  $f(\cdot)$  over  $\mathbb{S}^2$  using the non-uniform data from Figure 1 overlaid by three  $\mathcal{M}$ 's (dark solid lines) yielding  $\gamma_{\text{in},\theta}$  globally small (left), locally small (middle) and not small (right) respectively. Arrows are examples of vectors from  $T_{\mathbf{b}}\mathcal{M}^\perp \cap \langle \mathbf{b} \rangle^\perp$  in  $\gamma_{\text{in},\theta}$  (14).

#### 4.2 GUARANTEES ON GLOBAL MINIMIZERS

Using the quantities developed above, we are now ready to localize global minimizers of (OR<sup>2</sup>-ℓ<sup>1</sup>). To begin with, when  $\gamma_{\text{in},\theta}$  is small for all  $\theta \in (0, \frac{\pi}{2}]$  (e.g., Figure 3 left), we have the theorem below:

**Theorem 1.** *Let  $\gamma_{\text{in}} := \sup_{\theta \in (0, \frac{\pi}{2}]} \gamma_{\text{in},\theta}$ . Every global minimizer of (OR<sup>2</sup>-ℓ<sup>1</sup>) must be  $\pm \mathbf{b}^*$  if*

$$\frac{\frac{M^2}{N^2} \cdot \left[ (c_{\text{out,max}} - c_{\text{out,min}})^2 + \bar{\eta}_{\text{out}}^2 \right] + \gamma_{\text{in}}^2}{c_{\text{in}}^2} < 1. \quad (17)$$

Intuitively, Theorem 1 says the minimizers of (OR<sup>2</sup>-ℓ<sup>1</sup>) recover the ground truth  $\pm \mathbf{b}^*$  as long as  $N \gg M$ , inliers are well distributed with respect to  $\mathcal{M}$  ( $c_{\text{in}}$  large), outliers are well distributed ( $c_{\text{out,max}} - c_{\text{out,min}}$  and  $\bar{\eta}_{\text{out}}$  small), and  $\mathcal{M}$  does not curve too much against the inliers ( $\gamma_{\text{in}}$  small).

On the other hand,  $\gamma_{\text{in},\theta}$  could be small only for small  $\theta$  and increases drastically for large  $\theta$  (e.g., Figure 3 middle). In this case, we provide the following guarantee:

**Theorem 2.** *Every global minimizer of (OR<sup>2</sup>-ℓ<sup>1</sup>) must be  $\pm \mathbf{b}^*$  if*

$$\sin(\theta) > \frac{M}{N} \cdot \frac{c_{\text{out,max}} - c_{\text{out,min}}}{c_{\text{in}}}, \quad (18)$$

$$N^2 c_{\text{in}}^2 - N^2 \gamma_{\text{in},0}^2 \geq M^2 \bar{\eta}_{\text{out}}^2, \quad (19)$$

where  $\theta$  is the unique solution to

$$\cos^2(\theta) N^2 c_{\text{in}}^2 - N^2 \gamma_{\text{in},\theta}^2 = M^2 \bar{\eta}_{\text{out}}^2. \quad (20)$$

Note that the LHS of (20) is strictly decreasing in  $\theta$ , and non-positive with  $\theta = \frac{\pi}{2}$ ; the RHS of (20) is non-negative. Thus, with condition (19) in place, (20) is guaranteed to have a unique solution. Theorem 2 is tighter than Theorem 1: with some inspection one can see that the former entails the latter; if for some  $\theta = \bar{\theta}$  (18) holds and in (20) the LHS is greater than or equal to the RHS, then the minimizers of (OR<sup>2</sup>-ℓ<sup>1</sup>) are  $\pm \mathbf{b}^*$  regardless of the behavior of  $\gamma_{\text{in},\theta}$  for  $\theta > \bar{\theta}$ .

Finally, if  $\gamma_{\text{in},\theta}$  is large even for small  $\theta$  (e.g., Figure 3 right), we can only characterize the approximate location of the minimizer with the following theorem that does not depend on  $\gamma_{\text{in},\theta}$ .

