

Convex Relaxations of $SE(2)$ and $SE(3)$ for Visual Pose Estimation

Matanya B. Horowitz, Nikolai Matni, Joel W. Burdick

Abstract—This paper proposes a new method for rigid body pose estimation based on spectrahedral representations of the tautological orbitopes of $SE(2)$ and $SE(3)$. The approach can use dense point cloud data from stereo vision or an RGB-D sensor (such as the Microsoft Kinect), as well as visual appearance data as input. The method is a convex relaxation of the classical pose estimation problem, and is based on explicit linear matrix inequality (LMI) representations for the convex hulls of $SE(2)$ and $SE(3)$. Given these representations, the relaxed pose estimation problem can be framed as a robust least squares problem with the optimization variable constrained to these convex sets. Although this formulation is a *relaxation* of the original problem, numerical experiments indicates that it is indeed exact – i.e. its solution is a member of $SE(2)$ or $SE(3)$ – in many interesting settings. We additionally show that this method is *guaranteed* to be exact for a large class of pose estimation problems.

I. INTRODUCTION

The pose estimation problem in computer vision is to identify a transformation that, when applied to a known object model, yields the object as perceived through the system’s sensors. Pose estimation is fundamental in robotics as it allows systems to reason about the state of the environment from their sensory data. Applications arise in robotic manipulation [15], [18], [31] object tracking [24], and visual odometry [29], among others.

Many successful techniques have been proposed and validated. The typical procedure consists of identifying noteworthy features of the object model, and then identifying similar features from the sensed input. Given this correspondence, one of several methods estimates the transformation that aligns as many of these features as possible, for example [21]. Typically, potential mismatches in correspondence require the optimization to be repeated, giving rise to algorithms such as iterated closest point (ICP) [6]. These techniques are typically combined with probabilistic selection of features for matching using the RANSAC algorithm.

However, many existing methods for pose estimation are primarily local-search methods that perform iterative linearization of the coordinate transform from the object’s last known position via Levenberg-Marquardt [31], [20]. While such approximations are appropriate at high frame rates or slow object/robot movement, these methods require re-initialization if the object is not continually tracked successfully. These issues can be alleviated by techniques

that provide a strong prior on the object, such as those incorporating Extended Kalman Filters [24], [11], [27].

The problem of pose estimation is also often plagued by the presence of outliers in the sensory data stream, which may arise from multiple causes, such as erroneous visual depth discontinuities. Additionally, the construction of approximation algorithms further aggravates the problem of mis-correspondence. The common Best-Bin First (BBF) [4], used to rapidly compute approximate nearest neighbors, creates spurious correspondences in its quest for speed. Further, there is the tradeoff in feature-detection methods, for instance in the choice between SIFT [23] and SURF [3] features, with the former typically being more accurate and the latter more rapid. The system designer therefore faces a number of tradeoffs between robustness and speed, contributing to the presence of outliers and uncertainty [25], [19]. Indeed, it is the presence of such errors that limits the use of closed form solutions to the pose estimation problem in practice [16], [17].

Inspired by recent advances in convex algebraic geometry (reviewed in Section II), this paper takes a quite different approach to the pose estimation problem: pose estimation is framed as a (robust) least squares problem (Section III), where the optimization variable lies in the convex hull of either $SE(2)$ or $SE(3)$.

By relaxing the constraint on the optimization variable (i.e. the transformation matrix), and allowing it to lie in the convex hull of the Lie group, rather than the Lie group itself, the resulting optimization is a semidefinite program (SDP), and hence convex. Formulating the pose estimation problem as a convex optimization problem is appealing as there exist a number of off-the-shelf public and commercial solvers that allow for the practitioner to focus on the design of the system rather than the details of the numerical solution method. These solvers obtain second order convergence and may therefore be quite rapid in practice. This emphasis on a convex relaxation of the $SE(n)$ constraint has also been examined in [13], developed independently of the current work, where the authors relax the constraint of the special euclidean group by requiring the solution to lie in the convex hull of the orthogonal group $O(n)$.

Furthermore, if the relaxation can be shown to be exact, that is to say that the optimal solution to the problem is a member of $SE(2)$ or $SE(3)$, then the procedure is guaranteed to yield the *global optimum*, alleviating the need for any form of probabilistic sampling of the solution space. Indeed, numerical experiments show that under many standard settings the relaxation is exact, and our method yields the optimal solution to the pose estimation problem.

Matanya Horowitz, Nikolai Matni, and Joel Burdick are with the Department of Control and Dynamical Systems, Caltech, 1200 E California Blvd., Pasadena, CA. The corresponding author is available at mhorowit@caltech.edu

Matanya Horowitz is supported by a NSF Graduate Research Fellowship.

Encouraged by this empirical success, we also show that under mild assumptions a wide class of pose estimation problems are provably exact. In Section VI, we comment on both promising directions for expanding the class of problems for which this method is provably exact, and on heuristics for generating a transformation matrix given a solution that is not a member of $SE(2)$ or $SE(3)$.

