

# The Alignment of the Spheres: Globally-Optimal Spherical Mixture Alignment for Camera Pose Estimation

Dylan Campbell<sup>1</sup>, Lars Petersson<sup>1,2</sup>, Laurent Kneip<sup>3</sup>, Hongdong Li<sup>1</sup> and Stephen Gould<sup>1</sup>

<sup>1</sup>Australian National University

<sup>2</sup>Data61/CSIRO

<sup>3</sup>ShanghaiTech

firstname.lastname@anu.edu.au

lars.petersson@data61.csiro.au

lkneip@shanghaitech.edu.cn

## Abstract

Determining the position and orientation of a calibrated camera from a single image with respect to a 3D model is an essential task for many applications. When 2D–3D correspondences can be obtained reliably, perspective- $n$ -point solvers can be used to recover the camera pose. However, without the pose it is non-trivial to find cross-modality correspondences between 2D images and 3D models, particularly when the latter only contains geometric information. Consequently, the problem becomes one of estimating pose and correspondences jointly. Since outliers and local optima are so prevalent, robust objective functions and global search strategies are desirable. Hence, we cast the problem as a 2D–3D mixture model alignment task and propose the first globally-optimal solution to this formulation under the robust  $L_2$  distance between mixture distributions. We derive novel bounds on this objective function and employ branch-and-bound to search the 6D space of camera poses, guaranteeing global optimality without requiring a pose estimate. To accelerate convergence, we integrate local optimization, implement GPU bound computations, and provide an intuitive way to incorporate side information such as semantic labels. The algorithm is evaluated on challenging synthetic and real datasets, outperforming existing approaches and reliably converging to the global optimum.

## 1. Introduction

Estimating the pose of a calibrated camera given a single image and a 3D model, is useful for many applications, including localization and tracking [20, 33], augmented reality [41], motion segmentation [44] and object recognition [3]. The problem can be cast as a 2D–3D alignment problem in the image plane or on the unit sphere. The task is to find the rotation and translation that aligns the projection of a 3D model with the 2D image data, using points [16, 12], lines [7], silhouettes [14], or mixture models [4]. This is visualized in Figure 1 for mixture models on the unit sphere.

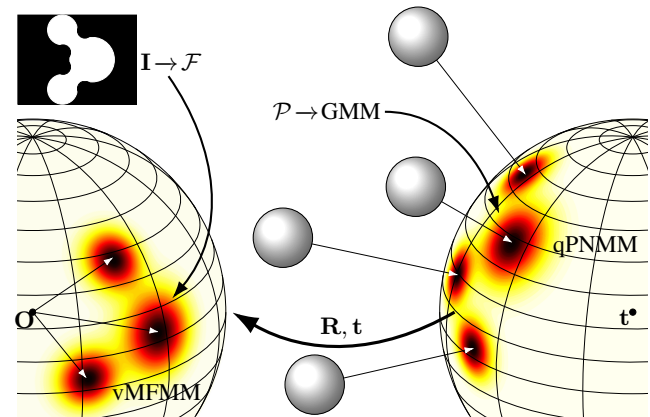


Figure 1. Spherical mixture alignment for estimating the 6-DoF absolute pose  $(\mathbf{R}, \mathbf{t})$  of a camera from a single image  $\mathbf{I}$ , relative to a 3D model (e.g. point-set  $\mathcal{P}$ ), without 2D–3D correspondences. Our algorithm recovers the transformation by generating mixture distributions from the data — a von Mises–Fisher Mixture Model (vMFMM) from the image via a bearing vector set  $\mathcal{F}$  and a Gaussian Mixture Model (GMM) from the 3D model, projected onto the sphere as a quasi-Projected Normal Mixture Model (qPNMM) — and applying branch-and-bound with tight novel bounds to find  $\mathbf{R}$  and  $\mathbf{t}$  that optimally aligns these spherical mixtures.

When 2D–3D correspondences are known, this becomes the well-studied Perspective- $n$ -Point (PnP) problem [37, 25]. However, correspondences between 2D and 3D modalities can be difficult to estimate, not least for the general case of aligning an image with a texture-less 3D model. Even when the model contains visual information, such as SIFT features [40], repetitive elements, occlusions, and appearance variations due to lighting and weather make the correspondence problem non-trivial. Methods that solve for pose and correspondences jointly avoid these problems. They include local optimization approaches [16, 43], which can only yield correct results when a good pose prior is provided, and randomized global search [20], which becomes computationally intractable as the problem size increases. In contrast, globally-optimal approaches [7, 12] obviate the need for pose priors and guarantee optimality.

This work is the first to propose a globally-optimal solution to the 2D–3D mixture alignment problem for camera pose estimation, depicted in Figure 1. The algorithm optimizes the robust  $L_2$  density distance and uses the branch-and-bound framework to guarantee global optimality, addressing the twin challenges of outliers and non-convexity. It provides a geometric solution without assuming that correspondences, pose priors, or training data are available.

The primary contributions are (i) a new closed-form mixture distribution on the sphere, the quasi-Projected Normal mixture, that approximates the projection of a 3D Gaussian mixture; (ii) a new robust objective function, the  $L_2$  distance between von Mises–Fisher and quasi-Projected Normal mixture distributions; (iii) an extension of the objective function to exploit information from deep networks to accelerate convergence; (iv) a fast local optimization algorithm using the objective function and closed-form gradient; (v) novel bounds on the objective function; and (vi) a globally-optimal algorithm for camera pose estimation, with bound computations implemented on the GPU.

An advantage of this approach is that aligning densities is closer to the fundamental 2D–3D problem of aligning physical and imaged surfaces than aligning discrete point samples, since densities model the underlying surfaces with arbitrarily accurate estimates [17], albeit at the limit. Another advantage is that it leverages the adaptive compression properties of mixture model clustering algorithms, enabling the processing of large noisy point-sets. In addition, the continuous objective function admits the use of local gradient-based optimization, which greatly expedites convergence. The algorithm can also be applied to a wide range of 3D data, including mesh and volume representations as well as point-sets. Finally, the approach solves the problem of extracting geometrically-meaningful elements in 2D and 3D by (optionally) using semantic information during optimization. This simple but effective extension reduces runtime and susceptibility to degenerate poses, using only easily-obtainable information.

## 2. Related Work

When 2D–3D correspondences are known,  $PnP$  solvers [37, 25] can accurately estimate the camera pose. However, outliers are almost always present in the correspondence set. When this is the case, the inlier set can be retrieved using RANSAC [20] or robust global optimization [19, 2, 18, 55]. Some of these approaches [20, 19] can be applied when correspondences are not available by providing all possible permutations of the correspondence set. However this hard combinatorial problem quickly becomes infeasible. Matching and filtering techniques have also been developed for large-scale localization problems to reduce the number of outliers in the initial set [49, 39, 61, 18, 55, 50]. These methods are only practical when 2D–3D correspondences

can be found and so are mostly used with Structure-from-Motion (SfM) point-sets. Each 3D point in these datasets is at a visually-distinctive location and is augmented with an image feature, simplifying the correspondence problem. This is not the case for standard point-sets, which contain only geometric information.

The problem is more complex when correspondences are not available at the outset. Local optimization approaches include SoftPOSIT [16], which iterates between solving for correspondences and solving for pose, and 2D/3D GMM registration [4], which projects 3D points into the camera plane then applies 2D Gaussian mixture alignment. This formulation treats points close to the camera identically to distant points and so neglects 3D scale information and creates false optima. Moreover, these methods only find locally-optimal solutions within the convergence basin of the provided pose prior. To alleviate this, global optimization approaches have been proposed, including random-start local search [16] and BlindPnP [43], which uses Kalman filtering to search over a probabilistic pose prior. RANSAC and variants [23] do not require a pose prior, but are only tractable for small numbers of points and outliers. Other approaches use regression forests or convolutional neural networks to learn 2D–3D correspondences from the data and thereby regress pose [52, 29, 5, 28]. These methods require a large training set of pose-labeled images, do not localize the camera with respect to an explicit 3D model, and cannot guarantee optimality.

