APPENDIX

A. Proof of the Minimum Distance to a Spherical Cap

Let $\mathbf{x}_i^0 = \mathbf{R}_{\mathbf{r}_0}\mathbf{x}_i$ and let $\mathbf{x}_{ij}^* = \mathbf{R}\mathbf{x}_i$ be the point on the origin-centred spherical cap such that $\angle\left(\mathbf{x}_i^0, \mathbf{x}_{ij}^*\right) = \beta$ and $\mathbf{x}_{ij}^*$ is coplanar with $\mathbf{x}_i^0$ and $\mathbf{y}_j' = (\mathbf{y}_j - \mathbf{t}_0)$, that is $\mathbf{x}_{ij}^* \cdot \left(\mathbf{x}_i^0 \times \mathbf{y}_j'\right) = 0$.

Theorem 3. *(Spherical cap distance) For the spherical cap defined by the vector $(\mathbf{x}_i' + \mathbf{t}_0)$ constrained by $\angle(\mathbf{x}_i', \mathbf{x}_i^0) \leqslant \beta$, the minimum distance from a point $\mathbf{y}_j$ to the spherical cap is given by*

$$\min \|\mathbf{x}_i' + \mathbf{t}_0 - \mathbf{y}_j\| = \begin{cases} \left| \|\mathbf{x}_i\| - \|\mathbf{y}_j - \mathbf{t}_0\| \right| & \text{for } \alpha \leqslant \beta \\ \|\mathbf{x}_{ij}^* - (\mathbf{y}_j - \mathbf{t}_0)\| & \text{for } \alpha > \beta \end{cases} \tag{1}$$

where α and β are shown in Figure A.1 and are given by

$$\alpha = \angle\left(\mathbf{x}_i^0, \mathbf{y}_j - \mathbf{t}_0\right) = \arccos \frac{\mathbf{x}_i^0 \cdot (\mathbf{y}_j - \mathbf{t}_0)}{\|\mathbf{x}_i\| \|\mathbf{y}_j - \mathbf{t}_0\|} \tag{2}$$

$$\beta = \angle\left(\mathbf{x}_i^0, \mathbf{x}_{ij}^*\right) = \min(\sqrt{3}\delta_r, \pi). \tag{3}$$

Proof. Dropping the subscripts and translating everything by $(-\mathbf{t}_0)$, an arbitrary point $\mathbf{x}' = \mathbf{R}\mathbf{x}$ on the spherical cap can be expressed as the rotation of the point $\mathbf{x}^0$ about the sphere centre towards $\mathbf{y}'$ by an angle $\gamma \in [0, \beta]$, followed by a rotation of this intermediate vector (denoted $\mathbf{x}''$) about the axis $\mathbf{x}^0$ by θ, as shown in Figure A.2. Note that $\mathbf{x}''$ is coplanar with $\mathbf{x}^0$ and $\mathbf{y}'$, and $\mathbf{x}^*$ is a special case of $\mathbf{x}''$ when $\gamma = \beta$. Also note that $\|\mathbf{x}^0\| = \|\mathbf{x}'\| = \|\mathbf{x}''\| = \|\mathbf{x}^*\| = \|\mathbf{x}\|$ and $\gamma \in [0, \beta]$ follows from Lemma 3. The first axis of rotation, perpendicular to the plane formed by $\mathbf{x}^0$ and $(\mathbf{y} - \mathbf{t}_0)$, is given by

$$\hat{\mathbf{u}} = \frac{\mathbf{x}^0 \times \mathbf{y}'}{\|\mathbf{x}\| \|\mathbf{y}'\| \sin\alpha}. \tag{4}$$

Therefore, by the Rodrigues' rotation formula,

$$\mathbf{x}'' = \mathbf{x}^0 \cos\gamma + \left(\hat{\mathbf{u}} \times \mathbf{x}^0\right) \sin\gamma + \hat{\mathbf{u}}(\hat{\mathbf{u}} \cdot \mathbf{x}^0)(1 - \cos\gamma) \tag{5}$$

$$= \mathbf{x}^0 \cos\gamma + \left(\hat{\mathbf{u}} \times \mathbf{x}^0\right) \sin\gamma \tag{6}$$

$$= \mathbf{x}^0 \cos\gamma + \frac{(\mathbf{x}^0 \cdot \mathbf{x}^0)\mathbf{y}' - (\mathbf{x}^0 \cdot \mathbf{y}')\mathbf{x}^0}{\|\mathbf{x}\| \|\mathbf{y}'\| \sin\alpha} \sin\gamma \tag{7}$$

$$= \frac{\sin(\alpha - \gamma)}{\sin\alpha}\mathbf{x}^0 + \frac{\|\mathbf{x}\| \sin\gamma}{\|\mathbf{y}'\| \sin\alpha}\mathbf{y}' \tag{8}$$

