

A Certifiable Algorithm for Simultaneous Shape Estimation and Object Tracking

Lorenzo Shaikewitz, Samuel Ubellacker, Luca Carlone
 Laboratory for Information & Decision Systems (LIDS)
 Massachusetts Institute of Technology
 Cambridge, USA
 Email: {lorenzos, subella, lcarlone}@mit.edu

Abstract—Applications from manipulation to autonomous vehicles rely on robust and general object tracking to safely perform tasks in dynamic environments. We propose the first certifiably optimal category-level approach for simultaneous shape estimation and pose tracking of an object of known category (e.g., a car). Our approach uses 3D semantic keypoint measurements extracted from an RGB-D image sequence, and phrases the estimation as a fixed-lag smoothing problem. Temporal constraints enforce the object’s rigidity (fixed shape) and smooth motion according to a constant-twist motion model. The solutions to this problem are the estimates of the object’s state (poses, velocities) and shape (parameterized according to the *active shape model*) over the smoothing horizon. Our key contribution is to show that despite the non-convexity of the fixed-lag smoothing problem, we can solve it to *certifiable optimality* using a small-size semidefinite relaxation. We also present a fast outlier rejection scheme that filters out incorrect keypoint detections with shape and time compatibility tests, and wrap our certifiable solver in a graduated non-convexity scheme. We evaluate the proposed approach on synthetic and real data, showcasing its performance in a table-top manipulation scenario and a drone-based vehicle tracking application.

I. INTRODUCTION

Target tracking is a well-studied problem across multiple research communities, including robotics, computer vision, and aerospace. Early work models the target as a point mass and is concerned with estimating its location from measurements (e.g., bearing vectors) while resolving data association, e.g., [1]. In modern robotics applications, robots navigate in close proximity and possibly interact with nearby objects. For those applications, the robot also needs to estimate the object’s shape. This leads to the coupled problem of *shape estimation and pose tracking*, which is crucial for autonomous vehicles [2], table-top manipulation [3], monitoring and surveillance [4], among other applications.

A significant body of work is dedicated to object pose and shape estimation in the *single-frame case* [5–8]. However, using a single image for detection and estimation sacrifices temporal information which is readily available. A good tracking algorithm leverages temporal consistency to operate in noisy, occlusion-rich, or highly dynamic environments. When the object shape is known exactly, *instance-level* tracking algorithms leverage known geometry [4, 9] or extensive

This work was partially supported by the AFOSR “Certifiable and Self-Supervised Category-Level Tracking” program, Carlone’s NSF CAREER award, the ONR RAPID program, and by MathWorks.




Fig. 1. **Overview of CAST#.** We estimate the shape and track the pose of an object from a sequence of images picturing the object. Given 3D keypoint measurements obtained via a learning-based detector, we formulate a non-convex fixed-lag smoothing problem where the shape is parametrized using an active shape model and motion smoothness is enforced using a constant-twist motion model. We solve this problem via a tight and small-size semidefinite relaxation and wrap the method in an outlier rejection scheme to robustly estimate shape and pose over a fixed time horizon.

training on the specific object instance [3, 10, 11]. However, in practical applications the object is rarely known exactly. Autonomous vehicles must track any of the thousands of car models in their surroundings; even if all were to be cataloged, the vehicle must still reason over intra-category variations.

Recent work has made considerable progress towards *category-level* and *category-free* object tracking using end-to-end learning [12, 13] or learning-based segmentation combined with local solvers for estimation [14–16]. While often effective in practice, learning-based approaches have unpredictable failure modes when used outside their training domain, while local-solver-based approaches are prone to converge to local minima corresponding to poor estimates. For safety critical applications, it is important to have interpretable and predictable models that produce provably optimal estimates.

Contribution. We propose the first certifiably optimal algorithm for category-level shape estimation and pose tracking (Fig. 1). In our problem formulation (Section III), we consider as input the 3D position of keypoints detected in an RGB-D image sequence. These inputs are fed to a fixed-lag smoother that performs estimation of the states over a receding horizon, while enforcing smoothness of the motion using a constant-twist motion model. We parametrize the object’s shape using an *active shape model* [17, 18], which describes shape as a linear combination of 3D models from a library of representative models. Our first key contribution (Section IV) is to show that despite the non-convexity of the estimation problem, we can develop an empirically tight semidefinite relaxation that computes certifiably optimal object poses, velocities, and shape without the need for an initial guess. We name the resulting approach *CAST** (*Certifiable Algorithm for Shape estimation and Tracking*). Our second contribution (Section V) is to extend *CAST** to the practical case where some of the keypoint measurements are outliers; we handle this case by combining an outlier-pruning method with a robust estimator based on graduated non-convexity [19]. We call the resulting outlier-robust approach *CAST#*. Our final contribution (Section VI) is to evaluate our methods in extensive experiments, and show they produce accurate estimates when tested in simulation, on the YCBInEOAT dataset [3], and in a drone-based vehicle tracking scenario.

II. RELATED WORK

Instance-Level Object Tracking. Traditional target tracking approaches circumvent shape estimation by assuming the object to be a point mass [20] or assuming full knowledge of the object shape [3, 4, 9–11, 21, 22]. Early approaches used handcrafted features, such as points, edges [23], or planes [24] to compute relative poses. The set of pose estimates could then be smoothed via Kalman filtering [4, 21]. More recently, the use of handcrafted features has given way to learned features [3] and edge detection [10], and new approaches based on point cloud registration [9], particle filtering [11], or unscented Kalman filters [22] have emerged.

Category-Level and Category-Free Object Tracking. In practical settings, instance-level information is rarely available. Recent approaches investigate pose and shape estimation for objects within a known category [12–14, 16, 25] or at least similar enough to the training data [15, 26]. These approaches generally extract a sparse representation of the object to estimate relative motion between frames. Wang et al. [14] focus on an attention mechanism for extracting frame-to-frame keypoints in a self-supervised manner, leaving the work of relative pose estimation to point cloud registration, which is unable to use temporal information beyond two frames. Wen and Bekris [25] use a similar architecture but take a SLAM-inspired approach, using dense frame-to-frame feature correspondences and multi-frame pose graph optimization to refine the estimate. Other methods use learned keypoint correspondences for the Iterative Closest Point (ICP) method [13] or learning-based regression to estimate relative motion in the small pose regime [12]. Even

with keyframe selection [15], frame-to-frame back-ends require a separate tool to obtain object pose relative to a camera or world frame, which is often useful in applications. In contrast, we propose an optimization back-end that produces *certifiably optimal* shape and pose estimates from category-level keypoints without relying on local solvers. This gives useful world-frame poses directly and allows the use of a motion model to mitigate the impact of measurement noise.

Certifiable Algorithms. Our work extends the body of work on *certifiable perception algorithms*: a certifiable algorithm solves an optimization problem and either provides a certificate of optimality or a bound on the suboptimality of the produced solution [9]. Certifiable algorithms are typically derived using semidefinite relaxations, and are usually based on Shor’s relaxation of Quadratically Constrained Quadratic Programs (QCQP) or Lasserre’s relaxation of polynomial optimization problems [27–29]. Certifiable algorithms have been proposed for rotation averaging [30, 31], pose graph optimization [32–34], 3D registration [9, 35], 2-view geometry [36–38], perspective-n-point problems [39], and single-frame pose and shape estimation [5]. Recent work has extended certifiable solvers to cope with outliers [27] and anisotropic noise [40]. Our approach extends [5] to tracking over a receding time horizon using a motion model.

III. PROBLEM FORMULATION

This section formalizes the *category-level shape estimation and pose tracking* problem. Given a sequence of RGB-D images picturing an object of known category (*e.g.*, a car), and assuming the availability of a 3D semantic keypoint detector, we seek an estimate of the time-independent shape and time-dependent pose (*i.e.*, position and orientation) of the object. Below we describe our choice of motion model, shape representation, and measurement model.

Object State and Motion Model. We represent the target object’s state using its pose and velocity at a particular time t . Denote the position and orientation of the target object in the world frame as $\mathbf{p}_t \in \mathbb{R}^3$ and $\mathbf{R}_t \in \text{SO}(3)$, respectively. Similarly, denote the target’s *body-frame* change in position and change in rotation between time steps with $\mathbf{v}_t \in \mathbb{R}^3$ and $\mathbf{\Omega}_t \in \text{SO}(3)$. These state variables are the discrete time analog to velocity and rotation rate. Any object’s motion obeys the following discrete-time first-order dynamics:

$$\mathbf{p}_{t+1} = \mathbf{p}_t + \mathbf{R}_t \mathbf{v}_t, \quad \mathbf{R}_{t+1} = \mathbf{R}_t \cdot \mathbf{\Omega}_t \quad (1)$$

The model (1) is quite general, since by choosing suitable values of $\mathbf{v}_t, \mathbf{\Omega}_t$ we can produce arbitrary trajectories.