**Theorem 3.** *Every global minimizer of (OR<sup>2</sup>-ℓ<sup>1</sup>) has an angle away from  $\pm \mathbf{b}^*$  of at most*

$$\arcsin\left(\frac{M}{N} \cdot \frac{c_{\text{out,max}} - c_{\text{out,min}}}{c_{\text{in}}}\right). \quad (21)$$

#### 4.3 CONVERGENCE RATES OF RSGM AND IRLS

We now consider convergence analysis of the proposed algorithms. We begin with Algorithm 1:

**Theorem 4.** *With a suitable initialization and geometrically diminishing step size, the sequence  $\{\mathbf{b}_k\}_{k=0}^K$  given by Algorithm 1 converges at a linear rate to the true vector  $\pm \mathbf{b}^*$ . More precisely, with the step size  $\alpha_k = \beta^k \cdot \alpha_0$ , the principal angle  $\theta_k$  between  $\mathbf{b}_k$  and  $\mathbf{b}^*$  is upper bounded as per*

$$\sin(\theta_k/2) \leq \beta^k \sin(\theta_0/2) \quad (22)$$

if the following conditions are satisfied:

- (i) Initialization not too far from  $\pm \mathbf{b}^*$ :  $\sin(\theta_0/2) < \frac{s}{2c_2}$ ,
- (ii) Small initial step size:  $0 < \alpha_0 < \min\left(\frac{\sin(\theta_0/2)}{s-2c_2 \sin(\theta_0/2)}, \frac{4(s-2c_2 \sin(\theta_0/2))}{(c_1+1)\kappa^2/\sin(\theta_0/2)}\right)$ ,
- (iii) Suitable step size decay:  $1 > \beta \geq \sqrt{1 + \alpha_0\left(\alpha_0 \frac{(c_1+1)\kappa^2}{4\sin(\theta_0/2)^2} + 2c_2 - \frac{s}{\sin(\theta_0/2)}\right)}$ .

In the above,  $s := \frac{1}{2}Nc_{\text{in}} - Mc_{\text{out,max}}$ ,  $\kappa$  is a Lipschitz constant of the objective  $f(\cdot)$ ,  $c_2 = \kappa \cdot c_3$ , and  $c_1$  and  $c_3$  are constants that depend only on  $\mathcal{M}$ .

Assuming condition (i) holds, one can verify that conditions (ii) and (iii) of Theorem 4 also holds true if the initial step size  $\alpha_0$  is small enough as  $\beta \rightarrow 1$ . Therefore, with  $\alpha_0$  and  $\beta$  suitably chosen, Algorithm 1 converges linearly if initialized in a neighborhood of  $\mathbf{b}^*$  of size  $s/c_2$  (Theorem 4).

We then proceed to Algorithm 2, an IRLS method of an entirely different style than Algorithm 1. The theoretical convergence rate guarantee, stated below, is also in a different style than Theorem 4:

**Theorem 5.** *The sequence  $\{\mathbf{b}_k\}_{k=0}^K$  produced by Algorithm 2 converges at a sublinear rate to a critical point of the following problem:*

$$\min_{\mathbf{b} \in \mathcal{M}} \sum_{j=1}^L h(\mathbf{y}_j^\top \mathbf{b}), \quad h(r) := \begin{cases} |r| & |r| > \delta \\ \frac{1}{2}(r^2/\delta + \delta) & |r| \leq \delta \end{cases} \quad (\text{OR}^2\text{-Huber})$$

More precisely, the Riemannian gradient of the objective of (OR<sup>2</sup>-Huber) is bounded above:

$$\min_{k=0,\dots,K} \left\| \text{rgrad} \sum_{j=1}^L h(\mathbf{y}_j^\top \mathbf{b}_k) \right\|_2 \leq \sqrt{\frac{2\sigma \cdot \sum_{j=1}^L h(\mathbf{y}_j^\top \mathbf{b}_0)}{K+1}} \quad (23)$$

In (23),  $\sigma$  is a finite positive constant depending on the manifold  $\mathcal{M}$ , data  $\mathbf{y}_j$ 's, and  $\delta$ .

Algorithm 2 converges to a critical point of a different Huber-type objective in (OR<sup>2</sup>-Huber) than the  $\ell^1$  objective of (OR<sup>2</sup>- $\ell^1$ ). The key is that the (rescaled and translated) Huber loss  $h(\cdot)$  in (OR<sup>2</sup>-Huber) tends to the  $\ell^1$  loss as  $\delta \rightarrow 0$ , and (OR<sup>2</sup>-Huber) can be thought of as a smoothed version of (OR<sup>2</sup>- $\ell^1$ ), easier to minimize. Theorem 5 reveals Algorithm 2 performs such smoothing implicitly with a sublinear convergence rate. Our proof of Theorem 5 extends the idea of Hong et al. (2017), which is specifically for the convex case, into the case with non-convex manifolds  $\mathcal{M}$ .