In the spirit of recent results in machine learning, we also show that a robust variant of the problem may be framed as an ℓ_1 -regularized optimization problem [10] (Section IV). The ability to incorporate regularization techniques further distinguishes our method from existing solutions. Application of the methodology to a practical example is demonstrated in Section V, along with a comparison to other approaches. This example shows that our approach provides consistently better estimates across all ranges of noise, and is particularly advantageous as the sensor noise level increases.

II. BRIEF INTRODUCTION TO ORBITOPES

The focus of this paper is on estimating the poses of rigid objects in $SE(n)$ and $SO(n)$ for $n = 2, 3$. Our analysis does not carry directly to $n > 3$, and for the remainder of the paper we use $SE(n)$ to refer to only $n = 2, 3$. Our approach relies upon the concept of an *orbitope*, which is formally the convex hull of the orbits given by the action of a compact algebraic group G acting linearly on a real vector space. Orbits arising from such groups have the structure of a real algebraic variety and thus, the orbitope is a *convex semi-algebraic* set. These objects, whose study was formalized and initiated in [30], are at the heart of the emerging field of convex algebraic geometry.

In the case of finite groups G , the study of such objects has a rich history. The platonic solids, the permutahedra, the Birkhoff polytope and the traveling salesman polytopes are examples of such discretely generated orbitopes. These well studied objects have been analyzed in depth in the context of combinatorial optimization [28], [34], [26]. Orbitopes of compact Lie groups have proved useful in the analysis of protein structure prediction [22], quantum information theory [1] and real algebraic geometry (in [2], certain $SO(n)$ orbitopes are used to show that there are many more non-negative polynomials than sums of squares).

We will be interested in the convex representation of orbitopes generated by a group acting on its identity element, that is to say the group's *tautological orbitope*. In particular, our method is based on recent results on the *spectrahedral representation* of tautological orbitopes for the groups $SO(n)$, $n = 2, 3$.

A set is said to admit a spectrahedral representation if it can be described as the intersection of the cone of positive semidefinite matrices and an affine subspace, i.e. if it can be written in terms of a linear matrix inequality (LMI). It is possible to then optimize affine and convex quadratic functions over these sets using the well established methods of *semi-definite programming* (SDP). Additionally, if a set is second-order cone (SOC) representable – a restricted class of spectrahedral constraints – then SOC programming

techniques may be leveraged. These have the advantage of being significantly faster than their SDP counterparts.

A. The tautological orbitope for $SO(2)$

We begin by noting that a simple parameterization of proper rotations in 2-dimensional space is given by

$$SO(2) = \left\{ \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} : \theta \in [0, 2\pi) \right\}. \quad (1)$$

Letting $x = \cos \theta$ and $y = \sin \theta$, we can equivalently rewrite this expression as

$$SO(2) = \mathcal{L} \cdot \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 = 1\} \quad (2)$$

for a linear mapping $\mathcal{L} : \mathbb{R}^2 \rightarrow \mathbb{R}^{2 \times 2}$. In this way, $SO(2)$ is expressed as a linear map of a simple set, i.e. the unit sphere. This is key as the convex hull operator, which we denote by $\text{conv}(\cdot)$, commutes with linear mappings.

Next, this constraint is *relaxed* by taking its convex hull. The unit sphere becomes the unit disk, and thus the constraint is replaced with $x^2 + y^2 \leq 1$. While simple, the computational gains realized by this step are significant.

$$\begin{aligned} \text{conv}(SO(2)) &= \left\{ \begin{bmatrix} x & y \\ -y & x \end{bmatrix} : x^2 + y^2 \leq 1 \right\} \\ &= \left\{ \begin{bmatrix} x & y \\ -y & x \end{bmatrix} : \begin{bmatrix} I & \begin{pmatrix} x \\ y \end{pmatrix} \\ (x \ y) & 1 \end{bmatrix} \succcurlyeq 0 \right\} \end{aligned} \quad (3)$$

where the last equality follows by applying the Schur complement [8] to the SOC constraint $x^2 + y^2 \leq 1$. The latter constraint in (3) is in fact an LMI, and the unit disk is therefore said to be *spectrahedrally representable*, and thus tractable to optimize over.

Remark 1. *The orbitope of $\text{conv}(SO(2))$ is SOC representable. These objects have significant computational advantages over general semidefinite programs and optimization over hundreds of variables may be done on time scales of several milliseconds [12].*

B. The tautological orbitope of $SO(3)$

An explicit representation of the tautological orbitope of $SO(3)$ is given by the following result from [30].

Proposition 2. *The tautological orbitope $\text{conv}(SO(3))$ is a spectrahedron whose boundary is a quartic hypersurface. In fact, a 3×3 -matrix $X = (x_{ij})$ lies in $\text{conv}(SO(3))$ if and only if it satisfies (4).*