Now, we assume that the velocities’ dynamics are approximately *constant twist*; *i.e.*, the body-frame velocity and the rotation rate are constant during short time intervals up to random perturbations $\mathbf{v}_t^\epsilon \in \mathbb{R}^3$ and $\mathbf{R}_t^\epsilon \in \text{SO}(3)$:

$$\mathbf{v}_{t+1} = \mathbf{v}_t + \mathbf{v}_t^\epsilon, \quad \mathbf{\Omega}_{t+1} = \mathbf{\Omega}_t \cdot \mathbf{R}_t^\epsilon \quad (2)$$

When $\mathbf{v}_t$ and $\mathbf{\Omega}_t$ are exactly constant the dynamical system in equations (1)–(2) models 3D spiral-shaped trajectories, including the corner cases of a straight line, circular trajectory, or in-place rotation. The random noise terms model




Fig. 2. **Active Shape Model.** Known 3D models in the *bottle* category and their averages computed according to the active shape model. Vertices are the original models and edges are the average of the two vertices. The active shape model can represent any 3D geometry in the convex hull of its shape library through a point-wise weighted average.

small deviations from these assumptions in the observed trajectory. The proposed constant-twist model is a 3D version of the popular constant-turn-rate model [21], generalizing it to allow an arbitrary axis of rotation. Such a model is expressive enough to capture the non-holonomic motion of a car and the unpredictable motions of a manipulated object.

In the following, we assume that the velocity noise follows an isotropic zero-mean Gaussian distribution: $\mathbf{v}_t^\epsilon \sim \mathcal{N}(0, \Sigma_v^\epsilon)$. Moreover, we assume that the relative rotation noise follows an isotropic Langevin distribution about the identity, following standard practice [32] for distributions over $\text{SO}(3)$: $\mathbf{R}_t^\epsilon \sim \mathcal{L}(\mathbf{I}_3, \kappa_t)$. In this equation, κ_t is the *concentration parameter* of the Langevin distribution (intuitively, this plays a similar role as the inverse of the variance).

Shape Parameterization We use the *active shape model* to represent intra-category shape variations. Given an object category (e.g., bottle), we assume a library of 3D models (e.g. specific bottle shapes) that span the category, where the objects in the library are denoted as $\mathcal{B}_k$, $k = 1, \dots, K$. Any instance, then, is just a pointwise linear combination of the models in the shape library (see Fig. 2). More formally, let $\mathbf{x}_i$ be any point on the instance object that corresponds to the point $\mathbf{b}_i^k \in \mathcal{B}_k$ in each library shape, $k = 1, \dots, K$. The active shape model is:

$$\mathbf{x}_i = \sum_{k=1}^K c_k \mathbf{b}_i^k \triangleq \mathbf{B}_i \mathbf{c} \quad (3)$$

where $c_k \in [0, 1]$ and $\sum_k c_k = 1$. Thus, the shape of the target object is fully specified by its shape coefficient $\mathbf{c} = [c_1, \dots, c_K]$ and the shape library for each point $\mathbf{B}_i = [\mathbf{b}_i^1, \dots, \mathbf{b}_i^K]$. This representation is simple and expressive: it captures any object in the convex hull of the shape library (including the library shapes themselves) via a linear combination described by a single vector of coefficients [17, 18]. Further, measurements of a small number of semantic *keypoints* are enough to resolve the object shape.

Measurement Model. The inputs to our estimator are measurements of the 3D positions of *semantic keypoints* on the target object. These keypoints correspond to semantically

meaningful features common to a specific object category, and are typically produced by a learning-based detector, as in [5, 41]. For instance, a set of keypoints on a bottle might be the locations of the cap, center-point of the base, label, etc. Such keypoints may be detected by a model trained on a category of bottles, not just a particular instance.

At each time t we are given the 3D position of N keypoints denoted $\mathbf{y}_t^1, \dots, \mathbf{y}_t^N$. These measurements obey the following generative model:

$$\mathbf{y}_t^i = \mathbf{R}_t \cdot (\mathbf{B}_i \mathbf{c}) + \mathbf{p}_t + \epsilon_t^i \quad (4)$$

Each measurement $\mathbf{y}_t^i$ is a rigid transformation $(\mathbf{R}_t, \mathbf{p}_t)$ of the keypoint's location in the object's frame $\mathbf{B}_i \mathbf{c}$ (expressed according to the active shape model) plus measurement noise ϵ_t . For now, we assume the measurement noise obeys an isotropic zero-mean Gaussian distribution: $\epsilon_t^i \sim \mathcal{N}(0, \Sigma_t^i)$.

Simultaneous Shape Estimation and Tracking. We are now ready to state the problem we tackle in this paper.

Problem 1 (Shape Estimation and Pose Tracking):

Consider an object of known category moving according to the dynamics in eqs. (1)-(2). Given measurements of N keypoints in the form (4) taken over T time steps, estimate the time-varying state $(\mathbf{R}_t, \mathbf{p}_t, \mathbf{v}_t, \Omega_t)$ and time-independent shape $\mathbf{c}$ of the object for $t = 1, \dots, T$.

Problem 1 may be interpreted as a fixed-lag smoother, where our primary goal is to estimate the state at time T using also the most recent $T - 1$ measurements.

IV. CAST*: CERTIFIABLE SHAPE ESTIMATION AND TRACKING IN THE OUTLIER-FREE SETTING

This section presents CAST*, a certifiably optimal estimator solving Problem 1 in the outlier-free setting. CAST* is also the basis for our outlier-robust extension in Section V.

We adopt a *maximum a posteriori* estimation framework that represents Problem 1 as an optimization problem. This framework minimizes the residual errors of the measurement and motion models over the time horizon T , possibly including priors. In our case, the only prior is that shape coefficients $\mathbf{c}$ are distributed according to a Gaussian with covariance $\frac{1}{\lambda} \mathbf{I}_3$ about the mean shape $\bar{\mathbf{c}} \triangleq \frac{1}{K} \mathbf{1}_K$. In practice, this prior regularizes the problem when the shape library is larger than the number of keypoints ($K > N$); see e.g. [5].

The maximum a posteriori estimator takes the form:

$$\begin{aligned} \min_{\substack{\mathbf{R}_t, \Omega_t \in \text{SO}(3), \\ \mathbf{p}_t, \mathbf{v}_t \in \mathbb{R}^3, \\ \mathbf{c} \in \mathbb{R}^K, \mathbf{1}_K^\top \mathbf{c} = 1 \\ t=1, \dots, T, \\ l=1, \dots, T-1}} \quad & \sum_{t=1}^T \sum_{i=1}^N w_t^i \|\mathbf{y}_t^i - \mathbf{R}_t \mathbf{B}_i \mathbf{c} - \mathbf{p}_t\|^2 + \lambda \|\Delta \mathbf{c}\|^2 \\ & + \sum_{t=1}^{T-1} \omega_t \|\mathbf{v}_{t+1} - \mathbf{v}_t\|^2 + \kappa_t \|\Omega_{t+1} - \Omega_t\|_F^2 \\ \text{s.t.} \quad & \mathbf{p}_{t+1} = \mathbf{p}_t + \mathbf{R}_t \mathbf{v}_t, \quad \mathbf{R}_{t+1} = \mathbf{R}_t \Omega_t \end{aligned} \quad (5)$$

In the previous expression, we used the shorthand $\Delta \mathbf{c} \triangleq \mathbf{c} - \bar{\mathbf{c}}$ and assumed isotropic covariances $\Sigma_t^i \triangleq \frac{1}{w_t^i} \mathbf{I}_3$ and $\Sigma_t^v \triangleq \frac{1}{\omega_t} \mathbf{I}_3$. We also relaxed the constraint $c_k \geq 0$. We observe that the objective is the sum of the shape prior with the negative log-likelihoods of the measurements (4) and dynamics (2). The constraints enforce the domains of the variables (e.g.

$\mathbf{R}_t \in \text{SO}(3)$ or $\mathbf{1}_K^\top \mathbf{c} = 1$) and the dynamics (1). We prove that (5) is a maximum a posteriori estimator in Appendix A.