where (6) follows, after substituting in (4), from the result that the scalar triple product is zero if any two vectors involved are equal, (7) follows from a vector triple product identity and (8) follows by expanding, simplifying and using $\mathbf{x}^0 \cdot \mathbf{y}' = \|\mathbf{x}\| \|\mathbf{y}'\| \cos\alpha$. Therefore, by the Rodrigues' rotation formula,

$$\mathbf{x}' = \mathbf{x}'' \cos\theta + \left(\frac{\mathbf{x}^0}{\|\mathbf{x}\|} \times \mathbf{x}''\right) \sin\theta + \frac{\mathbf{x}^0}{\|\mathbf{x}\|}\left(\frac{\mathbf{x}^0}{\|\mathbf{x}\|} \cdot \mathbf{x}''\right)(1 - \cos\theta) \tag{9}$$

$$= \left(\cos\gamma - \frac{\cos\alpha \sin\gamma \cos\theta}{\sin\alpha}\right)\mathbf{x}^0 + \frac{\sin\gamma}{\sin\alpha}\frac{\mathbf{x}^0 \times \mathbf{y}'}{\|\mathbf{y}'\|} + \frac{\sin\gamma \cos\theta}{\sin\alpha}\frac{\|\mathbf{x}\|}{\|\mathbf{y}'\|}\mathbf{y}' \tag{10}$$

where (10) follows from substituting in (8), expanding and simplifying. Now, the squared distance between point $\mathbf{y}'$ and an arbitrary point on the spherical cap is given by

$$\|\mathbf{x}' - \mathbf{y}'\|^2 = (\mathbf{x}' - \mathbf{y}') \cdot (\mathbf{x}' - \mathbf{y}') \tag{11}$$

$$= \mathbf{x}' \cdot \mathbf{x}' + \mathbf{y}' \cdot \mathbf{y}' - 2\mathbf{x}' \cdot \mathbf{y}' \tag{12}$$

$$= \|\mathbf{x}\|^2 + \|\mathbf{y}'\|^2 - 2\left(\cos\gamma - \frac{\cos\alpha \sin\gamma \cos\theta}{\sin\alpha}\right)\mathbf{x}^0 \cdot \mathbf{y}' - 2\frac{\|\mathbf{x}\| \sin\gamma \cos\theta}{\|\mathbf{y}'\| \sin\alpha}\mathbf{y}' \cdot \mathbf{y}' \tag{13}$$

$$= \|\mathbf{x}\|^2 + \|\mathbf{y}'\|^2 - 2\left(\cos\alpha\cos\gamma - \frac{\cos^2\alpha \sin\gamma \cos\theta}{\sin\alpha}\right)\|\mathbf{x}\| \|\mathbf{y}'\| - 2\frac{\sin\gamma\cos\theta}{\sin\alpha}\|\mathbf{x}\| \|\mathbf{y}'\| \tag{14}$$

$$= \|\mathbf{x}\|^2 + \|\mathbf{y}'\|^2 - 2\left(\cos\alpha\cos\gamma + \sin\alpha\sin\gamma\cos\theta\right)\|\mathbf{x}\| \|\mathbf{y}'\| \tag{15}$$

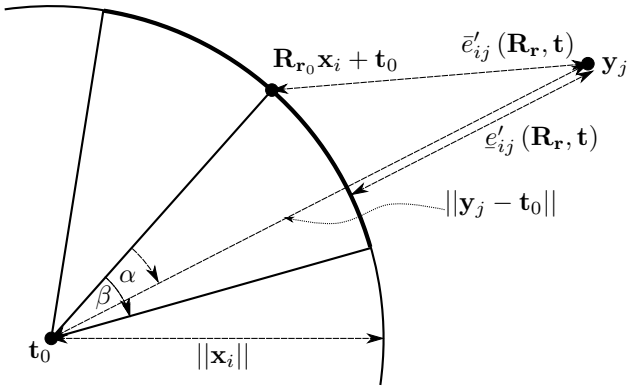

(a) Case 1: $\mathbf{y}_j$ is within the rotation cone ($\alpha \leqslant \beta$).

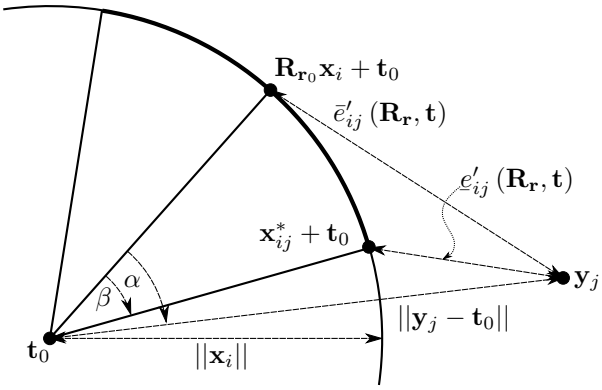

(b) Case 2: $\mathbf{y}_j$ is outside the rotation cone ($\alpha > \beta$).