Proof. The following is a slightly modified version of the original derivation suggested by Pablo Parrilo to the authors of [30], included here for completeness. An explicit parameterization of $SO(3)$ is given by its embedding into the space of pure quaternions (a subgroup of $SU(2)$) as

$$\begin{aligned} SO(3) &= \{U \in \mathbb{R}^{3 \times 3} \mid \quad (5) \\ U(u) &= \begin{bmatrix} 2(u_0^2 + u_1^2) - 1 & 2(u_1 u_2 - u_0 u_3) & 2(u_1 u_3 + u_0 u_2) \\ 2(u_1 u_2 + u_0 u_3) & 2(u_0^2 + u_2^2) - 1 & 2(u_2 u_3 - u_0 u_1) \\ 2(u_1 u_3 - u_0 u_2) & 2(u_2 u_3 + u_0 u_1) & 2(u_0^2 + u_3^2) - 1 \end{bmatrix} \\ u &\triangleq (u_0 \ u_1 \ u_2 \ u_3), \quad \|u\|_2^2 = 1 \end{aligned}$$

$$\begin{pmatrix} 1 + x_{11} + x_{22} + x_{33} & x_{32} - x_{23} & x_{13} - x_{31} & x_{21} - x_{12} \\ * & 1 + x_{11} - x_{22} - x_{33} & x_{21} + x_{12} & x_{13} + x_{31} \\ * & * & 1 - x_{11} + x_{22} - x_{33} & x_{32} + x_{23} \\ * & * & * & 1 - x_{11} - x_{22} + x_{33} \end{pmatrix} \succeq 0 \quad (4)$$

Fig. 1. Spectrahedral representation of $\text{conv}(SO(3))$ [30]. Omitted * elements indicate the symmetric completion of the matrix.

Noting that each term in $U(u)$ is quadratic in elements of u , we may define an auxiliary matrix

$$V := \begin{bmatrix} u_0^2 & u_0 u_1 & u_0 u_2 & u_0 u_3 \\ u_1 u_0 & u_1^2 & u_1 u_2 & u_1 u_3 \\ u_2 u_0 & u_2 u_1 & u_2^2 & u_2 u_3 \\ u_3 u_0 & u_3 u_1 & u_3 u_2 & u_3^2 \end{bmatrix} = uu^T. \quad (6)$$

We therefore see that if u is such that $U(u) \in SO(3)$, then V is a positive semi-definite (which we denote by $V \succcurlyeq 0$) rank-1 symmetric matrix satisfying $\text{trace}(V) = \|u\|_2^2 = 1$, and that there exists an invertible affine mapping \mathcal{A} such that $U(u) = \mathcal{A}(V)$. Equivalently, we may define the set

$$\mathcal{V} := \{V \succcurlyeq 0 : \text{rank}V = 1, \text{trace}V = 1\} \subset \mathbb{R}^{4 \times 4}$$

and express $SO(3)$ as

$$SO(3) = \mathcal{A} \cdot \mathcal{V} \quad (7)$$

for the appropriately chosen affine map $\mathcal{A} : \mathbb{R}^{4 \times 4} \rightarrow \mathbb{R}^{3 \times 3}$. Taking the convex hull of (7) we obtain

$$\begin{aligned} \text{conv}(SO(3)) &= \mathcal{A} \cdot \text{conv}\mathcal{V} \\ &= \mathcal{A} \cdot \text{conv}\{V \succcurlyeq 0 \mid \text{rank}V = 1, \text{trace}V = 1\} \\ &= \mathcal{A} \cdot \{V \succcurlyeq 0 \mid \text{trace}V = 1\}. \end{aligned} \quad (8)$$

It then suffices to invert \mathcal{A} to obtain the spectrahedral representation for $\text{conv}(SO(3))$ as given in (4). \square

C. The tautological orbitope for $SE(n)$

A spectrahedral representation of $SE(n)$ follows naturally from one for $SO(n)$. The space of the translations is simply \mathbb{R}^n , which of course is already a convex set. Thus, noting that when $SE(n)$ is expressed in homogeneous coordinates, it is none other than a linear mapping of elements of $SO(n) \times \mathbb{R}^n$, we are free to take convex hulls of the individual components of the transformation matrix. A parameterization of $S \in \text{conv}(SE(n))$ in homogeneous coordinates is therefore given by

$$S = \begin{bmatrix} R & T \\ 0 & 1 \end{bmatrix} \quad (9)$$

where $T \in \mathbb{R}^n$ and $R \in \text{conv}(SO(n))$.

III. LEAST SQUARES ESTIMATION FORMULATION

We assume we are given a model of an object whose pose we wish to estimate based on sensor data. This model consists of a vector $(m_i)_{i=1}^N$ of N features of a model with coordinates in \mathbb{R}^n for $n = 2, 3$ in the object's body-fixed reference frame. We are then given an observation that consists of a vector $(o_i)_{i=1}^N$ of the corresponding feature coordinates in the observation frame. The observation frame itself may

consist of a projection P to a lower, two dimensional image frame as is the case for visual data, or the data may remain in three dimensions. In either case, we wish to solve for the object pose that is most consistent with this observation.

The problem is general and, as framed, comprises a central component of many pose estimation problems in robotics. All pose estimation problems consists of matching known features on a model to observed features, whether these observations arise from an image, point cloud, or other sources.

