Certifiably Optimal Anisotropic Rotation Averaging

Supplementary Material

A. First-Order Uncertainty Propagation

We leverage Laplace's approximation, which allows us to identify the Hessian of a non-linear least-squares minimization task as the precision matrix of the posterior $p(\tilde{\boldsymbol{\omega}}_{ij}|\text{image } i \text{ and } j)$ (up to a positive scale). In the following we apply first-order uncertainty propagation to move from the minimal axis-angle parametrization $\boldsymbol{\omega} \in \mathbb{R}^3$ to a direct parametrization of rotation matrices $R_{ij} \in \mathbb{R}^{3\times 3}$.

For $\Delta\boldsymbol{\omega} \approx \mathbf{0}$ we have $\exp([\Delta\boldsymbol{\omega}]_\times) \approx \mathtt{I} + [\Delta\boldsymbol{\omega}]_\times$. The Jacobian of the map $\Delta\boldsymbol{\omega} \to \operatorname{vec}(\mathtt{I} + [\Delta\boldsymbol{\omega}]_\times)$ is the 9x3 matrix

$$\mathtt{J} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Therefore the Gaussian transformed by the mapping $\Delta\boldsymbol{\omega} \to \operatorname{vec}(\mathtt{I} + [\Delta\boldsymbol{\omega}]_\times)$ has covariance matrix

$$\Sigma_{[\Delta\boldsymbol{\omega}]_\times} = \mathtt{J}\Sigma_{\Delta\boldsymbol{\omega}}\mathtt{J}^\top = \mathtt{J}\mathtt{H}^{-1}\mathtt{J}^\top. \tag{A.1}$$

Since $\mathtt{J}$ has rank 3, the above relation only fixes 6 out of the 45 degrees of freedom in $\Sigma_{[\Delta\boldsymbol{\omega}]_\times}$, hence there are 39 d.o.f. available in $\Sigma_{[\Delta\boldsymbol{\omega}]_\times}$ to achieve the following properties: (i) $\Sigma_{[\Delta\boldsymbol{\omega}]_\times}$ is invertible and (ii) the maximum-likelihood objective

$$\begin{aligned} &\operatorname{vec}(R - \tilde{R})^\top \mathtt{M} \operatorname{vec}(R - \tilde{R}) \\ &\doteq \operatorname{vec}(R)^\top \mathtt{M} \operatorname{vec}(R) - 2\operatorname{vec}(\tilde{R})^\top \mathtt{M} \operatorname{vec}(R) \end{aligned} \tag{A.2}$$

is *linear* in R and matches $\Delta\boldsymbol{\omega}^\top \mathtt{H} \Delta\boldsymbol{\omega}$ to first order for a suitable matrix $\mathtt{M}$ to be determined. We consider

$$\begin{aligned} \Delta\boldsymbol{\omega}^\top \mathtt{H} \Delta\boldsymbol{\omega} &\overset{!}{=} \operatorname{tr}([\Delta\boldsymbol{\omega}]_\times^\top \mathtt{M} [\Delta\boldsymbol{\omega}]_\times) \\ &= \operatorname{vec}([\Delta\boldsymbol{\omega}]_\times)^\top \operatorname{vec}(\mathtt{M}[\Delta\boldsymbol{\omega}]_\times) \\ &= (\mathtt{J}\Delta\boldsymbol{\omega})^\top \operatorname{vec}(\mathtt{M}[\Delta\boldsymbol{\omega}]_\times) \\ &= \Delta\boldsymbol{\omega}^\top \mathtt{J}^\top (\mathtt{I} \otimes \mathtt{M}) \operatorname{vec}([\Delta\boldsymbol{\omega}]_\times) \\ &= \Delta\boldsymbol{\omega}^\top \mathtt{J}^\top (\mathtt{I} \otimes \mathtt{M}) \mathtt{J} \Delta\boldsymbol{\omega} \end{aligned} \tag{A.3}$$

for all $\Delta\boldsymbol{\omega}$, which implies

$$\mathtt{J}^\top (\mathtt{I} \otimes \mathtt{M}) \mathtt{J} = \mathtt{J}^\top \begin{pmatrix} \mathtt{M} & & \\ & \mathtt{M} & \\ & & \mathtt{M} \end{pmatrix} \mathtt{J} = \mathtt{H}. \tag{A.4}$$