## 5 EXPERIMENTS: ROBUST ESSENTIAL MATRIX ESTIMATION

Recall that our major claim is that enforcing manifold constraints would lead to theoretically more meaningful and practically more robust estimation than relaxing, if not ignoring, the manifold constraints. While previous sections focus on providing theoretical insights, here we are interested in consolidating our claim using numerical experiments. In particular, this section presents experimental results on robust essential matrix estimation, while in the interest of space we refer the reader to the Appendix for more experiments and details.

In order to validate our claim, we consider the following setup. We follow Kneip & Furgale (2014) to generate random rotations  $\mathbf{R}^* \in \text{SO}(3)$ , translations  $\mathbf{t}^* \in \mathbb{R}^3$ , and  $L$  correspondences, of which  $N$  are inliers and  $M$  are outliers. The correspondences are used to produce  $\mathcal{Y} \in \mathbb{R}^{9 \times L}$  following (4) and a normalization of columns. We test the IRLS algorithm for solving (1) and (OR<sup>2</sup>- $\ell^1$ ). For (1), the weighted least-squares step (10) can be implemented via SVD; the resulting IRLS algorithm was proposed in Tsakiris & Vidal (2018b) and here we call it DPCP<sup>3</sup> (*dual principal component pursuit*). For (OR<sup>2</sup>- $\ell^1$ ), with  $\mathcal{M}$  being the essential manifold  $\mathcal{M}_E$ , the weighted least-squares step (10) can be implemented via an efficient SDP solver (Zhao, 2022); here we call it OR<sup>2</sup>. Each method is run on  $\mathcal{Y}$  to obtain an estimate  $\hat{\mathbf{b}}$ , while we need to project the estimate of DPCP onto  $\mathcal{M}_E$  as it does not enforce the essential manifold constraint. We compute the angular error between  $\hat{\mathbf{b}}$  and the ground truth  $\mathbf{b}^*$ . We further decompose each estimated essential matrix into rotations and translations, and calculate angular rotation and translation error with  $\mathbf{R}^*, \mathbf{t}^*$ .

<sup>3</sup>We refer the reader to Tsakiris & Vidal (2018b); Ding et al. (2020b); Zhao (2022), where DPCP or IRLS really stood out with excellent performance in comparison to other competitive methods including RANSAC.



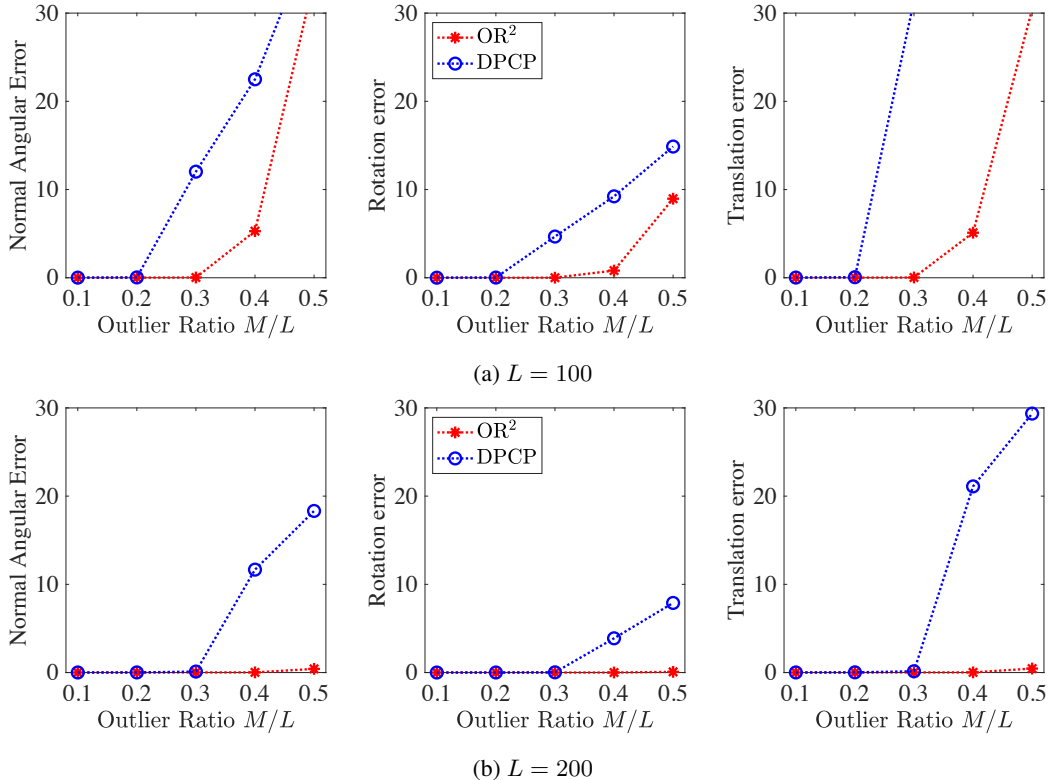


Figure 4: Robust essential matrix estimation over 50 randomly generated scene data.