Classical Least Squares Estimation. The problem of data fitting is a common one, with the most popular approach being that least-squares regression. In our case, the data and the model are related by a transformation element of $SE(n)$, with $n \in \{2, 3\}$. For $n = 2$, this leads to the optimization problem

$$\begin{aligned} \min & \sum_{i=1}^N \|o_i - Sm_i\|_2^2 \\ \text{s.t.} & S \in SE(2) \end{aligned} \quad (10)$$

The three-dimensional visual feature problem is similar but must also incorporate a projection, P , onto the camera frame. The general form of the least squares optimization problem is then

$$\begin{aligned} \min_S & \sum_{i=1}^N \|o_i - PSm_i\|_2^2 \\ \text{s.t.} & S \in SE(n) \end{aligned} \quad (11)$$

where $n = 2, 3$, and P may be set to the identity, or the camera matrix, as appropriate. Problems (10) and (11) are non-convex optimization problems, and typically difficult to solve. Beyond the issue of only finding local minima, the optimization can be slow because $SE(n)$ is difficult to represent in a manner that is amenable to optimization, with parametrizations typically incorporating trigonometric functions.

Convex Least Squares Estimation over $SE(n)$. The difficulties presented by parameterizations of $SE(n)$ in the least squares problem can be alleviated by relaxing the problem, and replacing the Lie group constraints with their respective orbitope constraints. In particular, by using the spectrahedral representations of the orbitopes of $SE(n)$ as discussed, we may write:

$$\begin{aligned} \min_S & \sum_{i=1}^N \|o_i - PSm_i\|_2^2 \\ \text{s.t.} & S \in \text{conv}(SE(n)) \end{aligned} \quad (12)$$

for $n = 2, 3$, where $\text{conv}(SE(n))$ now admits a semidefinite representation, as in equations (3), (4) and (9).

To complete the formulation, the problem data is represented in a vector form: let $O \triangleq [o_1, \dots, o_N]^T$, $M \triangleq [m_1, \dots, m_N]^T$, $\bar{P} = I_{N \times N} \otimes P$, and $\bar{S} = I_{N \times N} \otimes S$, where \otimes denotes the Kronecker Product, allowing us to rewrite optimization (12) as:

$$\begin{aligned} \min_{\bar{S}} \quad & \|O - \bar{P}\bar{S}M\|_2^2 \\ \text{s.t.} \quad & \bar{S} \in \text{conv}(I \otimes SE(3)) \end{aligned} \quad (13)$$

Remark 3. *It is possible to include additional feature weighting information in the optimization problem (13) by simply weighting $\bar{P} = C \otimes P$, and $\bar{S} = C \otimes S$ for C the diagonal matrix of weights $\{c_1, c_2, \dots, c_N\}$. This may arise, for instance, in SIFT feature matching wherein there is a score that is calculated for each feature pair, but any similar heuristic would be admissible.*

In order to formulate the pose estimation problem as a SDP, we have relaxed the constraint on the transformation matrix – in particular, rather than searching over $SE(n)$, we now optimize over its convex hull. Thus, in general, it is not true that the solution to (13) will be an admissible transformation, as Euclidean transformations are not closed under convex combinations. Thus an important, albeit obvious, question arises: when can we guarantee that the solution of our *convex relaxation* (13) is in fact an element of $SE(n)$.

A. Guarantee of Boundary Solution for $SO(n)$

Theorem 4 (Sufficient conditions for exactness). *Given $n = 2, 3$ $O, M \in \mathbb{R}^{4N}$, and $t \in \mathbb{R}^n$, if $P \in O(n+1)$, with $O(n+1)$ the orthogonal group, then the solution R to the optimization problem*

$$\begin{aligned} \min_R \quad & \|O - \bar{P}\bar{S}M\|_2^2 \\ \text{s.t.} \quad & \bar{S} = I_{N \times N} \otimes S \\ & S = \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix} \\ & R \in \text{conv}(SO(n)) \end{aligned} \quad (14)$$

lies in $SO(n)$.

Proof. Consider the non-convex optimization:

$$\begin{aligned} \min_R \quad & \|O - \bar{P}\bar{S}M\|_2^2 \\ \text{s.t.} \quad & \bar{S} = I_{N \times N} \otimes S \\ & S = \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix} \\ & R \in SO(n) \end{aligned} \quad (15)$$

Our approach will be to show that this optimization has an equivalent convex reformulation given by (14). The optimization is examined in the equivalent summation form (11). We first rewrite the objective so as to isolate the rotational component R :

$$\|o_i - PSm_i\|_2^2 = \|o_i - P\tau - P\rho m_i\|_2^2$$

where

$$\rho = \begin{bmatrix} R & 0 \\ 0 & 1 \end{bmatrix}$$

and $\tau = [t^T, 1]^T$. It is well known that the optimal value for t is the translation that aligns the centroids of M and O [16] and it may therefore be calculated a-priori.