Figure A.1. Upper and lower bounds of the pairwise residual error, neglecting translation. A 2D cross-section in the plane defined by points $\{\mathbf{R}_{\mathbf{r}_0}\mathbf{x}_i + \mathbf{t}_0, \mathbf{y}_j, \mathbf{t}_0\}$ is shown. The spherical cap cross-section is depicted as a bold curve.

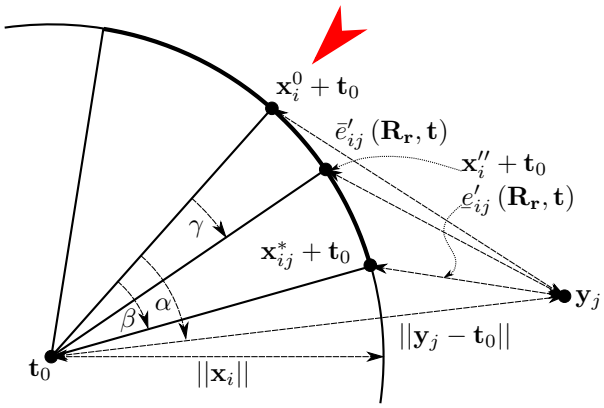

(a) Viewpoint A: the arbitrary vector $(\mathbf{x}''_i + \mathbf{t}_0)$ lies on the spherical cap and is coplanar with $(\mathbf{x}^0_i + \mathbf{t}_0)$, $\mathbf{y}_j$ and $\mathbf{t}_0$. $\mathbf{x}''_i$ is defined as the rotation of the vector $\mathbf{x}^0_i$ towards $(\mathbf{y}_j - \mathbf{t}_0)$ by the angle $\gamma \in [0, \beta]$.

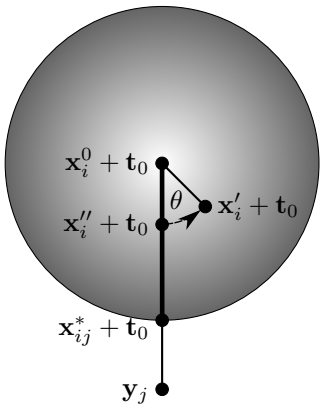

(b) Viewpoint B: the arbitrary vector $(\mathbf{x}'_i + \mathbf{t}_0)$ lies on the spherical cap and is not coplanar with the other points. $\mathbf{x}'_i$ is defined as the rotation of the vector $\mathbf{x}''_i$ about the axis $\mathbf{x}^0_i$ by the angle θ.

Figure A.2. Constructions for the proof of Theorem 3. (a) The spherical cap (the bold curve) is viewed in the plane defined by points $\{\mathbf{R}_{\mathbf{r}_0}\mathbf{x}_i + \mathbf{t}_0, \mathbf{y}_j, \mathbf{t}_0\}$. (b) The spherical cap (the grey shaded circle) is viewed in the direction of the large red arrow in (a).

where (13) follows from substituting in (10) and noting that the scalar triple product is zero if any two vectors involved are equal and (15) follows from the identity $\cos^2\alpha = 1 - \sin^2\alpha$. The squared distance is minimised when $\theta = 0$ and is given by

$$\min \|\mathbf{x}' - \mathbf{y}'\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}'\|^2 - 2\left(\cos\alpha\cos\gamma + \sin\alpha\sin\gamma\right)\|\mathbf{x}\|\|\mathbf{y}'\| \tag{16}$$

$$= \|\mathbf{x}\|^2 + \|\mathbf{y}'\|^2 - 2\cos(\alpha - \gamma)\|\mathbf{x}\|\|\mathbf{y}'\|. \tag{17}$$

When $\alpha \leqslant \beta$ (Case 1), (17) is minimised when $\gamma = \alpha$, giving

$$\min \|\mathbf{x}' - \mathbf{y}'\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}'\|^2 - 2\|\mathbf{x}\|\|\mathbf{y}'\| \tag{18}$$

$$= \left(\|\mathbf{x}\| - \|\mathbf{y}'\|\right)^2 \tag{19}$$

$$\therefore \min \|\mathbf{x}' - \mathbf{y}'\| = \left|\|\mathbf{x}\| - \|\mathbf{y}'\|\right| \text{ for } \alpha \leqslant \beta. \tag{20}$$

When $\alpha > \beta$ (Case 2), (17) is minimised when $\gamma = \beta$, giving

$$\min \|\mathbf{x}' - \mathbf{y}'\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}'\|^2 - 2\cos(\alpha - \beta)\|\mathbf{x}\|\|\mathbf{y}'\| \tag{21}$$

$$= \|\mathbf{x}^* - \mathbf{y}'\|^2 \tag{22}$$

$$\therefore \min \|\mathbf{x}' - \mathbf{y}'\| = \|\mathbf{x}^* - \mathbf{y}'\| \text{ for } \alpha > \beta. \tag{23}$$