By e.g. using a CAS the relations between $\mathtt{M}$ and $\mathtt{H}$ can be derived as

$$\mathtt{M}_{11} + \mathtt{M}_{22} = \mathtt{H}_{33} \quad \mathtt{M}_{11} + \mathtt{M}_{33} = \mathtt{H}_{22} \quad \mathtt{M}_{22} + \mathtt{M}_{33} = \mathtt{H}_{11} \tag{A.5}$$

and $\mathtt{M}_{ij} = -\mathtt{H}_{ij}$ for $i \neq j$. These relations can be more compactly written as

$$\mathtt{H} = \operatorname{tr}(\mathtt{M})\mathtt{I} - \mathtt{H} \tag{A.6}$$

in accordance with Section 4 of the main paper. We relate $[\Delta\boldsymbol{\omega}]_\times$ and R via $[\Delta\boldsymbol{\omega}]_\times \approx R\tilde{R}^\top - \mathtt{I}$ and therefore obtain

$$\begin{aligned} \operatorname{tr}([\Delta\boldsymbol{\omega}]_\times^\top \mathtt{M} [\Delta\boldsymbol{\omega}]_\times) &\approx \operatorname{tr}\Big((R\tilde{R}^\top - \mathtt{I})^\top \mathtt{M} (R\tilde{R}^\top - \mathtt{I})\Big) \\ &= \operatorname{tr}(\tilde{R}R^\top \mathtt{M} R\tilde{R}^\top) - 2\operatorname{tr}(\mathtt{M}R\tilde{R}^\top) + \operatorname{tr}(\mathtt{M}) \\ &= \operatorname{tr}(\mathtt{M}R\tilde{R}^\top \tilde{R}R^\top) - 2\operatorname{tr}(\tilde{R}^\top \mathtt{M} R) + \operatorname{tr}(\mathtt{M}) \\ &= 2\operatorname{tr}(\mathtt{M}) - 2\operatorname{tr}(\tilde{R}^\top \mathtt{M} R) \\ &= 2\operatorname{tr}(\mathtt{M}) - 2\left\langle \mathtt{M}\tilde{R}, R \right\rangle_F, \end{aligned} \tag{A.7}$$

which is equivalent to the cost (9) in the main text.

Since $[\Delta\boldsymbol{\omega}]_\times$ is (skew-)symmetric, a seemingly different solution $\mathtt{M}'$ can be obtained by factoring $\operatorname{tr}([\Delta\boldsymbol{\omega}]_\times^\top \mathtt{M} [\Delta\boldsymbol{\omega}]_\times)$ differently,

$$\begin{aligned} \Delta\boldsymbol{\omega}^\top \mathtt{H} \Delta\boldsymbol{\omega} &\overset{!}{=} \operatorname{tr}([\Delta\boldsymbol{\omega}]_\times^\top \mathtt{M}' [\Delta\boldsymbol{\omega}]_\times) \\ &= \operatorname{vec}(\mathtt{M}'[\Delta\boldsymbol{\omega}]_\times)^\top \operatorname{vec}([\Delta\boldsymbol{\omega}]_\times) \\ &= \operatorname{vec}(\mathtt{M}'[\Delta\boldsymbol{\omega}]_\times)^\top \mathtt{J}\Delta\boldsymbol{\omega} \\ &= \left((\mathtt{M}' \otimes \mathtt{I}) \operatorname{vec}([\Delta\boldsymbol{\omega}]_\times)\right)^\top \mathtt{J}\Delta\boldsymbol{\omega} \\ &= \Delta\boldsymbol{\omega}^\top \mathtt{J}^\top (\mathtt{M}' \otimes \mathtt{I}) \mathtt{J} \Delta\boldsymbol{\omega} \end{aligned} \tag{A.8}$$

for all $\Delta\boldsymbol{\omega}$, leading to the condition $\mathtt{J}^\top(\mathtt{M}' \otimes \mathtt{I})\mathtt{J} = \mathtt{H}$. Using a CAS it can be seen that these conditions on $\mathtt{M}'$ are the same as for $\mathtt{M}$ in (A.5), and therefore $\mathtt{M} = \mathtt{M}'$.

B. Parameterization of R

There are different ways of parametrizing rotations. In the main paper we use the $R = e^{[\Delta\boldsymbol{\omega}]_\times}\tilde{R}$ which lead to the objective of the form $-\langle M\tilde{R}, R\rangle$. In this section, we investigate the effects of switching to other possible parameterizations. In particular, we start by letting $R = e^{[\boldsymbol{\omega}]_\times} = e^{[\tilde{\boldsymbol{\omega}} + \Delta\boldsymbol{\omega}]_\times}$ and look at its approximations.