Notice that (5) is non-convex due to the constraint set $\text{SO}(3)$ and the quadratic equality constraints. Thus, local search methods such as gradient descent or Gauss-Newton are prone to local minima that result in bad estimates.

In the following we present our approach to solving (5) to certifiable optimality via a semidefinite relaxation. In Section IV-A we simplify the problem by analytically solving for the optimal shape coefficient. Using a change of variables, we rewrite (5) as a non-convex Quadratically Constrained Quadratic Program (QCQP) in Section IV-B and apply a semidefinite relaxation in Section IV-C. This relaxed problem can be solved using traditional convex optimization techniques and is shown to be empirically *tight* (i.e., the relaxation solves (5) to optimality) in Section VI.

A. Closed-Form Solution for Shape

Observe that (5) is a linearly constrained convex quadratic program in the variable $\mathbf{c}$. Thus, we can solve for the optimal shape coefficient $\mathbf{c}^*$ in closed form as a function of the other unknown variables. We formalize this observation below.

Proposition 1 (Optimal Shape): For any choice of positions and rotations $(\mathbf{p}_t, \mathbf{R}_t)$, the optimal shape coefficient solving (5) is

$$\mathbf{c}^* = 2\mathbf{G} \left(\mathbf{B}^\top \sum_{t=1}^T \mathbf{W}_t \begin{bmatrix} \mathbf{R}_t^\top (\mathbf{y}_t^1 - \mathbf{p}_t) \\ \vdots \\ \mathbf{R}_t^\top (\mathbf{y}_t^N - \mathbf{p}_t) \end{bmatrix} + \lambda \bar{\mathbf{c}} \right) + \mathbf{g} \quad (6)$$

where we defined the following symbols:

$$\begin{aligned} \mathbf{W}_t &\triangleq \text{blkdiag}(w_1^t \mathbf{I}_3, \dots, w_N^t \mathbf{I}_3) && \in \mathbb{R}^{3N \times 3N} \\ \mathbf{B} &\triangleq [\mathbf{B}_1^\top, \dots, \mathbf{B}_T^\top]^\top && \in \mathbb{R}^{3N \times K} \\ \mathbf{H} &\triangleq \frac{1}{2} \left(\mathbf{B}^\top \left(\sum_{t=1}^T \mathbf{W}_t \right) \mathbf{B} + \lambda \mathbf{I}_K \right)^{-1} && \in \mathbb{R}^{K \times K} \\ \mathbf{G} &\triangleq \mathbf{H} - \frac{\mathbf{H} \mathbf{1}_K \mathbf{1}_K^\top \mathbf{H}}{\mathbf{1}_K^\top \mathbf{H} \mathbf{1}_K}, && \mathbf{g} \triangleq \frac{\mathbf{H} \mathbf{1}_K}{\mathbf{1}_K^\top \mathbf{H} \mathbf{1}_K} \end{aligned} \quad (7)$$

Proposition 1 is proven in Appendix B. This insight is similar to the one used in [5] for the case of static object.

B. Change of Variables to Quadratic Program

Problem (5) remains non-convex in the state variables $(\mathbf{R}_t, \mathbf{p}_t, \mathbf{v}_t, \Omega_t)$ due to quadratic equality constraints. We aim to relax this problem into a convex semidefinite program. Towards this goal, we show how (5) can be rewritten as a quadratically constrained quadratic program (QCQP).

First, we observe that the constraint set is already quadratic: the dynamics (1) are quadratic equalities and $\text{SO}(3)$ constraints on rotations can be written as quadratic equality constraints, see, e.g., [42]. Moreover, the expression of $\mathbf{c}^*$ we found in (6) is also quadratic. Some of the terms in the objective, however, are not quadratic. Particularly, $\mathbf{R}_t \mathbf{B}_i \mathbf{c}^*$ is quadratic and appears in a squared norm, leading to a quartic expression. Fortunately, the rotational invariance of the ℓ_2 norm allows a simple reparameterization of the position variable as $\mathbf{s}_t \triangleq \mathbf{R}_t^\top \mathbf{p}_t$, turning the problem into a QCQP. This result is summarized below.

Proposition 2 (QCQP Formulation): The shape estimation and tracking problem can be equivalently written as a quadratically constrained quadratic program (QCQP):

$$\begin{aligned} \min_{\substack{\mathbf{R}_t, \Omega_t \in \text{SO}(3), \\ \mathbf{s}_t, \mathbf{v}_t \in \mathbb{R}^3, \\ t=1, \dots, T, \\ l=1, \dots, T-1}} \quad & \sum_{t=1}^T \sum_{i=1}^N w_t^i \|\mathbf{R}_t^\top \mathbf{y}_t^i - \mathbf{B}_i \mathbf{c} - \mathbf{s}_t\|^2 + \lambda \|\Delta \mathbf{c}\|^2 \\ & + \sum_{t=1}^{T-1} \omega_t \|\mathbf{v}_{t+1} - \mathbf{v}_t\|^2 + \kappa_t \|\Omega_{t+1} - \Omega_t\|_F^2 \\ \text{s.t.} \quad & \Omega_t \mathbf{s}_{t+1} = \mathbf{s}_t + \mathbf{v}_t, \quad \mathbf{R}_{t+1} = \mathbf{R}_t \Omega_t \end{aligned} \quad (9)$$

where $\mathbf{c}$ is defined as in (6) and is a linear function of $\mathbf{R}_t$ and $\mathbf{s}_t$.

The proof of the proposition is given in Appendix C.

We rewrite (9) in canonical form, separating the quadratically constrained variables $(\mathbf{s}, \mathbf{R}, \Omega)$ from the linearly constrained ones $(\mathbf{v})$:

$$\begin{aligned} f^* = \min_{\substack{\mathbf{x} \in \mathbb{R}^{21T-8} \\ \mathbf{v} \in \mathbb{R}^{3T-3}}} \quad & \mathbf{x}^\top \mathbf{Q} \mathbf{x} + \mathbf{v}^\top \mathbf{P} \mathbf{v} \\ \text{s.t.} \quad & \mathbf{x}^\top \mathbf{A}_i \mathbf{x} + \mathbf{d}_i^\top \mathbf{v} + f_i = 0, \quad i = 1, \dots, m \end{aligned} \quad (10)$$

In this equation, $\mathbf{x}$ is a vector in homogeneous form stacking all the unknowns in (9) except for $\mathbf{v}_t$ terms, which are stacked in $\mathbf{v}$; $\mathbf{Q}$, $\mathbf{P}$, and $\mathbf{A}_i$ are known symmetric matrices governing the quadratic objective and constraints, and the vectors $\mathbf{d}_i$ and scalars f_i capture the linear and constant portions of the m constraints, respectively.

C. Convex Semidefinite Relaxation

While the QCQP in (10) is still non-convex in the variable $\mathbf{x}$, it admits a standard semidefinite relaxation [32, 43, 44]. Instead of solving for $\mathbf{x}$ directly, we reparametrize the problem using $\mathbf{X} = \mathbf{x} \mathbf{x}^\top$ (a rank-1 positive semidefinite matrix), and drop the rank-1 constraint on $\mathbf{X}$ to obtain a convex problem that may be solved by off-the-shelf solvers such as MOSEK [45]. This is the well-known Shor's relaxation [29].

Corollary 3 (Shor's Semidefinite Relaxation): The following semidefinite program (SDP) is a convex relaxation of (10):

$$\begin{aligned} f_{\text{SDP}}^* = \min_{\substack{\mathbf{X} \in \mathbb{S}^{21T-8} \\ \mathbf{v} \in \mathbb{R}^{3T-3}}} \quad & \text{trace}(\mathbf{Q} \mathbf{X}) + \mathbf{v}^\top \mathbf{P} \mathbf{v} \\ \text{s.t.} \quad & \text{trace}(\mathbf{A}_i \mathbf{X}) + \mathbf{d}_i^\top \mathbf{v} + f_i = 0, \\ & \mathbf{X} \succeq 0, \quad i = 1, \dots, m \end{aligned} \quad (11)$$

Further, when the solution $\mathbf{X}^*$ of (11) is rank-1 we can recover exactly the solution to the non-convex QCQP (10) by factorizing $\mathbf{X}^* = \mathbf{x}^* (\mathbf{x}^*)^\top$.