□

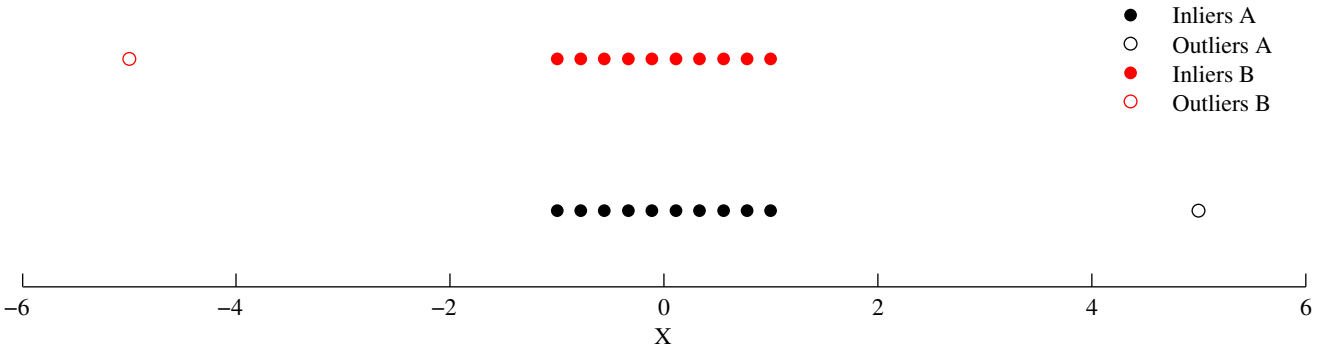

(a) Two 1D point-sets, both with a single outlier

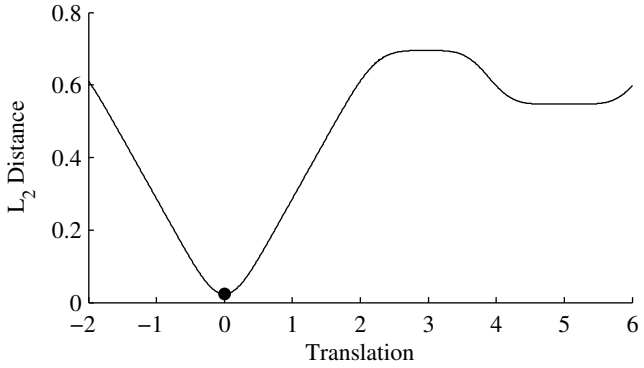

(b) L_2 Distance and L_2E Estimator ($\sigma = 0.2$)

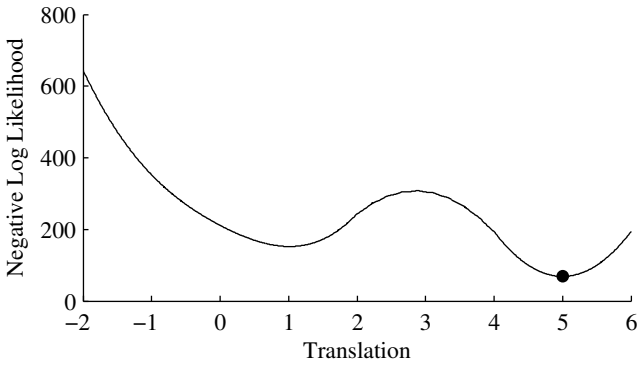

(c) Negative Log-Likelihood and Maximum Likelihood Estimator ($\sigma = 0.2$)

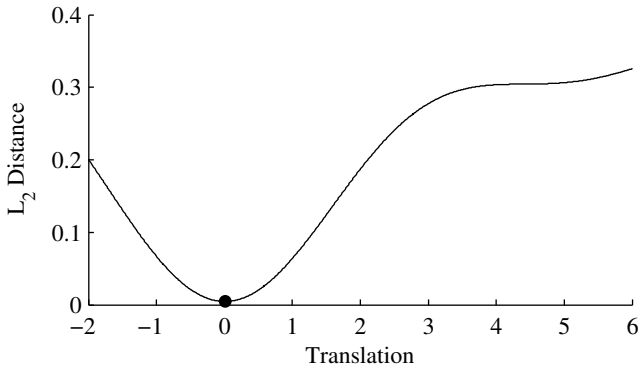

(d) L_2 Distance and L_2E Estimator ($\sigma = 1.0$)