Globally-optimal approaches can provide this guarantee without needing a pose estimate. They certify that the computed camera pose is a global optimizer of the objective function. The Branch-and-Bound (BB) [35] algorithm has been widely used for this purpose, with tractability continuing to be a significant impediment. For example, BB has been used for 2D–2D registration [6], relative pose estimation [24], 3D–3D rotational registration [38], 3D–3D registration with known correspondences [45], full 3D–3D registration [60], and robust 3D–3D registration [10].

For 2D–3D registration, Brown *et al.* [7] proposed a globally-optimal method using BB with a geometric error. Trimming was used to make the objective function robust to outliers. However this requires knowing the true outlier fraction in advance; if incorrectly specified, the optimum may not occur at the correct pose. Campbell *et al.* [11, 12] proposed a globally-optimal inlier set cardinality maximization solution to the problem. While robust, this objective function is discrete and challenging to optimize, and operates on sampled points instead of the underlying surfaces.

Our work is the first globally-optimal  $L_2$  density distance minimization solution to the camera pose estimation problem. It removes the assumptions that correspondences, training data or pose priors are available and is guaranteed to find the optimum of a robust objective function.

Table 1. Probability distributions in  $\mathbb{R}^3$  and  $\mathbb{S}^2$ .

Distribution	Notation	Parameters	Manifold
Gaussian	$\mathcal{N}$	$\boldsymbol{\mu}, \sigma^2$	$\mathbb{R}^3$
Projected Normal	PN	$\boldsymbol{\mu}, \sigma^2$	$\mathbb{R}^3$
quasi-Projected Normal	qPN	$\boldsymbol{\mu}, \sigma^2$	$\mathbb{S}^2$
von Mises–Fisher	vMF	$\hat{\boldsymbol{\mu}}, \kappa$	$\mathbb{S}^2$

### 3. Probability Distributions on the Sphere

2D directional data such as bearing vectors can be represented as points on the unit 2-sphere. These can be treated as samples from an underlying probability distribution in  $\mathbb{S}^2$ . For images, this distribution models the projection of visible surfaces onto the sphere. In this section, we will outline the probability distributions used in this work and derive a closed-form approximation for the last. The distributions referred to in this paper are summarized in Table 1.

The von Mises–Fisher distribution (vMF) [21], visualized in Figure 2, is the spherical analog of the isotropic Gaussian distribution and has a closed form in 3D, unlike more expressive non-isotropic distributions [30]. The probability density function of the vMF distribution in 3D is

$$\text{vMF}(\mathbf{f} \mid \hat{\boldsymbol{\mu}}, \kappa) = \frac{\exp(\kappa \hat{\boldsymbol{\mu}}^\top \mathbf{f})}{2\pi Z(\kappa)} \quad (1)$$

for the random unit bearing vector  $\mathbf{f}$ , mean direction  $\hat{\boldsymbol{\mu}}$ , and concentration  $\kappa > 0$ , and where

$$Z(\kappa) = (\exp(\kappa) - \exp(-\kappa)) \kappa^{-1}. \quad (2)$$

The Projected Normal (PN) distribution [42, 59, 58] is the projection of a Gaussian distribution  $\mathcal{N}$  onto the sphere. That is, if a random variable  $\mathbf{p}$  follows a Gaussian distribution, then the bearing vector  $\mathbf{f} = \mathbf{p}/\|\mathbf{p}\|$  follows a PN distribution. For a Gaussian mixture that models the distribution of 3D surfaces in a scene, the associated PN mixture models the scene as observed by a 2D sensor, albeit without visibility constraints. The probability density function of the isotropic PN distribution in 3D [46] is

$$\text{PN}(\mathbf{f} \mid \boldsymbol{\mu}, \sigma^2) = \frac{e^{-\frac{\rho^2}{2}}}{2\pi} \left[ \frac{\alpha}{\sqrt{2\pi}} + \Phi(\alpha) e^{\frac{\alpha^2}{2}} (1 + \alpha^2) \right] \quad (3)$$

for the bearing vector  $\mathbf{f}$ , mean position  $\boldsymbol{\mu} \in \mathbb{R}^3$ , and variance  $\sigma^2$ , and where  $\rho = \|\boldsymbol{\mu}\|/\sigma$ ,  $\alpha = \rho \boldsymbol{\mu}^\top \mathbf{f} / \|\boldsymbol{\mu}\|$ , and  $\Phi(\cdot)$  is the cumulative distribution function of  $\mathcal{N}$ .

While PN is the true distribution, it does not have a closed form. Moreover, similarity measures between PN distributions, such as the  $L_2$  distance, are not tractable to compute, since they do not simplify to a closed form when integrated over the sphere and would therefore require time-consuming numerical integration. As a result, it is impractical for alignment problems. Instead, we propose a new




Figure 2. 2D visualization of a 3D von Mises–Fisher distribution as the concentration parameter  $\kappa$  increases. As  $\kappa \rightarrow \infty$ , the distribution approaches a delta function on the sphere.




Figure 3. Comparison of the qPN and PN distributions. (a) The qPN and PN probability density functions are plotted with respect to the angle  $\angle(\mathbf{f}, \boldsymbol{\mu})$  for  $\rho = \|\boldsymbol{\mu}\|/\sigma = 1$ . The distributions are very similar, even for this small value of  $\rho$ . (b) The Mean Absolute Error (MAE) across the entire angular range is plotted with respect to  $\rho$  and is less than 0.01 for all  $\rho \geq 1$ .

closed-form distribution, the *quasi-Projected Normal* (qPN) distribution, that approximates a PN with a vMF distribution. Its probability density function is given by

$$\text{qPN}(\mathbf{f} \mid \boldsymbol{\mu}, \sigma^2) = \text{vMF}\left(\mathbf{f} \mid \frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|}, \left(\frac{\|\boldsymbol{\mu}\|}{\sigma}\right)^2 + 1\right). \quad (4)$$

This was derived by equating the vMF and PN density functions at  $\mathbf{f} = \hat{\boldsymbol{\mu}} = \boldsymbol{\mu}/\|\boldsymbol{\mu}\|$ , since they should evaluate to the same value in the direction of the mean vector. This gives

$$\frac{\kappa}{2\pi(1 - e^{-2\kappa})} = \frac{e^{-\frac{\rho^2}{2}}}{2\pi} \left[ \frac{\rho}{\sqrt{2\pi}} + \Phi(\rho) e^{\frac{\rho^2}{2}} (1 + \rho^2) \right] \quad (5)$$

which simplifies as  $\kappa \rightarrow \infty$  and  $\rho = \|\boldsymbol{\mu}\|/\sigma \rightarrow \infty$  to

$$\kappa = \left(\frac{\|\boldsymbol{\mu}\|}{\sigma}\right)^2 + 1. \quad (6)$$

While this derivation only proves equality in the limit in the direction of the mean vector, the empirical results in Figure 3 show that the distributions are very similar across the entire angular range, even for low values of  $\rho$ .