Recalling that for a matrix $T \in O(n+1)$, we have that $\|Tx\|_2 = \|x\|_2$, and letting $v_i = o_i - P\tau$, we may write for each $i \in \{1, \dots, N\}$,

$$\begin{aligned} \|v_i - P\rho m_i\|_2^2 &= \langle v_i - P\rho m_i, v_i - P\rho m_i \rangle \\ &= \|v_i\|_2^2 - 2 \langle v_i, P\rho m_i \rangle \\ &\quad + \langle P\rho m_i, P\rho m_i \rangle \\ &= \|v_i\|_2^2 - 2 \langle v_i, P\rho m_i \rangle + \|P\rho m_i\|_2^2 \\ &= \|v_i\|_2^2 - 2 \langle v_i, P\rho m_i \rangle + \|m_i\|_2^2, \end{aligned}$$

where the last equality follows from $P\rho \in O(n+1)$. As constant terms in the objective do not affect the optimal solution, optimization (15) can be rewritten in terms of the *linear* objective function $-2 \langle O, \bar{P}\bar{S}M \rangle$:

$$\begin{aligned} \min_R \quad & -2 \langle O, \bar{P}\bar{S}M \rangle \\ \text{s.t.} \quad & \bar{S} = I_{N \times N} \otimes S \\ & S = \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix} \\ & R \in SO(n) \end{aligned} \quad (16)$$

It is a standard result of convex optimization that linear objective functions attain their minima at extreme points of the feasible set [8] – thus there is no loss in replacing the constraint $R \in SO(n)$ with $R \in \text{conv}(SO(n))$ in (16), as the set of extreme points of $\text{conv}(SO(n))$ is precisely $SO(n)$. Rewriting the objective in its quadratic, rather than linear, form then yields (14), concluding the proof. \square

B. Projection of Interior Solutions

Although we have shown that the least-squares estimation over the Special Euclidean group will always produce border solutions, we will subsequently modify the optimization and lose this guarantee. In the case where the optimal solution is not an element of $SE(n)$, a practical algorithm will require a method of quickly generating a “nearby” Euclidean transformation. As a potential heuristic, we propose using a projection of the inadmissible transformation onto the admissible set. When this projection is taken with respect to the Frobenius norm, a simple singular value decomposition (SVD) based solution exists [5].

In particular, let a rotation $S \in \text{conv}(SO(n))$ have the SVD

$$S = U\Sigma V^T. \quad (17)$$

Then $S' = UV^T$ is the projection of S onto $SO(n)$, i.e. $S' = \arg \min\{\|T - S\|_{\text{Frobenius}} : T \in SO(n)\}$. Future work will look to quantify the performance of such heuristics.

IV. ROBUST FORMULATION

It is well known that the least-squares methodology is sensitive to the presence of outliers. The orbitope optimization approach can be modified to be more resilient to outliers by framing it as a robust least squares problem.

An advantage of using convex optimization lies in its flexibility. The optimization problem may be augmented with various features to improve performance beyond that obtainable by simple least squares estimation. Typical of vision problems are the existence of spurious outliers that distort the data, causing the estimators to diverge significantly [24]. It is therefore desirable that we remove the effects of such outliers automatically.

Suppose that we wish to estimate the parameters β such that $\|o - \beta m\|_2^2$ is minimized for data o, m . We may augment the traditional ℓ_2 objective with an ℓ_1 penalty term, or *regularizer*, as

$$\min_{\beta} \|o - \beta m\|_2^2 + \lambda \|\beta\|_1$$

where $\lambda > 0$ is a tuning parameter, resulting in the so-called LASSO least squares problem [32]. Interestingly, the effect of this regularizer is to drive many of the coefficients of β to zero, resulting in a parameter vector that is sparse. There are many applications where it is reasonable to expect that the true solutions are in fact sparse. This is the driving force behind the field of *compressed sensing* [9]. In the pose estimation problem it is reasonable to expect that there will exist outliers, but such outliers will be sparse in the data, i.e. we expect most correspondences to be accurate.

Inspired by these methods, we present an ℓ_1 -penalty based modification of optimization (13). In particular, we wish to minimize $\|O - \bar{P}\bar{S}M\|_2^2$ while having a solution that is robust to outliers, i.e. X contains some elements we wish to ignore. We separate the error into two components: an intrinsic error Z_2 due to sensor noise, and an error Z_1 due to incorrect correspondence (or outliers):

$$O - \bar{P}\bar{S}M = Z_1 + Z_2$$

The regularized problem then becomes

$$\begin{aligned} \min_{\bar{S}, Z_1, Z_2} \quad & \|Z_2\|_2^2 + \lambda \|Z_1\|_1 \\ \text{s.t.} \quad & \bar{S} \in \text{conv}(I \otimes SE(n)) \\ & O - \bar{P}\bar{S}M = Z_1 + Z_2, \end{aligned}$$

or, after solving for $Z_2 = (O - \bar{P}\bar{S}M) - Z_1$,

$$\begin{aligned} \min_{\bar{S}, Z_1} \quad & \|(O - \bar{P}\bar{S}M) - Z_1\|_2^2 + \lambda \|Z_1\|_1 \\ \text{s.t.} \quad & \bar{S} \in \text{conv}(I \otimes SE(3)) \end{aligned} \quad (18)$$

for some user determined $\lambda > 0$. Thus in principle, Z_1 will act as a correction term in the least squares component of the objective function, removing the effect of outliers – by adding the ℓ_1 penalty and tuning parameter λ , we control the sparsity and norm of Z_1 , and thus ensure that it does not overwhelm the correct data. Unfortunately, the lack of a linear objective causes the optimization program to lose the