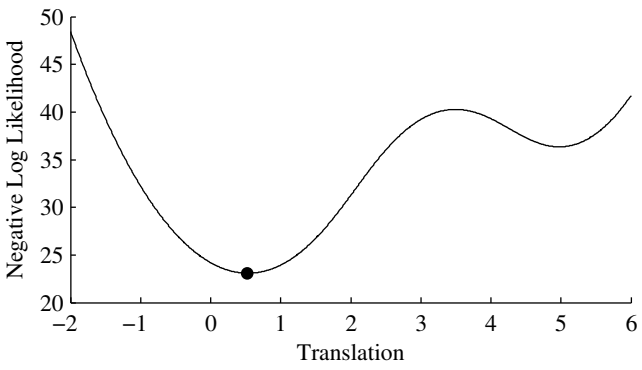

(e) Negative Log-Likelihood and Maximum Likelihood Estimator ($\sigma = 1.0$)

Figure B.3. Toy example demonstrating the robustness of the L_2E estimator. (a) Two 1D point-sets A and B which overlap exactly, except for a single outlier in each. As point-set B translates with respect to point-set A, the L_2 distance between Gaussian mixtures (constructed from the point-sets using kernel density estimation) and the negative log-likelihood is evaluated and plotted for different scales σ. (b) At a scale of $\sigma = 0.2$, the L_2E estimator is globally-optimal and multiple local minima exist. (c) At the same scale (and below), the Maximum Likelihood Estimator (MLE) is severely biased by the outliers and finds the incorrect translation. It also has multiple local minima. (d, e) At larger scales (such as $\sigma = 1.0$), the MLE is still biased, but less so. As the scale increases further, both estimates converge towards aligning the centres-of-mass of the point-sets. The translation estimates (L_2E and MLE) and their values are marked as black dots.

B. Robustness of the L_2E Estimator

The robustness of the L_2E estimator to outliers has been demonstrated both empirically and from its connection with M-estimators [1, 6]. While counter-intuitive, it arises from the Gaussian attenuation of outlying values. Scott [6] demonstrates that L_2E is an "inherently robust" estimator that has the advantage of not requiring any additional tuning factors, unlike many other robust functions. A toy example demonstrating that the L_2E estimator is not biased by systematic outliers, unlike the estimator from Maximum Likelihood Estimation (MLE), is shown in Figure B.3. In contrast, Figure B.4 shows that, in the absence of outliers, the alignment task is adequately (and perhaps preferably) handled by MLE.

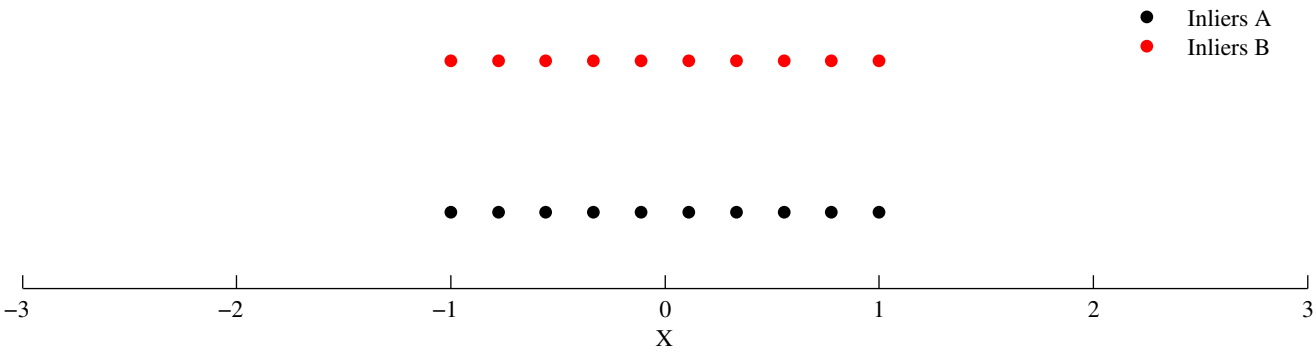

(a) Two 1D point-sets, both without any outliers

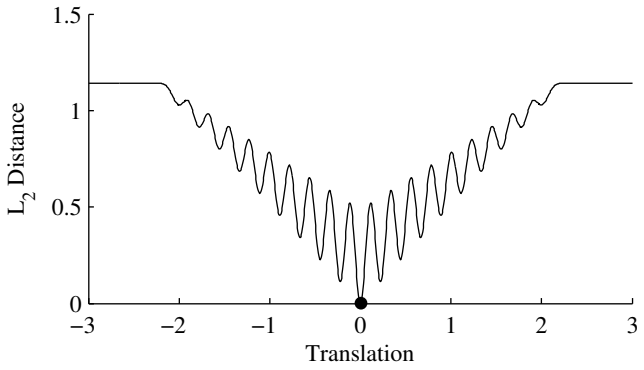

(b) L_2 Distance and L_2E Estimator ($\sigma = 0.05$)