### 4. Spherical Mixture Alignment

The alignment of mixture distributions to estimate relative sensor pose is a well-studied problem in  $\mathbb{R}^2$ ,  $\mathbb{R}^3$  [15, 57, 27, 10], and the sphere  $\mathbb{S}^2$  [54]. For 2D–3D camera pose estimation, we require a 3D positional and a 2D

directional mixture distribution to model the input data. We model the distribution of 3D points  $\mathbf{p} \in \mathbb{R}^3$  in the set  $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^{N_1}$  as a Gaussian Mixture Model (GMM). Let  $\theta_1 = \{\boldsymbol{\mu}_{1i}, \sigma_{1i}^2, \phi_{1i}\}_{i=1}^{n_1}$  be the parameter set of the  $n_1$ -component GMM with means  $\boldsymbol{\mu}_{1i} \in \mathbb{R}^3$ , variances  $\sigma_{1i}^2$ , and mixture weights  $\phi_{1i} \geq 0$ , where  $\sum \phi_{1i} = 1$ , with density

$$p(\mathbf{p} | \theta_1) = \sum_{i=1}^{n_1} \phi_{1i} \mathcal{N}(\mathbf{p} | \boldsymbol{\mu}_{1i}, \sigma_{1i}^2). \quad (7)$$

We also require a projection of this distribution onto the sphere. For this, we use the qPN Mixture Model (qPNMM) associated with this GMM, with density

$$p(\mathbf{f} | \theta_1) = \sum_{i=1}^{n_1} \phi_{1i} \text{qPN}(\mathbf{f} | \boldsymbol{\mu}_{1i}, \sigma_{1i}^2). \quad (8)$$

Finally, we model the distribution of bearing vectors  $\mathbf{f} \in \mathbb{S}^2$  in the set  $\mathcal{F} = \{\mathbf{f}_i\}_{i=1}^{N_2}$  as a vMF Mixture Model (vMFMM) [22, 53]. Let  $\theta_2 = \{\hat{\boldsymbol{\mu}}_{2j}, \kappa_{2j}, \phi_{2j}\}_{j=1}^{n_2}$  be the parameter set of the  $n_2$ -component vMFMM with mean directions  $\hat{\boldsymbol{\mu}}_{2j} \in \mathbb{S}^2$ , concentrations  $\kappa_{2j} > 0$ , and mixture weights  $\phi_{2j} \geq 0$ , where  $\sum \phi_{2j} = 1$ , with density

$$p(\mathbf{f} | \theta_2) = \sum_{j=1}^{n_2} \phi_{2j} \text{vMF}(\mathbf{f} | \hat{\boldsymbol{\mu}}_{2j}, \kappa_{2j}). \quad (9)$$

The bearing vector  $\mathbf{f}$  corresponds to a 2D point imaged by a calibrated camera. That is,  $\mathbf{f} \propto \mathbf{K}^{-1} \hat{\mathbf{x}}$  where  $\mathbf{K}$  is the matrix of intrinsic camera parameters and  $\hat{\mathbf{x}}$  is the homogeneous image point. These mixture distributions admit arbitrarily accurate estimates of noisy surface densities [17] and can be computed efficiently from the data [9, 34, 53].

The  $L_2$  distance between probability densities is a robust objective function that can be used to measure the alignment of two sets of sensor data, given a specific transformation [27, 54]. Unlike the Kullback–Leibler divergence, it is inherently robust to outliers [51] and operates on statistical densities generated from the raw sensor data. The densities model the underlying surfaces of the scene, which is beneficial because the fundamental 2D–3D registration problem is a surface alignment problem, not a discrete sampled point alignment problem.

**Lemma 1.** ( *$L_2$  objective function*) The  $L_2$  distance between qPNMM and vMFMM models with rotation  $\mathbf{R} \in SO(3)$  and translation  $\mathbf{t} \in \mathbb{R}^3$  can be minimized using the function

$$f(\mathbf{R}, \mathbf{t}) = \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \frac{\phi_{1i} \phi_{2j} Z(K_{1i1j}(\mathbf{t}))}{Z(\kappa_{1i}(\mathbf{t})) Z(\kappa_{2j}(\mathbf{t}))} - 2 \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \frac{\phi_{1i} \phi_{2j} Z(K_{1i2j}(\mathbf{R}, \mathbf{t}))}{Z(\kappa_{1i}(\mathbf{t})) Z(\kappa_{2j}(\mathbf{t}))} \quad (10)$$

where

$$K_{1i1j}(\mathbf{t}) = \left\| \kappa_{1i}(\mathbf{t}) \frac{\boldsymbol{\mu}_{1i} - \mathbf{t}}{\|\boldsymbol{\mu}_{1i} - \mathbf{t}\|} + \kappa_{1j}(\mathbf{t}) \frac{\boldsymbol{\mu}_{1j} - \mathbf{t}}{\|\boldsymbol{\mu}_{1j} - \mathbf{t}\|} \right\| \quad (11)$$

$$K_{1i2j}(\mathbf{R}, \mathbf{t}) = \left\| \kappa_{1i}(\mathbf{t}) \mathbf{R} \frac{\boldsymbol{\mu}_{1i} - \mathbf{t}}{\|\boldsymbol{\mu}_{1i} - \mathbf{t}\|} + \kappa_{2j} \hat{\boldsymbol{\mu}}_{2j} \right\| \quad (12)$$

$$\kappa_{1i}(\mathbf{t}) = \left( \frac{\|\boldsymbol{\mu}_{1i} - \mathbf{t}\|}{\sigma_{1i}} \right)^2 + 1 \quad (13)$$

and  $Z(\cdot)$  is defined as given in (2).

*Proof.* Given qPNMM and vMFMM models of the input data and a rigid transformation function  $T(\theta_1, \mathbf{R}, \mathbf{t}) = \{\mathbf{R}(\boldsymbol{\mu}_{1i} - \mathbf{t}), \sigma_{1i}^2, \phi_{1i}\}_{i=1}^{n_1}$ , the  $L_2$  distance between densities for a rotation  $\mathbf{R}$  and translation  $\mathbf{t}$  is given by

$$d_{L_2} = \int_{\mathbb{S}^2} [p(\mathbf{f} | T(\theta_1, \mathbf{R}, \mathbf{t})) - p(\mathbf{f} | \theta_2)]^2 d\mathbf{f} \quad (14)$$

$$= \int_{\mathbb{S}^2} [p(\mathbf{f} | T(\theta_1, \mathbf{R}, \mathbf{t}))]^2 + [p(\mathbf{f} | \theta_2)]^2 - 2p(\mathbf{f} | T(\theta_1, \mathbf{R}, \mathbf{t}))p(\mathbf{f} | \theta_2) d\mathbf{f}. \quad (15)$$

The function (10) is obtained by removing constant summands and factors, substituting (8), (9), (4) and (1) into (15), and replacing integrals of the form  $\int_{\mathbb{S}^2} \exp(\mathbf{x}^T \mathbf{f}) d\mathbf{f}$  with the normalization constant of a vMF density with  $\kappa = \|\mathbf{x}\|$  and  $\hat{\boldsymbol{\mu}} = \mathbf{x}/\kappa$ ; see appendix for details.  $\square$

The objective is then to find a rotation and translation that minimizes the  $L_2$  distance between the densities

$$(\mathbf{R}^*, \mathbf{t}^*) = \arg \min_{\mathbf{R}, \mathbf{t}} f(\mathbf{R}, \mathbf{t}). \quad (16)$$

Furthermore, if semantic class labels are available, for example using CNN–based semantic segmentation for 2D images [47, 13, 48] and 3D point-sets [36, 26, 56], the optimization problem can be formulated as a joint  $L_2$  distance minimization over the semantic classes, providing semantic-aware alignment and accelerating convergence. That is, given a class label set  $\Lambda$ , one can construct separate mixture distributions for each class and solve

$$(\mathbf{R}^*, \mathbf{t}^*) = \arg \min_{\mathbf{R}, \mathbf{t}} \sum_{l \in \Lambda} \phi_l f_l(\mathbf{R}, \mathbf{t}). \quad (17)$$

where  $\phi_l \geq 0$  are the class weights and  $f_l$  is the per-class function value computed according to (10).