Similar to relaxations derived in related work [5, 32, 46] the rank of $\mathbf{X}^*$ is a *certificate* for the optimality of the solution. Moreover, we can bound the suboptimality of a feasible solution to (10) obtained from (11) using the objective. Given a feasible solution $(\hat{\mathbf{x}}, \hat{\mathbf{v}})$ achieving objective $\hat{f}$ in (10), we bound its suboptimality using $\hat{f} - f_{\text{SDP}}^* \geq \hat{f} - f^* \geq 0$. The condition $\hat{f} = f_{\text{SDP}}^*$ also certifies the optimality of the solution. The scalar $d = f^* - f_{\text{SDP}}^*$ is typically referred to as the *duality gap*, while the difference $\hat{f} - f_{\text{SDP}}^*$ is sometimes called the *relaxation gap*.




Fig. 3. **Outlier Pruning.** Most outliers are easy to identify via shape or time compatibility tests. Shape compatibility retains keypoints that are mutually within the convex hull of the known shape library. Time compatibility compares keypoint pairs over multiple observations and retains groups that preserve 3D distance over time, up to a tolerance ϵ . We determine the largest set of compatible measurements via a mixed integer linear program.

The relaxation (11) is relevant in practice because we observe it to be empirically tight in the case of low-to-moderate noise and no outliers; hence it can produce optimal solutions without needing an initial guess. Moreover, the size of the SDP (11) is independent of the size of the shape library, hence the relaxation is relatively efficient to solve in practice. We name the resulting approach CAST*: *Certifiable Algorithm for Shape estimation and Tracking*.

V. ADDING OUTLIER ROBUSTNESS

Real-world measurements are often corrupted by outliers. In particular, sparse keypoints are vulnerable to misdetections and incorrect depth measurements. Without modifications, outliers degrade the result of CAST*. To tackle this problem, we propose a preprocessing step in which we quickly identify and prune gross outliers, and a wrapper for CAST* that iteratively converges to the inlier set. We name this approach CAST# and show empirical robustness to 50-60% outliers.

A. Compatibility Checks to Remove Gross Outliers

Inspired by [5], we introduce compatibility tests to identify gross outliers. These tests rely on the assumptions of rigid-body motion of the object and the active shape model. The most likely inlier set is thus the largest set of compatible measurements, found via a fast mixed-integer linear program.

Shape Compatibility. Recall that any observed object must lie within the convex hull of the shape library by assumption. Framed as pairwise compatibility, the true distance between any two keypoints i and j must lie somewhere between the minimum and maximum distance between i and j in the shape library models. Therefore, any two keypoint

measurements that are outside this bound cannot simultaneously be inliers; one or both must be outliers. Allowing for keypoint noise expands these bounds as summarized in Proposition 4. Refer to [46] for a full proof.

Proposition 4 (Shape Compatibility Test): Let ϵ be the maximum error for a measurement to be considered an inlier. If a pair of measurements $\mathbf{y}_t^i$ and $\mathbf{y}_t^j$ are both inliers, then:

$$b_{ij}^{\min} - 2\epsilon \leq \|\mathbf{y}_t^i - \mathbf{y}_t^j\| \leq b_{ij}^{\max} + 2\epsilon \quad (12)$$

where $b_{ij}^{\{\min, \max\}}$ are the minimum and maximum distances between keypoints i and j in the shape library:

$$b_{ij}^{\{\min, \max\}} \triangleq \{\min, \max\}_{\mathbf{c} \geq 0, \mathbf{1}^T \mathbf{c} = 1} \|(\mathbf{B}_i - \mathbf{B}_j)\mathbf{c}\| \quad (13)$$

Time Compatibility. For a rigid body the distance between two points is constant over time. This forms the basis for a compatibility test between pairs of points at two times.

Proposition 5 (Time Compatibility Test): Let ϵ be the maximum error for a measurement to be considered an inlier. Consider the measurements of keypoints i and j at times l and m . If these measurements are all inliers then:

$$\|\mathbf{y}_l^i - \mathbf{y}_l^j\| - \|\mathbf{y}_m^i - \mathbf{y}_m^j\| \leq 4\epsilon \quad (14)$$

We provide the proof in Appendix D.

Outlier Pruning. Any set of inliers must satisfy the compatibility conditions presented above. To prune gross outliers, we select the largest set of compatible measurements. Finding this set can be cast as a mixed-integer linear program which we solve using COPT [47].

Proposition 6 (Largest Set of Compatible Measurements): Let $\mathcal{S}$ be the set of measurement pairs that do not satisfy the shape compatibility condition (12) and $\mathcal{T}$ be the set of groups of four measurements that do not satisfy the time compatibility condition (14). The largest set of measurements that satisfy both shape and time compatibility is given by the following mixed integer linear program:

$$\begin{aligned} & \underset{\theta \in \{0,1\}^{N \times T}}{\operatorname{argmin}} \quad \sum_{t=1}^T \sum_{i=1}^N \theta_t^i \\ & \text{s.t.} \quad \theta_t^i + \theta_t^j \leq 1 \quad \forall (t, i, j) \in \mathcal{S} \\ & \quad \theta_l^i + \theta_l^j + \theta_m^i + \theta_m^j \leq 3 \quad \forall (l, m, i, j) \in \mathcal{T} \end{aligned} \quad (15)$$

where $\theta_t^i = 1$ denotes including measurement $\mathbf{y}_t^i$ in the set. The proof follows from Propositions 4 and 5.

B. Graduated Non-Convexity for Robustness.

While consistency checks can remove a significant proportion of outliers, they may miss a number of difficult-to-detect outliers. To remove these remaining outliers we use CAST* as a non-minimal solver for *graduated non-convexity* (GNC) [19]. We use the truncated least squares loss in GNC and follow the implementation and parameter choices of [19]. In our experiments, we show the combination of GNC and our compatibility checks is robust to 50-60% of outliers.




Fig. 4. **Performance of CAST* and CAST# in Synthetic Experiments.** Using the PASCAL3D+ aeroplane shape library, we generate synthetic measurements to test the robustness of CAST* and CAST# to measurement noise, process noise, and outliers. Plots show median and IQR of 500 runs.

VI. EXPERIMENTS

This section characterizes CAST* and CAST# in synthetic and real settings. Synthetic experiments (Section VI-A) show that the semidefinite relaxation in CAST* is empirically tight and returns accurate estimates in the presence of measurement and process noise, while CAST# is robust to 50-60% outliers. Section VI-B shows that CAST# is competitive with other category-level approaches on YCBInEOAT [3], a manipulation focused dataset for object tracking. Lastly, we show that CAST# remains accurate in a real-world drone-based vehicle tracking scenario (Section VI-C).

A. Optimality and Robustness in Synthetic Dataset

Dataset. We generate keypoint measurements according to the measurements and motion models in Section III. Ground truth trajectories follow the constant twist model (2) with Gaussian velocity noise and Langevin rotation rate noise (process noise). The ground truth trajectory and randomly generated shape determine the measured keypoint positions without regard for occlusion, subject to Gaussian perturbations (measurement noise) and outliers. We use the realistic PASCAL3D+ aeroplane shape library [48] (with characteristic length $l = 0.2$ m) to generate a ground truth shape

vector. In each experiment, we fix measurement noise to 5% of the characteristic length, and the process noise to 0.01m and 0.01rad. For the measurement noise experiment, we set the velocity weights $\omega_t = 1$ to improve tightness; we provide results with standard weights in Appendix E.

Baselines. We compare CAST* against PACE [5], a certifiably optimal solver for single-frame pose estimation, and PACE+EKF, an approach that filters the pose estimate from PACE at each time using an extended Kalman filter (EKF) while using a constant-twist motion model [49]. We test CAST* with a time horizon of 4, 8, or 12 frames and label the corresponding results as CAST-4, CAST-8, and CAST-12; we also report CAST-U, which is a variant of CAST-12 with no velocity or rotation smoothing ($\omega_t = 0$, $\kappa_t = 0$). Finally, to test CAST#, we replace a fraction of the measurements with random outliers normally distributed about the centroid of the object with a standard deviation equal to the characteristic length. For the tests with outliers, we compare against ablations of CAST# with only GNC or only compatibility-based (MILP) outlier rejection without GNC and use $T = 12$. In each experiment we show the median and interquartile range of the error of the last estimated state over 500 runs.