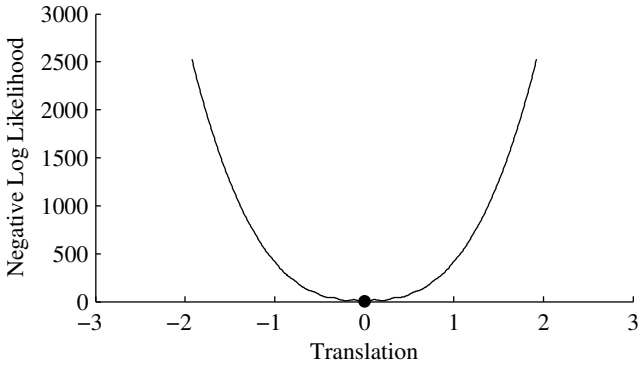

(c) Negative Log-Likelihood and Maximum Likelihood Estimator ($\sigma = 0.05$)

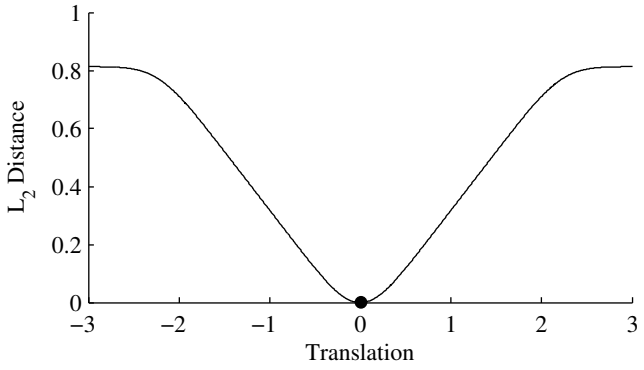

(d) L_2 Distance and L_2E Estimator ($\sigma = 0.2$)

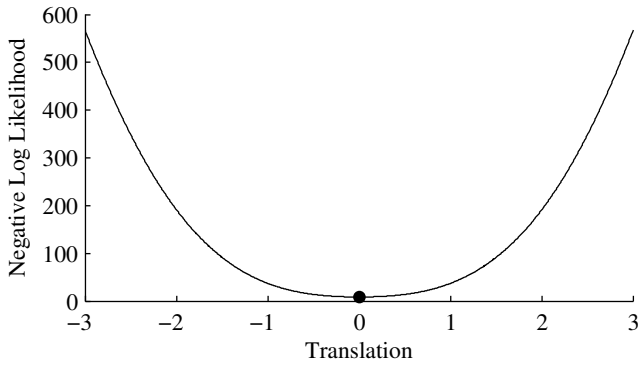

(e) Negative Log-Likelihood and Maximum Likelihood Estimator ($\sigma = 0.2$)

Figure B.4. Toy example demonstrating the alignment of two point-sets without outliers. (a) Two 1D point-sets A and B which overlap exactly. As point-set B translates with respect to point-set A, the L_2 distance between Gaussian mixtures (constructed from the point-sets using kernel density estimation) and the negative log-likelihood is evaluated and plotted for different scales σ. (b) At a scale of $\sigma = 0.05$, the L_2E estimator is globally-optimal and the profile of the L_2 distance function contains many local minima. (c) At the same scale, the MLE estimator is also globally-optimal, but the negative log-likelihood profile has fewer, much shallower local minima. (d, e) At larger scales ($\sigma = 0.2$), both profiles smooth out and both still find the global-optimum. The translation estimates (L_2E and MLE) and their values are marked as black dots.

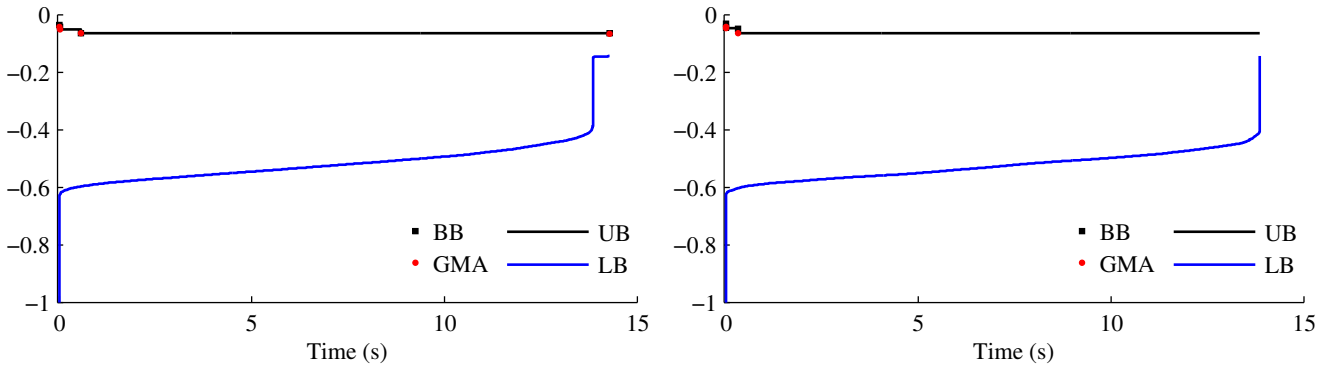

(a) DRAGON (b) BUNNY

Figure C.5. Evolution of the upper and lower bounds for the reconstructed DRAGON and BUNNY models. The normalised objective function value is plotted against time.