In Figure 4 we report the median errors for varying outlier ratio  $M/L \in \{0.1 : 0.1 : 0.5\}$  and total number of points  $L \in \{100, 200\}$ . As expected, the two methods, DPCP and  $\text{OR}^2$ , give almost 0 error with a small (e.g.,  $\leq 20\%$ ) outlier ratio. Notably, also can be seen is that  $(\text{OR}^2-\ell^1)$  is much more robust than (1): when  $L = 100$ ,  $(\text{OR}^2-\ell^1)$  has a translation error of  $5.2^\circ$  with 40% outliers, while (1) yields that of more than  $30^\circ$  with only 30% outliers, a sharp contrast. Moreover, with more points ( $L = 200$ ),  $(\text{OR}^2-\ell^1)$  achieves less than  $1^\circ$  for any error metric even with 50% outliers, while in that case (1) has, e.g.,  $7.8^\circ$  rotation error already. The experiments therefore confirm our claim that directly tackling the manifold constraints (e.g., essential manifolds) yield visible computational benefits than relaxing the constraints (e.g., into a unit sphere).

## 6 CONCLUSION AND DISCUSSION

In this paper, we presented a meaningful generalization of the seminal work of Späth & Watson (1987) on  $\ell^1$  orthogonal regression (on the sphere) to outlier-robust orthogonal regression on smooth manifolds, thereby bringing together robustness with the proper manifold constraints that are inherent in geometric vision (e.g., essential manifolds). Our experiments suggested that taking manifold constraints into consideration does improve the performance of the estimation algorithms. Even though handling manifold constraints directly is arguably more challenging, our core contribution consists of performing theoretical analysis for the proposed approach and derived several results regarding the optimization landscape and convergence of associated algorithms; all these are ultimately useful guarantees for outlier-robust orthogonal regression on manifolds.

Our analysis is not without limitations. First of all, to obtain conditions for global minimizers (§4.2), we defined several quantities in §4.1 that are not directly computable. However, for future work it is possible to derive probabilistic bounds for them, thereby obtaining more informative conditions. The second, related, question to consider is on specializing our conditions for manifolds of interest (e.g., essential manifolds). Finally, we focused on orthogonal regression in the paper which is associated with the  $\ell^1$  (and Huber) loss. It is an important future direction to generalize our analysis for other common loss functions (e.g., Logcosh, Geman-McClure, truncated quadratic).