TABLE I
COMPARISON OF METHODS ON YCBINEOAT DATASET

Method	Cracker		Sugar		Mustard	
	ADD	ADD-S	ADD	ADD-S	ADD	ADD-S
6-PACK	-	-	-	-	34.49	80.76
TEASER++	84.76	92.14	83.26	91.27	86.02	93.43
MaskFusion	79.74	88.28	36.18	45.62	11.55	13.11
BundleTrack	85.07	89.41	85.56	90.22	92.26	95.35
BundleSDF	81.44	90.76	86.55	92.85	90.83	95.48
CAST#-8	86.93	93.14	81.97	89.45	84.67	92.41
CAST#-GT	89.00	94.09	91.05	95.27	92.18	96.17

Results. Fig 4 reports the median position error (as a percent of length scale), rotation error (in degrees), shape error, and suboptimality gap or solve time for increasing measurement noise (normalized by length scale), process noise (reported as a multiple of 5%), and outlier ratio. In all experiments, CAST* and CAST# achieve the lowest median position, rotation, and shape error. In Figs. 4(a) and 4(b), CAST* is consistently tight (suboptimality gap $< 10^{-4}$) in low to moderate noise, and still gives an accurate estimate when not tight. Interestingly, while CAST* outperforms its unsmoothed variant CAST-U, the latter remains tight for higher measurement noise. CAST-U benefits over PACE from access to additional measurements, and CAST* benefits over CAST-U from additional information about the object’s motion. The primary cost of CAST* compared to PACE is its runtime. In our tests, the runtime ranged from 0.1 to 7 Hz depending on the time horizon (Appendix E). We also note the poor performance of PACE-EKF in both experiments. The EKF provides some benefit for very low noise but quickly diverges for higher noise as the distribution of PACE measurements deviates from Gaussian and the dynamics are nonlinear. We provide a comparison to an EKF using perturbed ground truth poses for the update step in Appendix E.

The outlier experiment in Fig. 4(c) shows robustness to 50-60% of outliers using CAST#. Compatibility tests alone are robust to 40-50% of outliers, while GNC only tolerates 20-30% of outliers. The data show GNC and MILP-based outlier rejection are complimentary, with the fast MILP solve time being unaffected by GNC in the low outlier regime.

B. YCBInEOAT Dataset

Dataset. The YCBInEOAT dataset [3] includes 9 RGB-D videos of a robotic manipulator interacting with 5 YCB objects [50]. It includes in-hand manipulation, pick-and-place, and handovers. Because keypoint detection is not the focus of this work, we train a simple RGB keypoint detector for each object using their CAD models and pre-defined semantic points. The detector has a ResNet18 backbone [51] and is trained on the BOP YCB-V synthetic dataset [52]. We report the ADD and ADD-S area under the precision-accuracy curve (AUC) scores at 0.1 m; see [53].

Baselines. We compare against state-of-the-art instance and category-level tracking approaches for the cracker, sugar, and mustard objects. We omit the small soup object and bleach because it matches the background and gripper colors and our simple keypoint detector is unable to detect reasonable keypoints. TEASER++ [9] is the only instance-level approach and uses the same keypoints given to CAST#.

TABLE II
QUANTITATIVE RESULTS OF DRONE EXPERIMENT

Method	ADD $\uparrow$	R_{err} (deg) $\downarrow$	p_{err} (cm) $\downarrow$	c_{err} $\downarrow$	FPS $\uparrow$
TEASER [9]	57.0	9.6 ± 23.2	4.3 ± 3.8	-	39.1
PACE [46]	52.0	12.1 ± 32.0	3.2 ± 2.4	0.79	3.94
CAST-4	56.6	7.6 ± 4.5	2.7 ± 1.3	0.84	3.65
CAST-8	58.0	7.0 ± 4.3	2.7 ± 1.4	0.76	1.44
CAST-12	58.8	6.5 ± 3.8	2.7 ± 1.4	0.71	0.67
CAST-U	58.2	8.8 ± 15.1	4.6 ± 20.0	0.71	0.91

For CAST#, we group mustard and bleach into the “bottle” category along with a CAD model of a ketchup bottle [54] (3 shapes, 65 keypoints). Similarly, we group cracker and sugar into the “box” category (2 shapes, 52 keypoints). Results from 6-PACK [14], MaskFusion [16], and BundleTrack [25] are taken from the results reported in [25]. BundleSDF [15] results were replicated using the open source implementation with ground truth segmentation masks. BundleSDF is the only category-free method. Since our keypoint detector is fairly simple, we also report CAST# evaluated on ground truth pixel keypoints with true depth and occlusions. For all methods we compute scores using the ground truth shape and initialize with the first frame ground truth pose. For CAST# and TEASER we do not use the ground truth to initialize.

Results. CAST# clearly outperforms baselines for the cracker object but underperforms for sugar and mustard (Table I). The results are encouraging: even with a simple keypoint detector, CAST# outperforms elaborate learning-based methods. The sugar and mustard results are not far behind baselines and mostly reflect the quality of the keypoint detector, which struggles with smaller objects (see video attachment). Given ground truth pixel keypoints, CAST# outperforms virtually all baselines, despite the low-quality depth data in the YCBInEOAT dataset.

C. Drone-based Vehicle Tracking

We use the drone platform described in [55] to evaluate CAST# under dynamic real-world conditions, see the video attachment. During the experiment we remotely piloted a mini racecar in an elliptical trajectory while the Soft Drone autonomously followed using the centroid and heading derived from raw keypoints estimated at 30 Hz. Our keypoint detector, like the YCBInEOAT experiments, used a ResNet-based architecture [51] with 7 keypoints, and was trained on images of a similar racecar. Offline, we used motion capture to transform the 3D position of each keypoint to the world frame (to compensate for the known motion of the drone) and ran CAST# to estimate the racecar’s shape and pose at each time step. For the category-level shape library we used scaled PASCAL3D+ car shapes and the racecar instance.

Quantitative results of TEASER, PACE, CAST#, and variants are given in Table II. TEASER and PACE operate on the same raw keypoint data as CAST# and are tuned for optimal performance. Across metrics, CAST# achieves the highest accuracy and lowest mean errors. In particular, the batch approach with motion priors significantly decreases the standard deviation of rotation and position errors. While the frames per second is not competitive with TEASER, CAST-4 is not significantly slower than PACE despite CAST# using an unoptimized MATLAB implementation.

VII. CONCLUSION

We propose CAST*, a certifiably optimal approach to simultaneously estimate the shape and track the pose of an object of a given category. Our approach uses a semidefinite relaxation to solve the tracking problem, avoiding assumptions of small motion and local minima from local solvers. In the presence of outliers, CAST# uses GNC and a fast shape and time compatibility scheme to reject gross measurements. Our experiments show the relaxation is empirically tight and is competitive against baselines on the YCBInEOAT dataset and in a real-world drone tracking setting.