C. Additional Experimental Results

In this section, we present additional figures from the experimental results section. In Figure C.5, Figure 4 has been re-plotted at a larger scale without using a logarithmic scale. This scale was necessary to visualise the conceptually critical BB and GMA steps, which were almost invisible in the figure otherwise. However, the bound evolution is clearer in Figure C.5.

In Figure C.6, the results from Tables 3 and 4 are shown as box-plots. It can be clearly seen that GOGMA [*] and GOGMA with refinement $[*]_+$ outperform Go-ICP [7] with a loose ϵ threshold $[52]_\text{a}$ and a tighter ϵ threshold $[52]_\text{b}$, particularly with regard to robustness. That is, GOGMA produced few outliers, all of which are in the vicinity of the correct transformation. In contrast, Go-ICP produced many outliers, most of which are incorrect even by the coarsest success criterion. This is likely due to the partially-overlapping nature of the point-sets. Under the Go-ICP framework, trimming would be required to handle these missing correspondences. However, any trimming made the runtime prohibitive for these datasets.

References

[1] A. Basu, I. R. Harris, N. L. Hjort, and M. Jones. Robust and efficient estimation by minimising a density power divergence. *Biometrika*, 85(3):549–559, 1998. 3

[2] D. Campbell and L. Petersson. An adaptive data representation for robust point-set registration and merging. In *Proc. 2015 Int. Conf. Comput. Vision*, pages 4292–4300, Santiago, Chile, 2015. IEEE.

[3] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, Jan. 1977.

[4] M. A. Figueiredo and A. K. Jain. Unsupervised learning of finite mixture models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(3):381–396, 2002.

[5] B. Jian and B. C. Vemuri. Robust point set registration using Gaussian mixture models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(8):1633–1645, 2011.

[6] D. W. Scott. Parametric statistical modeling by minimum integrated square error. *Technometrics*, 43(3):274–285, 2001. 3

[7] J. Yang, H. Li, and Y. Jia. GoICP: Solving 3D registration efficiently and globally optimally. In *Proc. 2013 Int. Conf. Comput. Vision*, pages 1457–1464, Sydney, Australia, 2013. IEEE. 5