## REFERENCES

- Yu Bai, Qijia Jiang, and Ju Sun. Subgradient descent learns orthogonal dictionaries. October 2018. [1](#)
- Daniel Barath and Jiří Matas. Graph-cut RANSAC. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6733–6741, 2018. [3](#)
- Daniel Barath, Maksym Ivashchkin, and Jiri Matas. Progressive NAPSAC: sampling from gradually growing neighborhoods. June 2019a. [3](#)
- Daniel Barath, Jiri Matas, and Jana Noskova. MAGSAC: Marginalizing sample consensus. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2019b. [3](#)
- Daniel Barath, Jana Noskova, Maksym Ivashchkin, and Jiri Matas. MAGSAC++, a fast, reliable and accurate robust estimator. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2020. [3](#)
- Nicolas Boumal, Bamdev Mishra, P-A Absil, and Rodolphe Sepulchre. Manopt, a matlab toolbox for optimization on manifolds. August 2013. [15](#)
- Nicolas Boumal, Pierre-Antoine Absil, and Coralia Cartis. Global rates of convergence for non-convex optimization on manifolds. *IMA Journal of Numerical Analysis*, 39(1):1–33, 2019. [20](#), [22](#)
- O Chum and J Matas. Matching with PROSAC — progressive sample consensus. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*. IEEE, 2005. [3](#)
- Ondřej Chum, Jiří Matas, and Josef Kittler. Locally optimized RANSAC. In *Pattern Recognition*, pp. 236–243. Springer Berlin Heidelberg, 2003. [3](#)
- Ondřej Chum, Tomáš Werner, and Jiří Matas. Two-view geometry estimation unaffected by a dominant plane. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2005*, volume 1, pp. 772–779, 2005. [3](#), [6](#)
- Tianjiao Ding, Yunchen Yang, Zhihui Zhu, Daniel P Robinson, René Vidal, Laurent Kneip, and Manolis C. Tsakiris. Robust homography estimation via dual principal component pursuit. In *IEEE Conference on Computer Vision and Pattern Recognition, 2020a*. [5](#)
- Tianjiao Ding, Yunchen Yang, Zhihui Zhu, Daniel P Robinson, René Vidal, Laurent Kneip, and Manolis C Tsakiris. Robust homography estimation via dual principal component pursuit. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6080–6089, 2020b. [2](#), [6](#), [8](#), [15](#)
- Tianyu Ding, Zhihui Zhu, Tianjiao Ding, Yunchen Yang, Daniel P Robinson, Manolis C Tsakiris, and René Vidal. Noisy dual principal component pursuit. In *International Conference on Machine Learning*, pp. 1617–1625, 2019a. [1](#)
- Tianyu Ding, Zhihui Zhu, Tianjiao Ding, Yunchen Yang, René Vidal, Manolis C. Tsakiris, and Daniel Robinson. Noisy dual principal component pursuit. In *International Conference on Machine Learning*, 2019b. [5](#)
- Tianyu Ding, Zhihui Zhu, Manolis Tsakiris, René Vidal, and Daniel Robinson. Dual principal component pursuit for learning a union of hyperplanes: Theory and algorithms. In *Artificial Intelligence and Statistics*, 2021. [1](#)
- Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24: 381–395, 1981. [3](#)
- Dar Gilboa, Sam Buchanan, and John Wright. Efficient dictionary learning with gradient descent. September 2018. [1](#)

- Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003. 6
- Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004. 2
- Mingyi Hong, Xiangfeng Wang, Meisam Razaviyayn, and Zhi-Quan Luo. Iteration complexity analysis of block coordinate descent methods. *Mathematical Programming*, 163:85–114, 2017. 8
- Maksym Ivashechkin, Daniel Barath, and Jiri Matas. USACv20: robust essential, fundamental and homography matrix estimation. *arXiv*. 3
- Maksym Ivashechkin, Daniel Barath, and Jiri Matas. VSAC: Efficient and accurate estimator for H and F. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, October 2021. 3
- Daichi Iwata, Michael Waechter, Wen-Yan Lin, and Yasuyuki Matsushita. An analysis of sketched irls for accelerated sparse residual regression. In *European Conference on Computer Vision*, 2020. 5
- Laurent Kneip and Paul Furgale. OpenGV: A unified and generalized approach to real-time calibrated geometric vision. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2014. 8
- Christian Kümmerle, Claudio Mayrunk Verdun, and Dominik Stöger. Iteratively reweighted least squares for basis pursuit with global linear convergence rate. *Advances in Neural Information Processing Systems*, 2021. 5
- Pierre-Yves Lajoie, Siyi Hu, Giovanni Beltrame, and Luca Carlone. Modeling perceptual aliasing in SLAM via Discrete-Continuous graphical models. October 2018. 3, 5
- Karel Lebeda, Jiri Matas, and Ondrej Chum. Fixing the locally optimized RANSAC. In *Proceedings of the British Machine Vision Conference 2012*, number 95. British Machine Vision Association, 2012. 3
- Xiao Li, Shixiang Chen, Zengde Deng, Qing Qu, Zhihui Zhu, and Anthony Man Cho So. Weakly convex optimization over stiefel manifold using riemannian Subgradient-Type methods. November 2019. 4, 20
- Yi Ma, Stefano Soatto, Jana Košecká, and S Shankar Sastry. *An Invitation to 3-D Vision*. Springer New York, 2003. 2
- Tyler Maunu and Gilad Lerman. Robust subspace recovery with adversarial outliers. April 2019. 3
- Lionel Moisan, Pierre Moulon, and Pascal Monasse. Fundamental matrix of a stereo pair, with a contrario elimination of outliers. *Image Processing On Line*, 5:89–113, 2016. 3
- David Nistér. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):756–770, 2004. 3
- Liangzu Peng, Mahyar Fazlyab, and René Vidal. Semidefinite relaxations of truncated Least-Squares in robust rotation search: Tight or not. In *Computer Vision – ECCV 2022*, pp. 673–691. Springer Nature Switzerland, 2022a. 3, 5, 15
- Liangzu Peng, Christian Kümmerle, and René Vidal. Global linear and local superlinear convergence of IRLS for non-smooth robust regression. In *Advances in Neural Information Processing Systems*, 2022b. 5
- Liangzu Peng, Manolis C. Tsakiris, and René Vidal. ARCS: Accurate rotation and correspondences search. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022c. 15
- Qing Qu, Ju Sun, and John Wright. *Finding a sparse vector in a subspace: Linear sparsity using alternating directions*, pp. 3401–3409. 2014. 1