REFERENCES

- [1] T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Multi-target tracking using joint probabilistic data association," in *Proc. 19th IEEE Conf. on Decision & Control*, 1980.
- [2] Y. Peng, "Deep learning for 3d object detection and tracking in autonomous driving: A brief survey," 2023.
- [3] B. Wen, C. Mitash, B. Ren, and K. E. Bekris, "se(3)-tracknet: Data-driven 6d pose tracking by calibrating image residuals in synthetic domains," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10367–10373.
- [4] V. Lepetit and P. Fua, *Monocular Model-Based 3D Tracking of Rigid Objects: A Survey*. Now Foundations and Trends, 2005.
- [5] J. Shi, H. Yang, and L. Carlone, "Optimal and robust category-level perception: Object pose and shape estimation from 2D and 3D semantic keypoints," *IEEE Trans. Robotics*, vol. 39, no. 5, pp. 4131–4151, 2023, [\(pdf\)](#).
- [6] J. Zhang, M. Wu, and H. Dong, "Genpose: Generative category-level object pose estimation via diffusion models," 2023.
- [7] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. Guibas, "Normalized object coordinate space for category-level 6d object pose and size estimation," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 2642–2651.
- [8] Y. Labbe, J. Carpentier, M. Aubry, and J. Sivic, "CosyPose: Consistent multi-view multi-object 6D pose estimation," in *European Conf. on Computer Vision (ECCV)*, 2020.
- [9] H. Yang, J. Shi, and L. Carlone, "TEASER: Fast and Certifiable Point Cloud Registration," *IEEE Trans. Robotics*, vol. 37, no. 2, pp. 314–333, 2020, extended arXiv version 2001.07715 [\(pdf\)](#).
- [10] L. Wang, S. Yan, J. Zhen, Y. Liu, M. Zhang, G. Zhang, and X. Zhou, "Deep active contours for real-time 6-dof object tracking," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [11] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, "Poserbpf: A rao-blackwellized particle filter for 6d object pose tracking," in *Robotics: Science and Systems (RSS)*, 2019.
- [12] Y. Weng, H. Wang, Q. Zhou, Y. Qin, Y. Duan, Q. Fan, B. Chen, H. Su, and L. J. Guibas, "Captra: Category-level pose tracking for rigid and articulated objects from point clouds," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Oct. 2021. [Online]. Available: <http://dx.doi.org/10.1109/ICCV48922.2021.01296>
- [13] J. Sun, Y. Wang, M. Feng, D. Wang, J. Zhao, C. Stachniss, and X. Chen, "Ick-track: A category-level 6-dof pose tracker using inter-frame consistent keypoints for aerial manipulation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Oct. 2022. [Online]. Available: <http://dx.doi.org/10.1109/IROS47612.2022.9982183>
- [14] C. Wang, R. Martín-Martín, D. Xu, J. Lv, C. Lu, L. Fei-Fei, S. Savarese, and Y. Zhu, "6-pack: Category-level 6d pose tracker with anchor-based keypoints," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10059–10066.
- [15] B. Wen, J. Tremblay, V. Blukis, S. Tyree, T. Müller, A. Evans, D. Fox, J. Kautz, and S. Birchfield, "Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 606–617.
- [16] M. Rünz, M. Buffier, and L. Agapito, "MaskFusion: Real-time recognition, tracking and reconstruction of multiple moving objects," in *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2018, pp. 10–20.
- [17] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models - their training and application," *Comput. Vis. Image Underst.*, vol. 61, no. 1, pp. 38–59, January 1995.
- [18] X. Zhou, S. Leonardos, X. Hu, and K. Daniilidis, "3D shape reconstruction from 2D landmarks: A convex formulation," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [19] H. Yang, P. Antonante, V. Tzoumas, and L. Carlone, "Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection," *IEEE Robotics and Automation Letters (RA-L)*, vol. 5, no. 2, pp. 1127–1134, 2020, arXiv preprint:1909.08605 (with supplemental material), [\(pdf\)](#).
- [20] Y. Bar-Shalom, *Multi-target multisensor tracking: Advanced applications*. Norwood, MA: Artech House, 1992.
- [21] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*, ser. Artech House radar library. Artech House, 1999. [Online]. Available: <https://books.google.com/books?id=ITfAQAIAAJ>
- [22] N. A. Piga, F. Bottarel, C. Fantacci, G. Vezzani, U. Pattacini, and L. Natale, "Maskukf: An instance segmentation aided unscented kalman filter for 6d object pose and velocity tracking," *Frontiers in Robotics and AI*, vol. 8, p. 594583, 2021.
- [23] G. Simon and M.-O. Berger, "A two-stage robust statistical method for temporal registration from features of various type," in *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, 1998, pp. 261–266.
- [24] —, "Pose estimation for planar structures," *IEEE Computer Graphics and Applications*, vol. 22, no. 6, pp. 46–53, 2002.
- [25] B. Wen and K. Bekris, "Bundletrack: 6d pose tracking for novel objects without instance or category-level 3d models," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8067–8074.
- [26] D. Rozumnyi, J. Matas, M. Pollefeys, V. Ferrari, and M. R. Oswald, "Tracking by 3d model estimation of unknown objects in videos," in *ICCV*, Paris, France, Oct. 2023.
- [27] H. Yang and L. Carlone, "Certifiably optimal outlier-robust geometric perception: Semidefinite relaxations and scalable global optimization," *IEEE Trans. Pattern Anal. Machine Intell.*, 2022, [\(pdf\)](#).
- [28] J. Lasserre, *Moments, positive polynomials and their applications*. World Scientific, 2010, vol. 1.
- [29] N. Shor, "Quadratic optimization problems," *Izv. Akad. Nauk SSSR Tekhn. Kibernet.*, vol. 1, pp. 128–139, 1987.
- [30] L. Brynte, V. Larsson, J. P. Iglesias, C. Olsson, and F. Kahl, "On the tightness of semidefinite relaxations for rotation estimation," *Journal of Mathematical Imaging and Vision*, vol. 64, no. 1, p. 57–67, Oct. 2021. [Online]. Available: <http://dx.doi.org/10.1007/s10851-021-01054-y>
- [31] J. Saunderson, P. Parrilo, and A. Willsky, "Semidefinite relaxations for optimization problems over rotation matrices," in *IEEE Conf. on Decision and Control (CDC)*, May 2014.
- [32] D. Rosen, L. Carlone, A. Bandeira, and J. Leonard, "SE-Sync: a certifiably correct algorithm for synchronization over the Special Euclidean group," *Intl. J. of Robotics Research*, 2018, arxiv preprint: 1611.00128, [\(pdf\)](#).
- [33] L. Carlone and F. Dellaert, "Duality-based verification techniques for 2D SLAM," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2015, pp. 4589–4596, [\(pdf\)](#), [\(code\)](#).
- [34] Y. Tian, K. Khosoussi, and J. P. How, "Block-coordinate descent on the Riemannian staircase for certifiably correct distributed rotation and pose synchronization," *Arxiv preprint, arXiv:1911.03721*, 2019.
- [35] J. Briales and J. Gonzalez-Jimenez, "Convex Global 3D Registration with Lagrangian Duality," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [36] J. Briales, L. Kneip, and J. Gonzalez-Jimenez, "A certifiably globally optimal solution to the non-minimal relative pose problem," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [37] J. Zhao, W. Xu, and L. Kneip, "A certifiably globally optimal solution to generalized essential matrix estimation," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [38] M. Garcia-Salguero, J. Briales, and J. Gonzalez-Jimenez, "Certifiable relative pose estimation," *Image and Vision Computing*, vol. 109, p. 104142, 2021.
- [39] L. Sun and Z. Deng, "Certifiably optimal and robust camera pose estimation from points and lines," *IEEE Access*, 2020.
- [40] C. Holmes, F. Dümbs, and T. D. Barfoot, "On semidefinite relaxations for matrix-weighted state-estimation problems in robotics," 2024.
- [41] G. Pavlakos, X. Zhou, A. Chan, K. Derpanis, and K. Daniilidis, "6-dof object pose from semantic keypoints," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2017.
- [42] R. Tron, D. Rosen, and L. Carlone, "On the inclusion of determinant constraints in lagrangian duality for 3D SLAM," in *Robotics: Science and Systems (RSS), Workshop "The problem of mobile sensors: Setting future goals and indicators of progress for SLAM"*, 2015, [\(pdf\)](#).
- [43] J. Briales and J. Gonzalez-Jimenez, "Fast global optimality verification in 3D SLAM," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 4630–4636.
- [44] A. Eriksson, C. Olsson, F. Kahl, and T.-J. Chin, "Rotation averaging and strong duality," *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [45] MOSEK ApS, *The MOSEK optimization toolbox for MATLAB manual. Version 8.1.*, 2017. [Online]. Available: <http://docs.mosek.com/8.1/toolbox/index.html>
- [46] J. Shi, H. Yang, and L. Carlone, "Optimal pose and shape estimation for category-level 3D object perception," in *Robotics: Science and Systems (RSS)*, 2021, arXiv preprint: 2104.08383, [\(pdf\)](#), [\(video\)](#).
- [47] D. Ge, Q. Huangfu, Z. Wang, J. Wu, and Y. Ye, "Cardinal Optimizer (COPT) user guide," <https://guide.coap.online/copt/en-doc>, 2022.
- [48] Y. Xiang, R. Mottaghi, and S. Savarese, "Beyond pascal: A benchmark for 3d object detection in the wild," in *IEEE Winter Conf. on Appl. of Computer Vision*. IEEE, 2014, pp. 75–82.
- [49] M. Brossard, A. Barrau, and S. Bonnabel, "A Code for Unscented Kalman Filtering on Manifolds (UKF-M)," in *2020 International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.
- [50] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The YCB object and Model set: Towards common benchmarks for manipulation research," in *Intl. Conf. on Advanced Robotics (ICAR)*, Jul. 2015, pp. 510–517.
- [51] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," pp. 770–778, 2016.
- [52] M. Denninger, D. Winkelbauer, M. Sundermeyer, W. Boerdijk, M. Knauer, K. H. Strobl, M. Humt, and R. Triebel, "Blenderproc2: A procedural pipeline for photorealistic rendering," *Journal of Open Source Software*, vol. 8, no. 82, p. 4901, 2023. [Online]. Available: <https://doi.org/10.21105/joss.04901>
- [53] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes," in *Robotics: Science and Systems (RSS)*, 2018.
- [54] A. Smit, "Heinz ketchup bottle," Oct 2020. [Online]. Available: <https://grabcad.com/library/heinz-ketchup-bottle-1>
- [55] S. Ubellacker, A. Ray, J. Bern, J. Strader, and L. Carlone, "High-speed aerial grasping using a soft drone with onboard perception," *Nature Robotics*, 2024, [\(pdf\)](#), [\(video\)](#), [\(web\)](#).