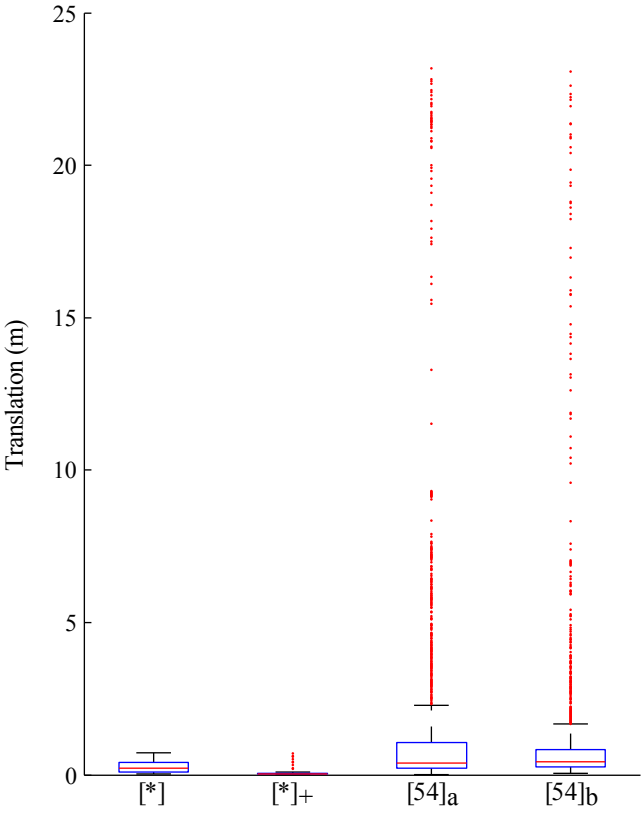

(a) Translation STAIRS

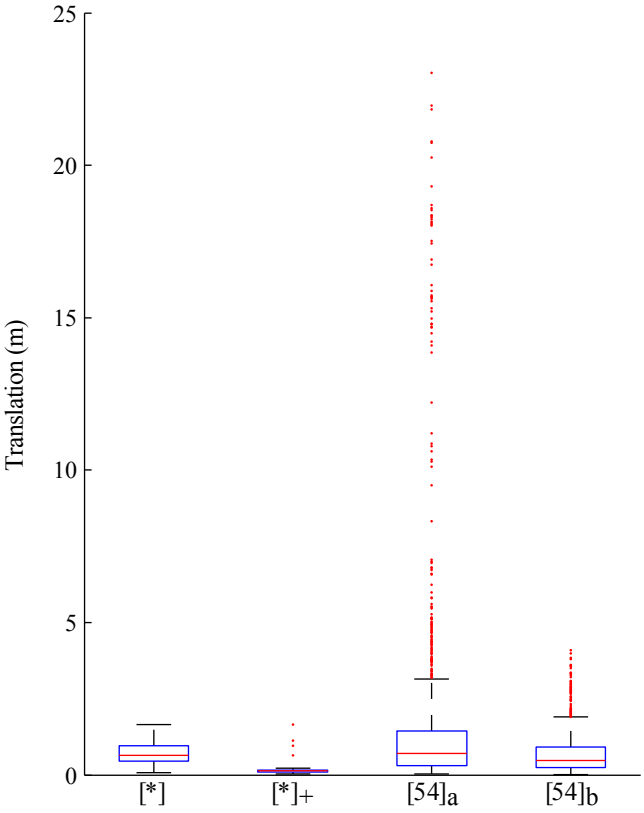

(b) Translation WOOD-SUMMER

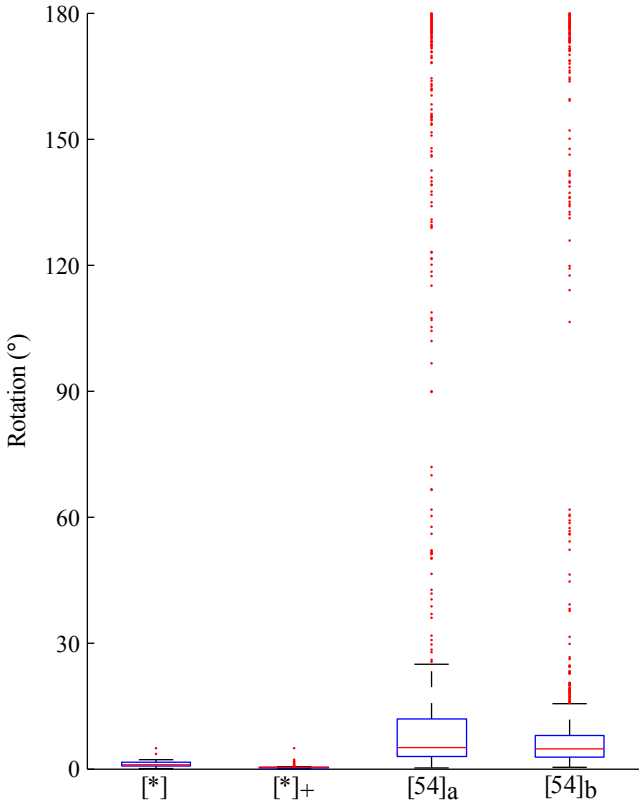

(c) Rotation STAIRS

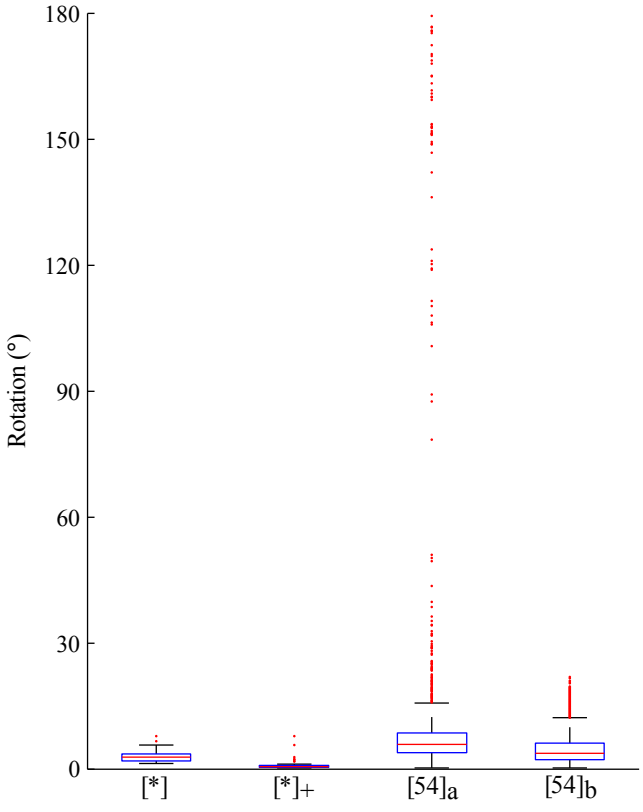

(d) Rotation WOOD-SUMMER

Figure C.6. Box plots of the translation and rotation errors for the STAIRS and WOOD-SUMMER datasets.