## 5. Branch-and-Bound

To solve the highly non-convex  $L_2$  distance problem (16), the Branch-and-Bound (BB) algorithm [35] may be applied. It requires an efficient way to branch the function domain and bound the function optimum for each branch,






Figure 4. Parameterization and branching of  $SE(3)$ . (a) Rotations are parameterized by angle-axis 3-vectors in a solid radius- $\pi$  ball. (b) Translations are parameterized by 3-vectors bounded by a cuboid with half-widths  $[\tau_x, \tau_y, \tau_z]$ . The joint domain is branched into 6D cuboids using an adaptive octree-like branching strategy.

such that the bounds converge as the branch size tends to zero. The efficiency of the algorithm depends on both the computational complexity of the bounds and how tight they are, since tighter bounds reduce the search space quicker by allowing suboptimal branches to be pruned.

### 5.1. Parameterizing and Branching the Domain

To find a globally-optimal solution, the  $L_2$  distance must be optimized over the domain of 3D motions, that is, the group  $SE(3) = SO(3) \times \mathbb{R}^3$ . For BB, the domain must be bounded, so we restrict the space of translations to the set  $\Omega_t$ , under the assumption that the camera is a finite distance from the 3D model. The domains are shown in Figure 4. We minimally parameterize rotation space  $SO(3)$  with angle-axis 3-vectors  $\mathbf{r}$  with rotation angle  $\|\mathbf{r}\|$  and rotation axis  $\hat{\mathbf{r}} = \mathbf{r}/\|\mathbf{r}\|$ . As a result, the space of all 3D rotations can be represented as a solid ball of radius  $\pi$  in  $\mathbb{R}^3$ . For ease of manipulation, we use the 3D cube circumscribing the  $\pi$ -ball as the rotation domain  $\Omega_r$  [38]. The notation  $\mathbf{R}_r \in SO(3)$  is used to denote the rotation matrix obtained from  $\mathbf{r}$  using Rodrigues' rotation formula. We parameterize translation space  $\mathbb{R}^3$  with 3-vectors in a bounded domain chosen as the cuboid  $\Omega'_t$  containing the bounding box of the model. To avoid the non-physical case where the camera is located within a small value  $\zeta$  of a 3D surface, the translation domain is restricted such that  $\Omega_t = \Omega'_t \cap \{\mathbf{t} \in \mathbb{R}^3 \mid \|\boldsymbol{\mu} - \mathbf{t}\| \geq \zeta, \forall \boldsymbol{\mu} \in \boldsymbol{\theta}_1\}$ . Finally, we branch the domain into 6D cuboids (6-orthotopes)  $\mathcal{C}_r \times \mathcal{C}_t$  using an adaptive branching strategy that chooses to subdivide the rotation or translation dimensions based on which has the greater angular uncertainty, reducing redundant branching.

### 5.2. Bounding the Branches

The success of a branch-and-bound algorithm is predicated on the quality of its bounds. For  $L_2$  distance minimization, we require bounds on the minimum of the objective function (10) within a transformation domain  $\mathcal{C}_r \times \mathcal{C}_t$ .

An upper bound can be found by evaluating the function at any transformation in the branch. A lower bound can be found using the bounds  $\psi_r$  and  $\psi_t$  on the rotation and translation uncertainty angles derived in Lemmas 3 and 5 in Campbell *et al.* [12], reproduced here as Lemmas 2 and 3.

**Lemma 2.** (Rotation uncertainty angle bound) Given a 3D point  $\mathbf{p}$  and a rotation cube  $\mathcal{C}_r$  centered at  $\mathbf{r}_0$  with surface  $\mathcal{S}_r$ , then  $\forall \mathbf{r} \in \mathcal{C}_r$ ,

$$\angle(\mathbf{R}_r \mathbf{p}, \mathbf{R}_{\mathbf{r}_0} \mathbf{p}) \leq \min \left\{ \max_{\mathbf{r} \in \mathcal{S}_r} \angle(\mathbf{R}_r \mathbf{p}, \mathbf{R}_{\mathbf{r}_0} \mathbf{p}), \pi \right\} \triangleq \psi_r(\mathbf{p}, \mathcal{C}_r). \quad (18)$$

**Lemma 3.** (Translation uncertainty angle bound) Given a 3D point  $\mathbf{p}$  and a translation cuboid  $\mathcal{C}_t$  centered at  $\mathbf{t}_0$  with vertices  $\mathcal{V}_t$ , then  $\forall \mathbf{t} \in \mathcal{C}_t$ ,

$$\angle(\mathbf{p} - \mathbf{t}, \mathbf{p} - \mathbf{t}_0) \leq \begin{cases} \max_{\mathbf{t} \in \mathcal{V}_t} \angle(\mathbf{p} - \mathbf{t}, \mathbf{p} - \mathbf{t}_0) & \text{if } \mathbf{p} \notin \mathcal{C}_t \\ \pi & \text{else} \end{cases} \triangleq \psi_t(\mathbf{p}, \mathcal{C}_t). \quad (19)$$

**Theorem 1.** (Objective function bounds) For the transformation domain  $\mathcal{C}_r \times \mathcal{C}_t$  centered at  $(\mathbf{r}_0, \mathbf{t}_0)$ , the minimum of the objective function (10) has an upper bound

$$\bar{d} \triangleq f(\mathbf{R}_{\mathbf{r}_0}, \mathbf{t}_0) \quad (20)$$

and a lower bound

$$\underline{d} \triangleq \sum_{i=1}^{n_1} \sum_{j=1}^{n_1} \phi_{1i} \phi_{1j} \min_{\mathbf{t} \in \mathcal{C}_t} \frac{Z(\underline{K}_{1i1j}(\mathbf{t}))}{Z(\kappa_{1i}(\mathbf{t}))Z(\kappa_{1j}(\mathbf{t}))} - 2 \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \phi_{1i} \phi_{2j} \max_{\mathbf{t} \in \mathcal{C}_t} \frac{Z(\overline{K}_{1i2j}(\mathbf{t}))}{Z(\kappa_{1i}(\mathbf{t}))Z(\kappa_{2j}(\mathbf{t}))} \quad (21)$$

where

$$\underline{K}_{1i1j}(\mathbf{t}) = \sqrt{\kappa_{1i}^2(\mathbf{t}) + \kappa_{1j}^2(\mathbf{t}) + 2\kappa_{1i}(\mathbf{t})\kappa_{1j}(\mathbf{t}) \cos A} \quad (22)$$

$$\overline{K}_{1i2j}(\mathbf{t}) = \sqrt{\kappa_{1i}^2(\mathbf{t}) + \kappa_{2j}^2(\mathbf{t}) + 2\kappa_{1i}(\mathbf{t})\kappa_{2j}(\mathbf{t}) \cos B} \quad (23)$$

$$A = \min\{\pi, \angle(\boldsymbol{\mu}_{1i} - \mathbf{t}_0, \boldsymbol{\mu}_{1j} - \mathbf{t}_0) + \psi_t(\boldsymbol{\mu}_{1i}, \mathcal{C}_t) + \psi_t(\boldsymbol{\mu}_{1j}, \mathcal{C}_t)\} \quad (24)$$

$$B = \max\{0, \angle(\boldsymbol{\mu}_{1i} - \mathbf{t}_0, \mathbf{R}_{\mathbf{r}_0}^{-1} \hat{\boldsymbol{\mu}}_{2j}) - \psi_t(\boldsymbol{\mu}_{1i}, \mathcal{C}_t) - \psi_r(\hat{\boldsymbol{\mu}}_{2j}, \mathcal{C}_r)\}. \quad (25)$$