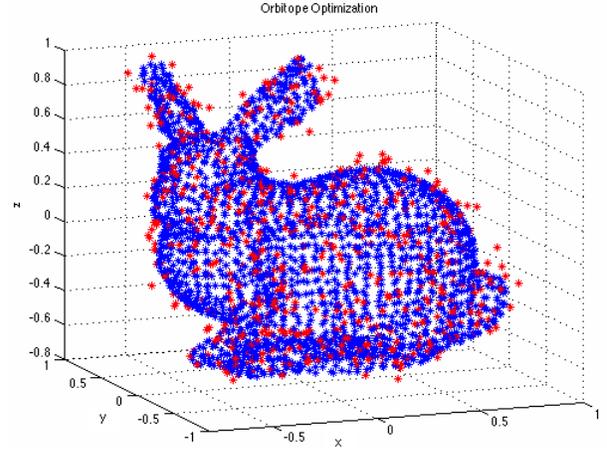


Fig. 2. Pose estimation for a corrupted Stanford Bunny mesh. Blue points are the optimized model points while red are the observation points. Gaussian noise with zero mean and standard deviation of $\delta = .06$ was used for the visualization.

guarantee of being on the boundary of $\text{conv}(SO(n))$, as was the case without the regularization.

Remark 5. Numerical experiments suggest that this method also typically gives solutions that are in $SE(n)$. As mentioned in Section III-B, we are exploring possible projection based heuristics for when our relaxation fails to be exact.

V. EXAMPLES

A benchmark pose estimation problem was used to compare the proposed method against some classical pose estimation procedures. To perform the optimization portion of our algorithm, we used the CVX parser [14] in conjunction with the publicly available SDPT3 semidefinite program solver [33].

A. Three Dimensional Pose Estimation

The method was tested on a pose estimation problem involving three dimensional point-cloud data. As a ground truth model, we used the Stanford Bunny from the Stanford Computer Graphics Laboratory, consisting of a mesh of 944 points. The model was normalized so that the bunny had maximal extent from the origin of one unit in the y -coordinate from the origin. The centroid of the model was placed at the origin, as all methods may be centered in this manner. To synthesize measurements, the model points were corrupted with varying amounts of Gaussian noise. This corruption had zero mean, standard deviation δ , and was applied to the (x, y, z) components of each model point individually. The correspondence between the points was maintained in order to negate the effects of this component of the optimization problem. A visualization of a typical result is shown in Figure 2 for a particular noise covariance value. Performance results with varying numbers of corrupted samples are shown in Figure 3. The errors reported were the sum of the squares of the individual error distances between corresponding points, where the error was taken with respect to the true, underlying model

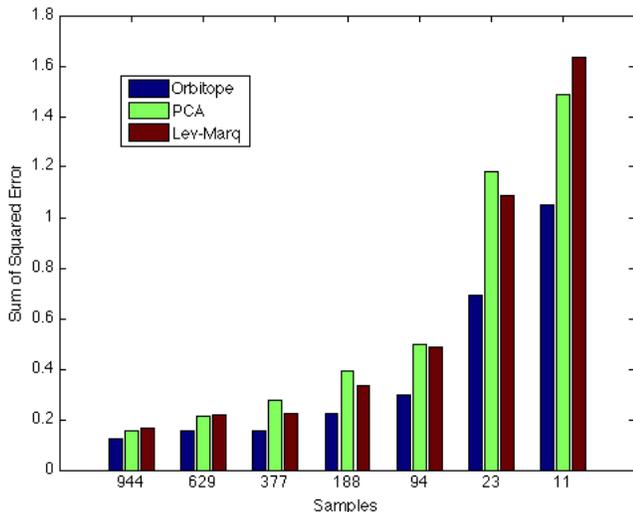


Fig. 3. Mean error over twenty trials when using Orbitope, PCA, and Levenberg-Marquardt optimizations in blue, green, and red respectively with varying numbers of randomly sampled observed points. Noise with standard deviation $\delta = 0.01$ was used.

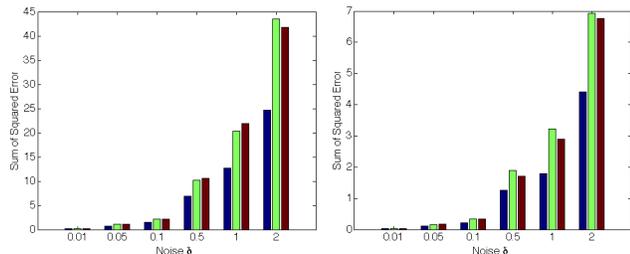


Fig. 4. Mean error over twenty trials when using Orbitope, PCA, and Levenberg-Marquardt optimizations in blue, green, and red respectively when varying noise standard deviation. Number of samples $N = 23$ and $N = 944$ were used in the left figures respectively.

and not the observed points. We compared our method to an implementation using Levenberg-Marquardt, as well as an implementation that aligns the principal components of the observed data to the model using Principle Component Analysis (PCA). These are both estimates to find a transformation to match a known model to observed points, given the correspondence between these two groups. These three methods would typically be used as part of a single iteration of ICP (in particular, usually PCA), where after the solution is computed the correspondence problem would again be solved, and the process repeated. The methods were also tested on varying levels of noise, with the results shown in Figure 4.