- Qing Qu, Xiao Li, and Zhihui Zhu. A nonconvex approach for exact and efficient multichannel sparse blind deconvolution. August 2019. 1
- Qing Qu, Zhihui Zhu, Xiao Li, Manolis C Tsakiris, John Wright, and René Vidal. Finding the sparsest vectors in a subspace: Theory, algorithms, and applications. January 2020. 1
- Rahul Raguram, Ondrej Chum, Marc Pollefeys, Jiri Matas, and Jan Michael Frahm. USAC: A universal framework for random sample consensus. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(8):2022–2038, 2013. 3
- H Spath and G A Watson. On orthogonal linear  $\ell_1$  approximation. *Numerische Mathematik*, 51:531–543, 1987. 1
- H Späth and GA Watson. On orthogonal linear  $\ell_1$  approximation. *Numerische Mathematik*, 51:531–543, 1987. 9
- C V Stewart. MINPRAN: a new robust estimator for computer vision. *IEEE Trans. Pattern Anal. Mach. Intell.*, 17(10):925–938, October 1995. 3
- Ju Sun, Qing Qu, and John Wright. Complete dictionary recovery over the sphere I: Overview and the geometric picture. *IEEE Trans. Inf. Theory*, 63(2):853–884, 2017a. 1
- Ju Sun, Qing Qu, and John Wright. Complete dictionary recovery over the sphere II: Recovery by riemannian Trust-Region method. *IEEE Trans. Inf. Theory*, 63(2):885–914, 2017b. 1
- P H S Torr and A Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Comput. Vis. Image Underst.*, 78(1):138–156, April 2000. 3
- Philip Hilaire Torr, Slawomir J Nasuto, and John Mark Bishop. Napsac: High noise, high dimensional robust estimation-it’s in the bag. In *British Machine Vision Conference (BMVC)*, 2002. 3
- Manolis C Tsakiris and Rene Vidal. Dual principal component pursuit. In *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*. IEEE, December 2015. 1
- Manolis C Tsakiris and René Vidal. Hyperplane clustering via dual principal component pursuit. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3472–3481. PMLR, 2017. 1
- Manolis C. Tsakiris and René Vidal. Dual principal component pursuit. *Journal of Machine Learning Research*, 19(18):1–50, 2018a. 5
- Manolis C Tsakiris and Rene Vidal. Dual principal component pursuit. *Journal of Machine Learning Research*, 19:1–49, 2018b. 1, 8
- Heng Yang and Luca Carlone. A quaternion-based certifiably optimal solution to the wahba problem with outliers. May 2019. 3, 5, 15, 16
- Heng Yang and Luca Carlone. Certifiably optimal Outlier-Robust geometric perception: Semidefinite relaxations and scalable global optimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, PP, May 2022. 3, 5
- Yunchen Yang, Xinyue Zhang, Tianjiao Ding, Daniel P Robinson, Rene Vidal, and Manolis C Tsakiris. Boosting RANSAC via dual principal component pursuit. October 2021. 3
- Ji Zhao. An efficient solution to Non-Minimal case essential matrix estimation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(4):1777–1792, April 2022. 4, 5, 8
- Ji Zhao, Wanting Xu, and Laurent Kneip. A certifiably globally optimal solution to generalized essential matrix estimation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2020. 5

Zhihui Zhu, Yifan Wang, Daniel Robinson, Daniel Naiman, Rene Vidal, and Manolis Tsakiris. Dual principal component pursuit: Improved analysis and efficient algorithms. In *Advances in Neural Information Processing Systems*, 2018. [1](#), [4](#), [6](#), [15](#), [19](#)

Zhihui Zhu, Tianyu Ding, Daniel Robinson, Manolis Tsakiris, and René Vidal. A linearly convergent method for non-smooth non-convex optimization on the grassmannian with applications to robust subspace and dictionary learning. In *Adv. Neural Inf. Process. Syst.*, volume 32. proceedings.neurips.cc, 2019. [4](#)