*Proof.* The validity of the upper bound follows from

$$f(\mathbf{R}_{\mathbf{r}_0}, \mathbf{t}_0) \geq \min_{\substack{\mathbf{r} \in \mathcal{C}_r \\ \mathbf{t} \in \mathcal{C}_t}} f(\mathbf{R}_r, \mathbf{t}). \quad (26)$$

That is, the function value at a specific point within the domain is greater than or equal to the minimum within the domain. For the lower bound, observe that  $\forall \mathbf{t} \in \mathcal{C}_t$ ,

$$K_{1i1j}(\mathbf{t}) = \sqrt{\frac{2\kappa_{1i}(\mathbf{t})\kappa_{1j}(\mathbf{t}) \cos \angle(\boldsymbol{\mu}_{1i}-\mathbf{t}, \boldsymbol{\mu}_{1j}-\mathbf{t})}{+\kappa_{1i}^2(\mathbf{t}) + \kappa_{1j}^2(\mathbf{t})}} \quad (27)$$

$$\geq \sqrt{\frac{\kappa_{1i}^2(\mathbf{t}) + \kappa_{1j}^2(\mathbf{t}) + 2\kappa_{1i}(\mathbf{t})\kappa_{1j}(\mathbf{t}) \cos A}{}} \quad (28)$$

$$= \underline{K}_{1i1j}(\mathbf{t}) \quad (29)$$

where (28) follows from the triangle inequality in spherical geometry and Lemma 3, since

$$\angle(\mathbf{a}, \mathbf{b}) \leq \angle(\mathbf{c}, \mathbf{d}) + \angle(\mathbf{a}, \mathbf{c}) + \angle(\mathbf{b}, \mathbf{d}) \quad (30)$$

$$\leq \angle(\mathbf{c}, \mathbf{d}) + \psi_t(\boldsymbol{\mu}_{1i}, \mathcal{C}_t) + \psi_t(\boldsymbol{\mu}_{1j}, \mathcal{C}_t) \quad (31)$$

for  $\mathbf{a} = \boldsymbol{\mu}_{1i}-\mathbf{t}$ ,  $\mathbf{b} = \boldsymbol{\mu}_{1j}-\mathbf{t}$ ,  $\mathbf{c} = \boldsymbol{\mu}_{1i}-\mathbf{t}_0$ , and  $\mathbf{d} = \boldsymbol{\mu}_{1j}-\mathbf{t}_0$ . Also observe that  $\forall(\mathbf{r}, \mathbf{t}) \in (\mathcal{C}_r \times \mathcal{C}_t)$ ,

$$K_{1i2j}(\mathbf{R}_r, \mathbf{t}) = \sqrt{\frac{2\kappa_{1i}(\mathbf{t})\kappa_{2j} \cos \angle(\boldsymbol{\mu}_{1i}-\mathbf{t}, \mathbf{R}_r^{-1}\hat{\boldsymbol{\mu}}_{2j})}{+\kappa_{1i}^2(\mathbf{t}) + \kappa_{2j}^2}} \quad (32)$$

$$\leq \sqrt{\frac{\kappa_{1i}^2(\mathbf{t}) + \kappa_{2j}^2 + 2\kappa_{1i}(\mathbf{t})\kappa_{2j} \cos B}{}} \quad (33)$$

$$= \bar{K}_{1i2j}(\mathbf{t}) \quad (34)$$

where (33) follows from the reverse triangle inequality in spherical geometry and Lemmas 3 and 2. With  $\underline{K}_{1i1j}$  and  $\bar{K}_{1i2j}$ , a valid lower bound for (10) can be constructed by observing that  $Z(x)$  (2) is a monotonically increasing function for  $x \geq 0$  and the dependency of  $\kappa_{1i}$  on  $\mathbf{t}$  can be optimized separately. See the appendix for the full proof.  $\square$

## 6. The GOSMA Algorithm

The Globally-Optimal Spherical Mixture Alignment (GOSMA) algorithm is outlined in Algorithm 1. It employs a depth-first search strategy using a priority queue (line 5) where the priority is inverse to the lower bound. The algorithm terminates with  $\epsilon$ -optimality, whereby the difference between the best function value  $d^*$  and the global lower bound  $\underline{d}$  is less than  $\epsilon$  (line 6). Branching and bounding is performed on the GPU (line 7), with each thread computing the bounds for a single branch.

We also developed a local optimization algorithm denoted as Spherical Mixture Alignment (SMA), which was integrated into GOSMA (line 9). We use the quasi-Newton L-BFGS algorithm [8] to minimize (10), with the gradient derived in the appendix. SMA is run whenever the BB algorithm finds a sub-cube  $\mathcal{C}_{ij}$  that has an upper bound less than the best-so-far function value  $d^*$  (line 9), initialized with the center transformation of  $\mathcal{C}_{ij}$ . In this way, BB and SMA collaborate, with SMA quickly converging to the closest local minimum and BB guiding the search into the convergence basins of better local minima. SMA accelerates

---

**Algorithm 1** GOSMA: a globally-optimal spherical mixture alignment algorithm for camera pose estimation

---

**Input:** GMM-vMFMM pair, threshold  $\epsilon$ , domain  $\Omega_r \times \Omega_t$

**Output:** optimal function value  $d^*$ , camera pose  $(\mathbf{r}^*, \mathbf{t}^*)$

- 1:  $d^* \leftarrow \infty$
  - 2: Add domain  $\Omega_r \times \Omega_t$  to priority queue  $Q$
  - 3: **loop**
  - 4:   Update lowest lower bound  $\underline{d}$  from  $Q$
  - 5:   Remove set of hypercubes  $\{\mathcal{C}_i\}$  from  $Q$
  - 6:   **if**  $d^* - \underline{d} \leq \epsilon$  **then** terminate
  - 7:   Evaluate  $\bar{d}_{ij}$  &  $\underline{d}_{ij}$  in parallel for subcubes of  $\{\mathcal{C}_i\}$
  - 8:   **for all** sub-cubes  $\mathcal{C}_{ij} \in \{\mathcal{C}_i\}$  **do**
  - 9:     **if**  $\bar{d}_{ij} < d^*$  **then**  $(d^*, \mathbf{r}^*, \mathbf{t}^*) \leftarrow \text{SMA}(\mathbf{r}_{0ij}, \mathbf{t}_{0ij})$
  - 10:    **if**  $\underline{d}_{ij} < d^*$  **then** add  $\mathcal{C}_{ij}$  to queue  $Q$
- 

convergence since reducing  $d^*$  early allows larger branches to be culled (line 10), greatly reducing the search space.


## 7. Results

The GOSMA algorithm, denoted GS, was evaluated with respect to the baseline algorithms SoftPOSIT [16], BlindPnP [43], RANSAC [20], and GOPAC [12], denoted SP, BP, RS and GP respectively, using both synthetic and real data. The RANSAC approach uses the OpenGV framework [31] and the P3P algorithm [32] with randomly-sampled correspondences. To generate GMMs and vMFMMs, we cluster the point-set with DP-means [34] and the bearing vector set with DP-vMF-means [53], and fit maximum likelihood mixture models to the clusters. These methods automatically select a parsimonious representation that adapts to the complexity of the scene geometry.