The examples were run on a 1.7Ghz Macbook Air. The Levenberg-Marquardt algorithm typically required 25ms for execution; aligning the principal components of the data typically required 2ms; and the orbitope optimization lasted 130ms. For the range of sample points used, the execution times did not vary significantly. All methods also did not change their execution time significantly with respect to the chosen noise levels. As the correspondence was perfect, all method performed well. However, the method proposed here

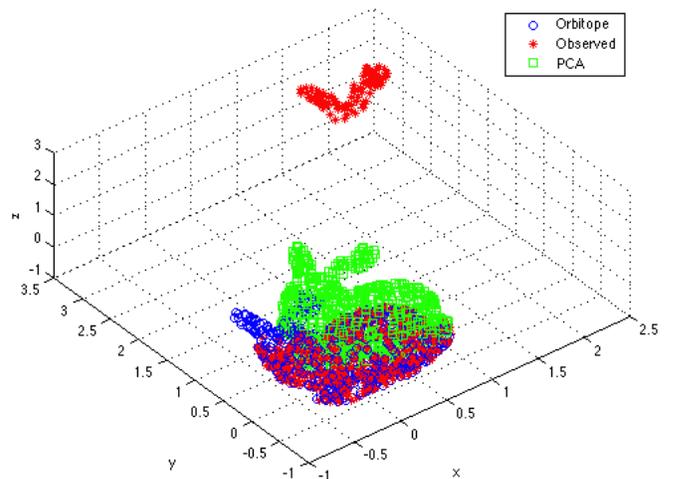


Fig. 5. Visualization of results for robust estimation problem. Noise of $\delta = 0.01$ is added to the normalized model, and the ears consisting of points with component $y \geq 0.6$ are translated by $t = (2, 2, 2)$ to create artificial outliers. The detected data is shown in red, while the matched model using the robust optimization is shown in blue, while the PCA match is shown in green.

provided superior solutions not only in aggregate but in the individual trials as well. For instance, in the 20 trials for which the noise was held at $\delta = 0.1$, and $N = 944$ points, the orbitope optimization error varied between 95% and 23% of the Levenberg-Marquardt error.

Next we tested the outlier rejection capabilities of the method by introducing artificial outliers into the model. The number of outlier points was relatively sparse compared to the total number of model points, and they were correctly detected as outliers by the algorithm for $\lambda = 0.1$. The results are visualized in Figure 5. The error residual for the robust method was $e = 0.7825$ while that of PCA was $e = 25.3098$. For these noise levels, the error of the orbitope optimization is typical (Figure 4) for the case where there are no outliers, indicating the shifting of the ears has only marginal impact on the solution. Indeed, it is possible to make this gap arbitrarily large by moving the ears further. Such situations are common in practice where, for example, collections of points may be erroneously associated with a model, or points of a model may not be visible due to occlusion.

VI. DISCUSSION

This paper introduced a substantially new formulation of the pose estimation problem based on spectrahedral representations of the convex hulls of $SO(n)$ and $SE(n)$, for $n \in \{2, 3\}$. This reformulation allowed for a convex relaxation of the pose estimation problem, resulting in a convex least squares problem. We also introduced an ℓ_1 penalized variant of our problem so as to automatically remove outliers.

We proved that under certain conditions, the solution is guaranteed to be a proper rigid body transformation, and is therefore a *globally optimal* solution to the pose estimation problem. As mentioned, the investigation of more general conditions under which boundary solutions are guaranteed is an active area of current research, as well as the quality

of projection based heuristics for when our relaxation fails to be exact.

Our numerical experiments showed that this new estimation method consistently produced more accurate estimates, particularly at high noise levels. Recent work [10] in the context of linear inverse problems has resulted in error bounds on regularized estimation problems. The application of these methods to the pose estimation problem may yield quantifiable bounds on estimation error, and an insight as to the types of disturbances (e.g., Gaussian vs. Poisson noise) which are best handled by the method.

The method proposed here was an order of magnitude slower than existing methods. There are a number of ways to greatly increase its speed. Our implementation relied on using CVX in conjunction with SDPT3: these are general solvers and parsers, and there is overhead accumulated. In a deployed implementation the solver would be coded directly in a language such as C++ with no parsing. Additionally, such least squares formulations are amenable to alternating direction method of multipliers (ADMM) solutions, allowing for parallelization of the optimization procedure [7]. Finally, in many applications it is appropriate to use a well-informed initial guess for the solver, such as the object's last known pose. Such warm-starts are not necessary for our optimization method, but can further improve execution time.

A. Acknowledgements

The authors would like to thank Venkat Chandresakaran for discussion on the topics present in this paper. The first author is grateful for the support of a National Science Foundation graduate fellowship. This work was partially supported by DARPA under the ARM-S and DRC programs.