APPENDIX A MAXIMUM A POSTERIORI DERIVATION

Here we show that (5) is a maximum a posteriori (MAP) estimator. We first restate the problem:

$$\begin{aligned} \min_{\substack{\mathbf{R}_t, \boldsymbol{\Omega}_t \in \text{SO}(3), \\ \mathbf{p}_t, \mathbf{v}_t \in \mathbb{R}^3, \\ \mathbf{c} \in \mathbb{R}^K, \mathbf{1}_K^\top \mathbf{c} = 1 \\ t=1, \dots, T, \\ l=1, \dots, T-1}} \quad & \sum_{t=1}^T \sum_{i=1}^N w_t^i \|\mathbf{y}_t^i - \mathbf{R}_t \mathbf{B}_i \mathbf{c} - \mathbf{p}_t\|^2 + \lambda \|\Delta \mathbf{c}\|^2 \\ & + \sum_{t=1}^{T-1} \omega_t \|\Delta \mathbf{v}_t\|^2 + \kappa_t \|\Delta \boldsymbol{\Omega}_t\|_F^2 \\ \text{s.t.} \quad & \mathbf{p}_{t+1} = \mathbf{p}_t + \mathbf{R}_t \mathbf{v}_t, \quad \mathbf{R}_{t+1} = \mathbf{R}_t \boldsymbol{\Omega}_t \\ & \Delta \mathbf{v}_t = \mathbf{v}_{t+1} - \mathbf{v}_t, \quad \Delta \boldsymbol{\Omega}_t = \boldsymbol{\Omega}_{t+1} - \boldsymbol{\Omega}_t \end{aligned} \quad (\text{A16})$$

where we introduced auxiliary variables $\Delta \mathbf{v}_t$ and $\Delta \boldsymbol{\Omega}_t$ for the velocity changes. We now show that the first summand in (A16) corresponds to the likelihood of our keypoint measurements (4), while the other terms describe our priors on the shape, velocity, and rotation rates. Denote the quantities to estimate by $\mathbf{z} \triangleq [\mathbf{c}, \{\mathbf{p}_t, \mathbf{R}_t\}_{t=1}^T, \{\mathbf{v}_t, \boldsymbol{\Omega}_t\}_{t=1}^{T-1}, \{\Delta \mathbf{v}_t, \Delta \boldsymbol{\Omega}_t\}_{t=1}^{T-1}]$ belonging to the domain $\mathbb{Z}$ which includes all relevant constraints in (A16). The MAP estimator takes the form:

$$\arg \max_{\mathbf{z} \in \mathbb{Z}} \mathbb{P}(\mathbf{z} | \{\mathbf{y}_t^i\}_{i,t=1}^{N,T}) = \arg \max_{\mathbf{z} \in \mathbb{Z}} \mathbb{P}(\{\mathbf{y}_t^i\}_{i,t=1}^{N,T} | \mathbf{z}) \mathbb{P}(\mathbf{z}) \quad (\text{A17})$$

where we expanded using Bayes rule. Assuming independent measurements, shape independence, and Markovian time-independence, we can rewrite (A17) as:

$$\arg \max_{\mathbf{z} \in \mathbb{Z}} \prod_{i,t=1}^{N,T} \mathbb{P}(\mathbf{y}_t^i | \mathbf{z}) \prod_{t=1}^{T-1} \mathbb{P}(\Delta \mathbf{v}_t) \mathbb{P}(\Delta \boldsymbol{\Omega}_t) \mathbb{P}(\Delta \mathbf{c}) \quad (\text{A18})$$

For the posterior $\mathbb{P}(\mathbf{y}_t^i | \mathbf{z})$ we assume a zero-mean Gaussian with covariance $\boldsymbol{\Sigma}_t^i = \frac{1}{w_t^i} \mathbf{I}_3$. Hence, using (4) :

$$\mathbb{P}(\mathbf{y}_t^i | \mathbf{z}) = \alpha_t^i \exp\left(-\frac{w_t^i}{2} \|\mathbf{y}_t^i - \mathbf{R}_t \mathbf{B}_i \mathbf{c} - \mathbf{p}_t\|^2\right) \quad (\text{A19})$$

with normalization constant α_t^i .

Similarly, for velocity and shape we assume a zero-mean Gaussian prior with covariance $\frac{1}{\omega_t} \mathbf{I}_3$ and $\frac{1}{\lambda} \mathbf{I}_3$ respectively:

$$\mathbb{P}(\Delta \mathbf{v}_t) = \alpha_t^v \exp\left(-\frac{\omega_t}{2} \|\Delta \mathbf{v}_t\|^2\right) \quad (\text{A20})$$

$$\mathbb{P}(\Delta \mathbf{c}) = \alpha_c \exp\left(-\frac{\lambda}{2} \|\Delta \mathbf{c}\|^2\right) \quad (\text{A21})$$

We also assume that the rotation rate follows a Langevin distribution with concentration parameter κ_t :

$$\mathbb{P}(\Delta \boldsymbol{\Omega}_t) = \alpha_t^o \exp\left(-\kappa_t \|\Delta \boldsymbol{\Omega}_t\|_F^2\right) \quad (\text{A22})$$

where α_t^v , α_c , and α_t^o are suitable normalization constants.

Replacing the maximum of the posterior with the minimum of the negative logarithm of the posterior and dropping multiplicative and additive constants we arrive at the result.

APPENDIX B PROOF OF PROPOSITION 1: CLOSED-FORM OPTIMAL SHAPE

Holding all other variables constant, (5) is a linearly constrained least squares program in $\mathbf{c}$. Thus, the minimum with respect to $\mathbf{c}$ is convex and admits a unique solution via the KKT conditions. If we drop objective terms that do not depend on $\mathbf{c}$ in (5), we get:

$$\min_{\substack{\mathbf{c} \in \mathbb{R}^K, \\ \mathbf{1}_K^\top \mathbf{c} = 1}} \sum_{t=1}^T \sum_{i=1}^N w_t^i \|\mathbf{y}_t^i - \mathbf{R}_t \mathbf{B}_i \mathbf{c} - \mathbf{p}_t\|^2 + \lambda \|\Delta \mathbf{c}\|^2 \quad (\text{A23})$$

Expanding the summation over keypoint indices i and moving the weights into the norm:

$$\begin{aligned} \sum_{t=1}^T \left\| \underbrace{\begin{bmatrix} \sqrt{w_t^1} \mathbf{I}_3 \\ \vdots \\ \sqrt{w_t^N} \mathbf{I}_3 \end{bmatrix}}_{\triangleq \mathbf{W}_t} \left(\underbrace{\begin{bmatrix} \mathbf{R}_t^\top (\mathbf{y}_t^1 - \mathbf{p}_t) \\ \vdots \\ \mathbf{R}_t^\top (\mathbf{y}_t^N - \mathbf{p}_t) \end{bmatrix}}_{\triangleq \mathbf{h}_t} - \underbrace{\begin{bmatrix} \mathbf{B}_1 \\ \vdots \\ \mathbf{B}_N \end{bmatrix}}_{\triangleq \mathbf{B}} \mathbf{c} \right) \right\|^2 \\ + \lambda \|\Delta \mathbf{c}\|^2 \end{aligned} \quad (\text{A24})$$

where we used the 2-norm rotational invariance to move $\mathbf{R}_t^\top$.