We report the median translation error (in metres), rotation error (degrees), and runtime (seconds) including on-the-fly mixture generation. We also report the success rate, a summary statistic defined as the fraction of experiments where the correct pose was found: an angular error less than 0.1 radians and a relative translation error less than 5%. Except where otherwise specified, the normalized  $L_2$  distance threshold  $\epsilon$  was set to 0.1, the point-to-camera limit  $\zeta$  was set to 0.5, the scale parameters for mixture model generation  $\lambda_p$  and  $\lambda_f$  were set to 0.25m and  $2^\circ$  respectively, and semantic information was used in the real data experiments only, with class weights  $\phi_l = |\Lambda|^{-1}$ , the inverse of the number of classes. All experiments were run on a 3.4GHz CPU and two GeForce GTX 1080Ti GPUs, and the C++ code is available on the first author's website.

### 7.1. Synthetic Data Experiments

To evaluate GOSMA under a range of perturbations, 25 independent Monte Carlo simulations were performed per parameter setting using the framework of BlindPnP [43]:



(a) 3D Results (b) 2D Results (c) Bound Evolution  
 Figure 5. Sample 2D and 3D results for the random point data. (a) 3D points, true and GOSMA-estimated camera fulcra (completely overlapping) and toroidal pose prior. (b) 2D points (dots) and 3D points projected using the GOSMA-estimated camera pose (circles), with 2D and 3D outliers shown in red. (c) Evolution over time of the upper (red) and lower (magenta) bounds, remaining unexplored volume (blue) and queue size (green) as a fraction of their maximum values.

$N_I$  random 3D point inliers and  $\omega_{3D}N_I$  outliers were generated from  $[-1, 1]^3$ ; the inliers were projected to a  $640 \times 480$  virtual image with a focal length of 800; normal noise with  $\sigma = 2$  pixels was added to the 2D points; and  $\omega_{2D}N_I$  random outlier points were added to the image. An example of the data and alignment results is shown in Figure 5.

The time evolution of the global upper and lower bounds is shown in Figure 5(c). The plot reveals how local and global optimization strategies collaborate to decrease the upper bound with BB guiding the search into better convergence basins and SMA jumping to the nearest local minimum (the staircase pattern). It also shows that the majority of the runtime is spent increasing the lower bound, indicating that GOSMA will often find the global optimum when terminated early, albeit without an optimality guarantee.

To facilitate fair comparison with the local methods SoftPOSIT and BlindPnP, a torus pose prior was used for these experiments. It constrains the camera center to a torus around the 3D point-set with the optical axis directed towards the model [43]. The torus prior was represented as a 50 component GMM for BlindPnP and 50 initial poses for SoftPOSIT. GOSMA and GOPAC were given a set of translation cubes that approximated the torus and were not given any rotation prior. RANSAC was set to explore correspondence space for up to 120s. The results are shown in Figure 6. Runtime values are clipped to an upper limit of 120s so that the scale is interpretable. GOSMA and GOPAC outperform the other methods, reliably finding the correct pose while still being relatively efficient. While GOSMA has longer runtimes in the first two experiments, it has much better behaviour than the other methods when 2D outliers are present. For example, when  $\omega_{2D} = 1$ , the median runtime of GOPAC (167s) is more than 30x higher than GOSMA (5s), while both always find the correct pose, with median angular errors below  $1^\circ$  and relative translation errors below 2%. In fact, this random point setup significantly favors point-based methods such as GOPAC at the expense of our approach. For real surfaces, GOSMA is able to lever-

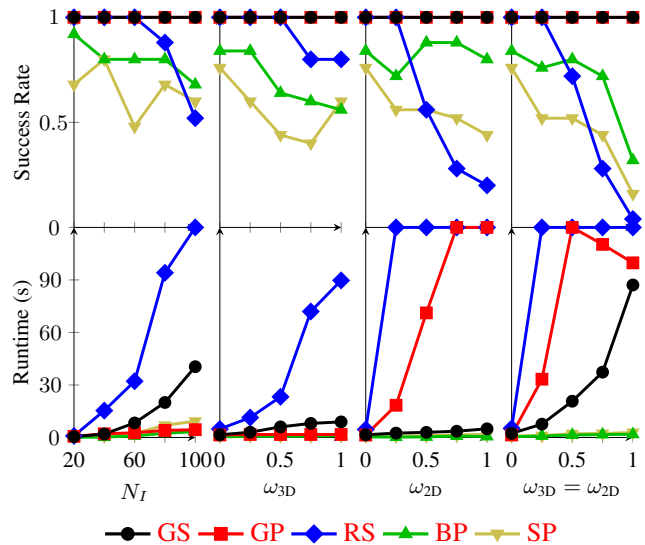


Figure 6. Results for the random points dataset with the torus prior. The success rates and median runtimes are plotted with respect to the number of inlier points ( $N_I$ ), the fraction of additional 3D outliers ( $\omega_{3D}$ ), 2D outliers ( $\omega_{2D}$ ), and both, with default parameters  $N_I = 30$  inlier points and  $\omega_{3D} = \omega_{2D} = 0$ , for 25 Monte Carlo simulations per parameter value.

age its ability to adaptively compress the data, allowing it to quickly process a very large number of points.

## 7.2. Real Data Experiments

The Stanford 2D-3D-Semantics (2D-3D-S) [1] dataset contains panoramic images, point-sets, and semantic annotations for both modalities. It is a large indoor dataset with approximately 1 million points per room and 8 million pixels per photo, collected using a structured-light RGBD camera. We evaluated our algorithm on area 3 of the dataset, which contains lounges, offices and a conference room. The test data has 33 panoramic images taken from distinct camera poses where the camera is at least 50cm from any surface, and covers 13 rooms. Each room is a separate point-set, which models visibility constraints but assumes that the camera’s position is known to the room level. Using this information, we set the translation domain to be the room size. Semantic information is used by all methods in these tests: GOPAC and RANSAC use the pre-processing strategy from Campbell *et al.* [12], selecting points and pixels from furniture classes only, whereas GOSMA uses class labels during optimization (17), making more effective use of the information. We also randomly downsample the point-sets and images to 100k points and pixels to reduce the mixture generation time. The mixture scale parameters  $\lambda_p$  and  $\lambda_f$  [34, 53] were automatically selected to yield approximately 10 components per semantic class, 60–100 components in total. For GOPAC, the inlier threshold  $\theta$  was set to



Table 2. Camera pose results for GOSMA (GS), GOSMA without class labels during optimization ( $GS^{-\Lambda}$ ), GOPAC (GP) and RANSAC (RS) for area 3 of the Stanford 2D-3D-S panoramic image dataset. Translation error, rotation error and runtime quartiles ( $Q_2$   $Q_3$   $Q_1$ ) and the success rate are reported.

Method	GS	$GS^{-\Lambda}$	GP	RS
Translation error (m)	<b>0.08</b> <sub>0.05</sub> <sup>0.15</sup>	0.14 <sub>0.09</sub> <sup>0.23</sup>	0.15 <sub>0.10</sub> <sup>0.27</sup>	0.56 <sub>0.39</sub> <sup>2.06</sup>
Rotation error (°)	<b>1.13</b> <sub>0.91</sub> <sup>2.18</sup>	2.38 <sub>1.25</sub> <sup>4.61</sup>	3.78 <sub>2.47</sub> <sup>5.10</sup>	18.3 <sub>8.94</sub> <sup>108</sup>
Runtime (s)	<b>1.8</b> <sub>1.4</sub> <sup>4.4</sup>	19.1 <sub>12.8</sub> <sup>43.7</sup>	902 <sub>448</sub> <sup>902</sup>	120 <sub>120</sub> <sup>120</sup>
Success rate	<b>1.00</b>	0.85	0.82	0.09


2.5° and the angular tolerance  $\eta$  was set to 0.25°.