REFERENCES

- [1] D. Avis, H. Imai, and T. Ito. On the relationship between convex bodies related to correlation experiments with dichotomic observables. *Journal of Physics A: Mathematical and General*, 39(36):11283, 2006.
- [2] A. Barvinok and G. Blekherman. Convex geometry of orbits. *Combinatorial and computational geometry*, 52:51–77, 2005.
- [3] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool. Speeded-up robust features (SURF). *Computer vision and image understanding*, 110(3):346–359, 2008.
- [4] J. S. Beis and D. G. Lowe. Shape indexing using approximate nearest-neighbour search in high-dimensional spaces. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1000–1006, 1997.
- [5] C. Belta and V. Kumar. Euclidean metrics for motion generation on SE (3). *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 216(1):47–60, 2002.
- [6] P. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, Feb 1992.
- [7] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [8] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [9] E. J. Candès, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [10] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. The Convex Geometry of Linear Inverse Problems. *Foundations of Computational Mathematics*, 12(6):805–849, Oct. 2012.
- [11] T. H. Chung and J. Burdick. A Decision-Making Framework for Control Strategies in Probabilistic Search. In *IEEE International Conference on Robotics and Automation*, pages 4386–4393, 2007.
- [12] A. Domahidi, E. Chu, and S. Boyd. ECOS: An SOCP solver for embedded systems. In *European Control Conference (ECC)*, pages 3071–3076, 2013.
- [13] J. R. Forbes and A. H. J. de Ruiter. Linear-matrix-inequality-based solution to Wahbas problem. *Journal of guidance, control, and dynamics*, 2013. To Appear.
- [14] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.0 beta. <http://cvxr.com/cvx>, Sept. 2013.
- [15] P. Hebert, N. Hudson, J. Ma, and J. Burdick. Fusion of stereo vision, force-torque, and joint sensors for estimation of in-hand object location. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5935–5941, 2011.
- [16] B. K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4(4):629–642, Apr 1987.
- [17] B. K. P. Horn, H. M. Hilden, and S. Negahdaripour. Closed-form solution of absolute orientation using orthonormal matrices. *Journal of the Optical Society of America A*, 5(7):1127–1135, Jul 1988.
- [18] N. Hudson, T. Howard, J. Ma, A. Jain, M. Bajracharya, S. Myint, C. Kuo, L. Matthies, P. Backes, P. Hebert, et al. End-to-end dexterous manipulation with deliberate interactive estimation. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 2371–2378. IEEE, 2012.
- [19] N. Hudson, J. Ma, P. Hebert, A. Jain, M. Bajracharya, T. Allen, R. Sharan, M. B. Horowitz, C. Kuo, T. Howard, et al. Model-based autonomous system for performing dexterous, human-level manipulation tasks. *Autonomous Robots*, 36(1-2):31–49, 2014.
- [20] M. Krainin, P. Henry, X. Ren, and D. Fox. Manipulator and object tracking for in-hand 3d object modeling. *The International Journal of Robotics Research*, 30(11):1311–1327, 2011.
- [21] V. Lepetit, F. Moreno-Noguer, and P. Fua. EPnP: An accurate O(n) solution to the PnP problem. *International journal of computer vision*, 81(2):155–166, 2009.
- [22] M. Longinetti, L. Sgheri, and F. Sottile. Convex hulls of orbits and orientations of a moving protein domain. *Discrete & Computational Geometry*, 43(1):54–77, 2010.
- [23] D. G. Lowe. Distinctive Image Features from Scale-Invariant Key-points. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [24] J. Ma, T. H. Chung, and J. Burdick. A probabilistic framework for object search with 6-DOF pose estimation. *The International Journal of Robotics Research*, 30(10):1209–1228, Sept. 2011.
- [25] M. Martinez, A. Collet, and S. S. Srinivasa. MOPED: A scalable and low latency object recognition and pose estimation system. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2043–2049, 2010.
- [26] N. McCarthy, D. Ogilvie, N. Zobin, and V. Zobin. Convex geometry of coxeter-invariant polyhedra. *Contemporary Mathematics*, 321:153–180, 2003.
- [27] A. I. Mourikis, S. I. Roumeliotis, and J. W. Burdick. Sc-kf mobile robot localization: a stochastic cloning kalman filter for processing relative-state measurements. *IEEE Transactions on Robotics*, 23(4):717–730, 2007.
- [28] S. Onn. Geometry, complexity, and combinatorics of permutation polytopes. *Journal of Combinatorial Theory, Series A*, 64(1):31–49, 1993.
- [29] S. T. Pfister, K. L. Kriebbaum, S. I. Roumeliotis, and J. W. Burdick. Weighted range sensor matching algorithms for mobile robot displacement estimation. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 1667–1674, 2002.
- [30] R. Sanyal, F. Sottile, and B. Sturmfels. Orbitopes. *arXiv preprint arXiv:0911.5436*, 2009.
- [31] S. S. Srinivasa, D. Ferguson, C. J. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and M. V. Weghe. HERB: a home exploring robotic butler. *Autonomous Robots*, 28(1):5–20, 2010.
- [32] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, pages 267–288, 1996.
- [33] K. C. Toh, M. Todd, and R. H. Tunc. SDPT3 – a MATLAB software package for semidefinite programming. *Optimization methods and software*, 11:545–581, 1999.
- [34] G. M. Ziegler. *Lectures on polytopes (graduate texts in mathematics)*. Springer, 1994.