We can now write the stationarity condition. Using the dual variable μ for the condition $\mathbf{1}^\top \mathbf{c} = 1$ and simplifying:

$$0 = 2\mathbf{B}^\top \left(\sum_{t=1}^T \mathbf{W}_t^2 \right) \mathbf{B} \mathbf{c} - 2\mathbf{B}^\top \left(\sum_{t=1}^T \mathbf{W}_t^2 \mathbf{h}_t \right) + 2\lambda \Delta \mathbf{c} + \mathbf{1} \mu \quad (\text{A25})$$

Putting this together with primal feasibility, we arrive at the following linear system:

$$\begin{bmatrix} \mathbf{H}^{-1} & \mathbf{1}_K \\ \mathbf{1}_K^\top & 0 \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ \mu \end{bmatrix} = \begin{bmatrix} 2 \left(\mathbf{B}^\top \sum_{t=1}^T \mathbf{W}_t^2 \mathbf{h}_t + \lambda \bar{\mathbf{c}} \right) \\ 1 \end{bmatrix} \quad (\text{A26})$$

where $\mathbf{H} \triangleq \frac{1}{2} \left(\mathbf{B}^\top \left(\sum_{t=1}^T \mathbf{W}_t^2 \right) \mathbf{B} + \lambda \mathbf{I}_K \right)^{-1}$.

Solving for $\mathbf{c}$ and substituting the definitions of $\mathbf{G}$ and $\mathbf{g}$, we arrive at the result. Crucially, notice that the matrix we must invert to get $\mathbf{H}$ is made up only of constants.

APPENDIX C PROOF OF PROPOSITION 2: QUADRATICALLY CONSTRAINED QUADRATIC PROGRAM

We focus on the measurement terms, the constraints, and the variable $\mathbf{c}$. The remaining objective terms contain norms of single-degree variables and are thus quadratic. The key idea is to let $\mathbf{s}_t \triangleq \mathbf{R}_t^\top \mathbf{p}_t$. Then, the measurement term of the objective may be rotated without changing its norm:

$$\|\mathbf{y}_t^i - \mathbf{R}_t \mathbf{B}_i \mathbf{c} - \mathbf{p}_t\|^2 = \|\mathbf{R}_t^\top \mathbf{y}_t^i - \mathbf{B}_i \mathbf{c} - \mathbf{s}_t\|^2 \quad (\text{A27})$$

Similarly, the optimal solution $\mathbf{c}^*$ may be rewritten as:

$$\mathbf{c}^* = 2\mathbf{G} \left(\mathbf{B}^\top \sum_{t=1}^T \mathbf{W}_t \begin{bmatrix} \mathbf{R}_t^\top \mathbf{y}_t^1 - \mathbf{s}_t \\ \vdots \\ \mathbf{R}_t^\top \mathbf{y}_t^N - \mathbf{s}_t \end{bmatrix} + \lambda \bar{\mathbf{c}} \right) + \mathbf{g} \quad (\text{A28})$$

to complete the changes needed for the objective. Notice that $\mathbf{c}$ is a linear function of $\mathbf{R}$ and $\mathbf{s}$. Thus, every term of the objective is quadratic in the new variables $\mathbf{s}_t$, $\mathbf{R}_t$, $\mathbf{v}_t$, $\boldsymbol{\Omega}_t$.




Fig. 5. **Performance of CAST* in synthetic experiments with increasing measurement noise.** Robustness to measurement noise with CAST* using the inverse of the simulated velocity covariance for the velocity weights ω_t . The key difference between this plot and Fig. 4(a) lies in the suboptimality gap figure, where CAST* loses tightness quickly. Despite losing its optimality certificate, CAST* maintains the lowest position, rotation, and shape errors.

The variable $\mathbf{p}$ still remains in the constraint $\mathbf{p}_{t+1} = \mathbf{p}_t + \mathbf{R}_t \mathbf{v}_t$. Multiplying both sides by $\mathbf{R}_t^\top$:

$$\mathbf{R}_t^\top \mathbf{p}_{t+1} = \mathbf{s}_t + \mathbf{v}_t \quad (\text{A29})$$

From the rotation rate constraint, $\mathbf{R}_{t+1} = \mathbf{R}_t \Omega_t \Rightarrow \mathbf{R}_t^\top = \Omega_t^\top \mathbf{R}_{t+1}^\top$. Plugging this in gives the desired constraint.

APPENDIX D PROOF OF PROPOSITION 5: TIME COMPATIBILITY TEST

Compare the distance between keypoints i and j at each time, rotating to align coordinates with the body frame:

$$\left| \|\mathbf{R}_l^\top (\mathbf{y}_l^i - \mathbf{y}_l^j)\| - \|\mathbf{R}_m^\top (\mathbf{y}_m^i - \mathbf{y}_m^j)\| \right| \quad (\text{A30})$$

Bound this using the reverse and forward triangle inequalities, noting that noise is isotropic ($\mathbf{R}\epsilon = \epsilon$):

$$\begin{aligned} (\text{A30}) &\leq \|(\mathbf{B}_i - \mathbf{B}_j + \epsilon_l^i - \epsilon_l^j) - (\mathbf{B}_i - \mathbf{B}_j + \epsilon_m^i - \epsilon_m^j)\| \\ &= \|\epsilon_y^{i,l} - \epsilon_y^{i,m} - \epsilon_y^{j,m} + \epsilon_y^{j,l}\| \leq 4\epsilon \end{aligned} \quad (\text{A31})$$

Since the 2-norm is invariant to rotations, we can remove the rotations from (A30) and obtain the result.

APPENDIX E ADDITIONAL EXPERIMENTAL RESULTS

Additional Synthetic Results. In Section VI-A we showed the robustness of CAST* to measurement and process noise, and the robustness of CAST# to outliers. Here we give the runtimes of each method (Table III), show results for the choice of weights resulting from MAP estimation (Fig. 5), and address the EKF results (Fig. 6).

TABLE III
SYNTHETIC EXPERIMENT RUNTIMES

Runtime (s)	PACE	CAST-			
		4	8	12	U
Meas. Noise	0.0028	0.483	2.15	5.49	5.25
Proc. Noise	0.0040	0.857	4.05	10.6	10.2

From Table III we observe CAST* is the slowest of the tested methods. The results are obtained with a non-optimized MATLAB implementation and we expect computational gains from further code optimization. This aside, the variable horizon length allows a trade-off between computational speed and accuracy. As computation improves, the benefits of certifiable optimality and increased accuracy make CAST* an attractive choice of tracking algorithm.

Recall that in the tests in Section VI we chose the velocity weights to be $\omega_t = 1$ instead of setting them as prescribed by

MAP estimation (where they should be taken as the inverse of the variance of the prior). This is equivalent to increasing the standard deviation of the velocity noise by a factor of 10; in other words, it reduces the effect of motion smoothing. Fig. 5 shows that using the true velocity covariance degrades tightness, although it does not have any visible effect on the accuracy results.

Lastly, we present additional results showing the performance of the EKF using perturbed ground truth data instead of PACE. Specifically, we perturb the ground truth poses according to a zero-mean Gaussian with standard deviation equal to 1/25th of the measurement noise for position and 1/50th for rotation (arbitrarily chosen as realistic values). Fig. 6 shows the median EKF estimate consistently beats the perturbed ground truth value (results are averaged over 500 independent trials for each noise value). The large interquartile range is likely because of errors due to linearization, particularly of the constant twist motion model. The EKF likely struggled when using PACE's poses in the measurement update because of the high variance and heavy-tailed distribution of the estimates.




Fig. 6. **Extended Kalman Filter with perturbed ground truth measurements.** With Gaussian-perturbed ground truth measurements, the extended Kalman filter outperforms the raw measurements in median error across measurement noise values. This supports our claim that the EKF performs poorly using pose estimate from PACE, likely due to the high variance and heavy-tailed distribution of the estimates.

Results for Bleach and Soup on YCBInEOAT. Table IV shows scores for all tested methods on the “soup” and “bleach” objects. As mentioned in the text, the soup object is particularly difficult because it is very small and cylindrically symmetric, which CAST* is not designed to

TABLE IV
ADDITIONAL YCBINEOAT RESULTS

Method	Bleach		Soup	
	ADD	ADD-S	ADD	ADD-S
6-PACK	4.18	18.00	12.82	60.32
TEASER++	35.39	46.40	65.85	81.53
MaskFusion	29.83	43.31	5.65	6.45
BundleTrack	89.34	94.72	86.00	95.13
BundleSDF	85.59	93.11	80.54	96.47
CAST#-8	47.53	45.82	27.61	41.70
CAST#-GT	62.19	75.14	37.07	63.29

handle (other approaches also achieve low scores, compared to the other objects). The bleach object is larger but matches the background color, making keypoint detection difficult.