Qualitative and quantitative results are given in Figure 7 and Table 2. Note that GOPAC and RANSAC were terminated at 900s and 120s respectively. GOSMA outperforms the other methods considerably, finding the correct pose in all cases with a small median runtime. We also tested GOSMA without semantic labels during optimization, only during pre-processing ( $GS^{-\Lambda}$ ), the same as for GOPAC and RANSAC. While this is more accurate and much faster than GOPAC, optimizing across the semantic classes provides another large accuracy and runtime gain. We would like to emphasize the difficulty of this problem setup: the algorithm is given an image, a point-set and semantic class labels, and is required to estimate the camera pose. Compared to the synthetic data experiments, the sheer number of points and pixels, many of which are outliers, precludes the use of traditional methods.

## 8. Discussion and Conclusion

In this paper, we have proposed a novel mixture alignment formulation for the camera pose estimation problem using the robust  $L_2$  density distance on the sphere. Furthermore, we have developed a novel algorithm to minimize this distance using branch-and-bound, guaranteeing optimality regardless of initialisation. To accelerate convergence, a local optimization algorithm was developed and integrated, GPU bound computations were implemented, and a principled way to incorporating side information such as semantic labels was devised. The algorithm outperformed other local and global methods on challenging synthetic and real datasets, finding the global optimum reliably.

This approach has several limitations, however. Firstly, it scales quadratically with the number of mixture components, which scales with surface complexity. Secondly, it is unable to resolve certain degenerate poses, such as when a wall fills the field-of-view of the camera. In this case, many camera poses satisfy the 2D information. Thirdly, it does not use a geometric objective function, which reduces its interpretability. A robust objective function in the image space such as intersection-over-union would be preferred,



(a) 3D point-set and ground-truth (gray), GOSMA (black), GOPAC (red) and RANSAC (blue) camera poses. Object points are highlighted in green.



(b) 3D points projected onto the image using the GOSMA (top), GOPAC (middle), and RANSAC (bottom) camera poses. For clarity, only object points are plotted.

Figure 7. Qualitative camera pose results for office 3 of the Stanford 2D-3D-S dataset, showing the pose of the camera when capturing the image and the projection of 3D object points onto it. Only GOSMA found the correct pose as defined in this section. Best viewed in color.

although it is not tractable for mixtures on the sphere. Finally, the quality of its pose estimate depends on how well the mixtures represent the physical and projected surfaces in the real scene and image. While they can represent these surfaces arbitrarily accurately, the number of components is limited by practical considerations. Anisotropic densities would be preferred for this reason, however only isotropic densities, which model surfaces as points with a symmetric field of influence, have a tractable closed form on the sphere. Hence, further investigation is warranted into aligning representations that model surfaces with fewer parameters, such as wireframes or meshes.



## References

- [1] I. Armeni, A. Sax, A. R. Zamir, and S. Savarese. Joint 2D-3D-semantic data for indoor scene understanding. *ArXiv e-prints*, Feb. 2017. [7](#)
- [2] E. Ask, O. Enqvist, and F. Kahl. Optimal geometric fitting under the truncated  $L_2$ -norm. In *Proceedings of the 2013 Conference on Computer Vision and Pattern Recognition*, pages 1722–1729. IEEE, June 2013. [2](#)
- [3] M. Aubry, D. Maturana, A. A. Efros, B. C. Russell, and J. Sivic. Seeing 3D chairs: exemplar part-based 2D-3D alignment using a large dataset of CAD models. In *Proceedings of the 2014 Conference on Computer Vision and Pattern Recognition*, pages 3762–3769. IEEE, June 2014. [1](#)
- [4] N. Baka, C. Metz, C. J. Schultz, R.-J. van Geuns, W. J. Niessen, and T. van Walsum. Oriented Gaussian mixture models for nonrigid 2D/3D coronary artery registration. *IEEE Transactions on Medical Imaging*, 33(5):1023–1034, 2014. [1](#), [2](#)
- [5] E. Brachmann, A. Krull, S. Nowozin, J. Shotton, F. Michel, S. Gumhold, and C. Rother. DSAC – Differentiable RANSAC for camera localization. In *Proceedings of the 2017 Conference on Computer Vision and Pattern Recognition*, pages 2492–2500, July 2017. [2](#)
- [6] T. M. Breuel. Implementation techniques for geometric branch-and-bound matching methods. *Computer Vision and Image Understanding*, 90(3):258–294, June 2003. [2](#)
- [7] M. Brown, D. Windridge, and J.-Y. Guillemaut. Globally optimal 2D-3D registration from points or lines without correspondences. In *Proceedings of the 2015 International Conference on Computer Vision*, pages 2111–2119, Dec. 2015. [1](#), [2](#)
- [8] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995. [6](#)
- [9] D. Campbell and L. Petersson. An adaptive data representation for robust point-set registration and merging. In *Proceedings of the 2015 International Conference on Computer Vision*, pages 4292–4300. IEEE, Dec. 2015. [4](#)
- [10] D. Campbell and L. Petersson. GOGMA: Globally-Optimal Gaussian Mixture Alignment. In *Proceedings of the 2016 Conference on Computer Vision and Pattern Recognition*, pages 5685–5694. IEEE, June 2016. [2](#), [3](#)
- [11] D. Campbell, L. Petersson, L. Kneip, and H. Li. Globally-optimal inlier set maximisation for simultaneous camera pose and feature correspondence. In *Proceedings of the 2017 International Conference on Computer Vision*, pages 1–10. IEEE, Oct. 2017. [2](#)
- [12] D. Campbell, L. Petersson, L. Kneip, and H. Li. Globally-optimal inlier set maximisation for camera pose and correspondence estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, preprint, June 2018. [1](#), [2](#), [5](#), [6](#), [7](#)
- [13] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the 2018 European Conference on Computer Vision*, Sept. 2018. [4](#)
- [14] G. K. Cheung, S. Baker, and T. Kanade. Visual hull alignment and refinement across time: A 3D reconstruction algorithm combining shape-from-silhouette with stereo. In *Proceedings of the 2003 Conference on Computer Vision and Pattern Recognition*, volume 2, pages II–375. IEEE, June 2003. [1](#)
- [15] H. Chui and A. Rangarajan. A feature registration framework using mixture models. In *Proceedings of the 2000 Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 190–197. IEEE, June 2000. [3](#)
- [16] P. David, D. Dementhon, R. Duraiswami, and H. Samet. SoftPOSIT: simultaneous pose and correspondence determination. *International Journal of Computer Vision*, 59(3):259–284, 2004. [1](#), [2](#), [6](#)
- [17] L. Devroye. *A course in density estimation*. Progress in Probability and Statistics. Birkhäuser Boston Inc., 1987. [2](#), [4](#)
- [18] O. Enqvist, E. Ask, F. Kahl, and K. Åström. Tractable algorithms for robust model estimation. *International Journal of Computer Vision*, 112(1):115–129, 2015. [2](#)
- [19] O. Enqvist and F. Kahl. Robust optimal pose estimation. In *Proceedings of the 2008 European Conference on Computer Vision*, pages 141–153. Springer, Oct. 2008. [2](#)
- [20] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. [1](#), [2](#), [6](#)
- [21] R. Fisher. Dispersion on a sphere. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 217, pages 295–305. Royal Society, May 1953. [3](#)
- [22] S. Gopal and Y. Yang. Von Mises-Fisher clustering models. In T. Jebara and E. P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 154–162. PMLR, June 2014. [4](#)
- [23] W. E. L. Grimson. *Object Recognition by Computer: The Role of Geometric Constraints*. MIT Press, Cambridge, MA, USA, 1990. [2](#)
- [24] R. I. Hartley and F. Kahl. Global optimization through rotation space search. *International Journal of Computer Vision*, 82(1):64–79, Apr. 2009. [2](#)
- [25] J. A. Hesch and S. I. Roumeliotis. A direct least-squares (DLS) method for PnP. In *Proceedings of the 2011 International Conference on Computer Vision*, pages 383–390. IEEE, Nov. 2011. [1](#), [2](#)
- [26] Q. Huang, W. Wang, and U. Neumann. Recurrent slice networks for 3D segmentation of point clouds. In *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2626–2635, 2018. [4](#)
- [27] B. Jian and B. C. Vemuri. Robust point set registration using Gaussian mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1633–1645, 2011. [3](#), [4](#)
- [28] A. Kendall and R. Cipolla. Geometric loss functions for camera pose regression with deep learning. In *Proceedings of the 2017 Conference on Computer Vision and Pattern Recognition*, pages 6555–6564, July 2017. [2](#)

- [29] A. Kendall, M. Grimes, and R. Cipolla. PoseNet: A convolutional network for real-time 6-DOF camera relocalization. In *Proceedings of the 2015 International Conference on Computer Vision*, pages 2938–2946, Dec. 2015. 2
- [30] J. T. Kent. The Fisher–Bingham distribution on the sphere. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 71–80, 1982. 3
- [31] L. Kneip and P. Furgale. OpenGV: A unified and generalized approach to real-time calibrated geometric vision. In *Proceedings of the 2014 International Conference on Robotics and Automation*, pages 1–8. IEEE, June 2014. 6
- [32] L. Kneip, D. Scaramuzza, and R. Siegwart. A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera position and orientation. In *Proceedings of the 2011 Conference on Computer Vision and Pattern Recognition*, pages 2969–2976. IEEE, June 2011. 6
- [33] L. Kneip, Z. Yi, and H. Li. SDICP: Semi-dense tracking based on iterative closest points. In M. W. J. Xianghua Xie and G. K. L. Tam, editors, *Proceedings of the 2015 British Machine Vision Conference*, pages 100.1–100.12. BMVA Press, Sept. 2015. 1
- [34] B. Kulis and M. I. Jordan. Revisiting k-means: New algorithms via Bayesian nonparametrics. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1131–1138. Omnipress, 2012. 4, 6, 7
- [35] A. H. Land and A. G. Doig. An automatic method of solving discrete programming problems. *Econometrica: Journal of the Econometric Society*, pages 497–520, 1960. 2, 4
- [36] L. Landrieu and S. Martin. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition*, Salt Lake City, USA, June 2018. 4
- [37] 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. 1, 2
- [38] H. Li and R. Hartley. The 3D-3D registration problem revisited. *Proceedings of the 2007 International Conference on Computer Vision*, pages 1–8, Oct. 2007. 2, 5
- [39] Y. Li, N. Snavely, D. Huttenlocher, and P. Fua. Worldwide pose estimation using 3D point clouds. In *Proceedings of the 2012 European Conference on Computer Vision*, pages 15–29. Springer-Verlag, Oct. 2012. 2
- [40] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 1
- [41] E. Marchand, H. Uchiyama, and F. Spindler. Pose estimation for augmented reality: a hands-on survey. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2633–2651, 2016. 1
- [42] K. Mardia. *Statistics of Directional Data*. Probability and Mathematical Statistics. Academic Press, 1972. 3
- [43] F. Moreno-Noguer, V. Lepetit, and P. Fua. Pose priors for simultaneously solving alignment and correspondence. In *Proceedings of the 2008 European Conference on Computer Vision*, pages 405–418. Springer, Oct. 2008. 1, 2, 6, 7
- [44] C. F. Olson. A general method for geometric feature matching and model extraction. *International Journal of Computer Vision*, 45(1):39–54, 2001. 1
- [45] C. Olsson, F. Kahl, and M. Oskarsson. Branch-and-bound methods for Euclidean registration problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):783–794, 2009. 2
- [46] T. M. Pukkila and C. R. Rao. Pattern recognition based on scale invariant discriminant functions. *Information Sciences*, 45(3):379–389, 1988. 3
- [47] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proceedings of the 2015 International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015. 4
- [48] S. Rota Bulò, L. Porzi, and P. Kotschieder. In-place activated batchnorm for memory-optimized training of DNNs. In *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 5639–5647, Salt Lake City, USA, June 2018. 4
- [49] T. Sattler, B. Leibe, and L. Kobbelt. Fast image-based localization using direct 2D-to-3D matching. In *Proceedings of the 2011 International Conference on Computer Vision*, pages 667–674. IEEE, Nov. 2011. 2
- [50] T. Sattler, B. Leibe, and L. Kobbelt. Efficient effective prioritized matching for large-scale image-based localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1744–1756, Sept. 2017. 2
- [51] D. W. Scott. Parametric statistical modeling by minimum integrated square error. *Technometrics*, 43(3):274–285, 2001. 4
- [52] J. Shotton, B. Glocker, C. Zach, S. Izadi, A. Criminisi, and A. Fitzgibbon. Scene coordinate regression forests for camera relocalization in RGB-D images. In *Proceedings of the 2013 Conference on Computer Vision and Pattern Recognition*, pages 2930–2937. IEEE, June 2013. 2
- [53] J. Straub, T. Campbell, J. P. How, and J. W. Fisher. Small-variance nonparametric clustering on the hypersphere. In *Proceedings of the 2015 Conference on Computer Vision and Pattern Recognition*, pages 334–342. IEEE, June 2015. 4, 6, 7
- [54] J. Straub, T. Campbell, J. P. How, and J. W. Fisher III. Efficient global point cloud alignment using Bayesian nonparametric mixtures. In *Proceedings of the 2017 Conference on Computer Vision and Pattern Recognition*, pages 2403–2412. IEEE, July 2017. 3, 4
- [55] L. Svärm, O. Enqvist, F. Kahl, and M. Oskarsson. City-scale localization for cameras with known vertical direction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(7):1455–1461, 2016. 2
- [56] M. Tatarchenko, J. Park, V. Koltun, and Q.-Y. Zhou. Tangent convolutions for dense prediction in 3D. In *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3887–3896, Salt Lake City, USA, June 2018. 4
- [57] Y. Tsin and T. Kanade. A correlation-based approach to robust point set registration. In *Proceedings of the 2004*

- European Conference on Computer Vision*, pages 558–569. Springer, May 2004. Point-set available at <http://www.cs.cmu.edu/~ytsin/KCReg/KCReg.zip>. 3
- [58] F. Wang and A. E. Gelfand. Directional data analysis under the general projected normal distribution. *Statistical Methodology*, 10(1):113–127, 2013. 3
- [59] G. Watson. *Statistics on Spheres*, volume 6 of *University of Arkansas Lecture Notes in the Mathematical Sciences*. Wiley, 1983. 3
- [60] J. Yang, H. Li, D. Campbell, and Y. Jia. Go-ICP: A globally optimal solution to 3D ICP point-set registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(11):2241–2254, Nov. 2016. 2
- [61] B. Zeisl, T. Sattler, and M. Pollefeys. Camera pose voting for large-scale image-based localization. In *Proceedings of the 2015 International Conference on Computer Vision*, pages 2704–2712. IEEE, Dec. 2015. 2