We first notice that we can expect $R\tilde{R}^\top \approx \mathtt{I}$ as $\tilde{R}$ is a noisy realization of R, therefore R and $\tilde{R}^\top$ approximately commute and

$$\begin{aligned} e^{[\Delta\boldsymbol{\omega}]_\times} &= e^{[\boldsymbol{\omega} - \tilde{\boldsymbol{\omega}}]_\times} = e^{-\alpha[\tilde{\boldsymbol{\omega}}]_\times + [\boldsymbol{\omega}]_\times - (1-\alpha)[\tilde{\boldsymbol{\omega}}]_\times} \\ &\approx e^{-\alpha[\tilde{\boldsymbol{\omega}}]_\times} e^{[\boldsymbol{\omega}]_\times} e^{-(1-\alpha)[\tilde{\boldsymbol{\omega}}]_\times} \\ &= (\tilde{R}^\top)^\alpha R (\tilde{R}^\top)^{1-\alpha} \end{aligned} \tag{B.9}$$

for any $\alpha \in [0, 1]$, i.e.

$$R \approx \tilde{R}^\alpha e^{[\Delta\boldsymbol{\omega}]_\times} \tilde{R}^{(1-\alpha)}. \tag{B.10}$$

The Hessian induced by the two-view optimization can then be computed in accordance with this mapping. The first-order Taylor expansion of (B.9) results in

$$[\Delta\boldsymbol{\omega}]_\times \approx (\tilde{R}^\top)^\alpha R (\tilde{R}^\top)^{1-\alpha} - \mathtt{I}. \tag{B.11}$$

Plugging this into (5) gives the corresponding linear cost

$$-\langle \tilde{R}^\alpha M \tilde{R}^{(1-\alpha)}, R \rangle. \tag{B.12}$$

The natural choices for α are 0, $1/2$ and 1. Setting $\alpha = 0$ yields the formulation employed in the main paper. We found that for the other values of α the solution of anisotropic rotation averaging is the same. However, when using the Hessian matrix computed from the Jacobian of the initial parameterization, $R = e^{[\tilde{\boldsymbol{\omega}}+\Delta\boldsymbol{\omega}]_\times}$, the results vary when optimizing the anisotropic cost for different values of α. In our synthetic experiments the best solution is often obtained with $\alpha = 1/2$, however, this setting does not outperform the proposed formulation overall.

C. Spectral method for anisotropic costs

We recall that our main objective is

$$\sum_{i,j} \left\langle M_{ij} \tilde{R}_{ij}, R_j R_i^\top \right\rangle. \tag{C.13}$$

Let us assume the noisefree setting, i.e. $\tilde{R}_{ij} = R_i R_j^\top$. The cost matrix $\mathtt{N}$ is then given by

$$\mathtt{N} = \begin{pmatrix} \vdots & & \\ \cdots & M_{ij} R_j R_i^\top & \cdots \\ \vdots & & \end{pmatrix}. \tag{C.14}$$

Consequently,

$$\begin{aligned} \mathtt{NR} &= \begin{pmatrix} \vdots \\ \sum_i M_{ij} R_j R_i^\top R_i \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \sum_i M_{ij} R_j \\ \vdots \end{pmatrix} \\ &= \underbrace{\begin{pmatrix} \ddots & & \\ & \sum_i M_{ij} & \\ & & \ddots \end{pmatrix}}_{=:\mathtt{D}} \mathtt{R}. \end{aligned} \tag{C.15}$$

Hence, $\mathtt{NR} = \mathtt{DR}$ or

$$\mathtt{D}^{-1}\mathtt{NR} = \mathtt{R}, \tag{C.16}$$

and $\mathtt{R}$ can be extracted as the eigenspace corresponding to the eigenvalue of 1. Since $\mathtt{N}$ does not necessarily have rank-3 (even in the noise-free case), we obtain $\mathtt{R}$ as the right singular vectors corresponding to the three smallest singular values of $\mathtt{N} - \mathtt{I}$. We include results for this method in Table D.1 introduced in the next section.

D. Experimental Details

We provide complementing results on synthetic and real datasets below.

D.1. Synthetic experiments

We set up a synthetic graph with three cameras as shown in Figure D.1. The estimated relative rotations corresponding to the black dashed edges are "certain", i.e. $\Delta\boldsymbol{\omega}_{ij}$ are drawn from $\mathcal{N}(0, \varepsilon\mathtt{I})$, where $\varepsilon = 0.001$. The estimated relative rotation of the gray edge has varying uncertainty around one of the three axes, e.g., for the x- (red) axis, the noise covariance is $\mathrm{diag}(\sigma, \varepsilon, \varepsilon)$, where σ varies from 0.01 to 0.3. The effects on the error are shown in Figure D.1. The proposed method relies on certain relative rotations and gives an accurate solution, while the standard isotropic approach is negatively affected by the single noisy relative rotation.

We also present complementing results of the synthetic experiments of Sec 5.1 in the paper. Figure D.2 shows the rotation error histograms and the dependency of the runtime on the number of cameras for the other fractions of the observed relative rotations. These results are in line with the analysis provided in the paper — using the proposed objective leads to lower errors and using the proposed constraints speeds up the employed SDP solver.

D.2. Real experiments

To account for the estimated uncertainties in the evaluation, we compared the methods using the Mahalanobis distance between the axis-angle vectors of the estimated and ground truth rotations. The axis-angle vector $\boldsymbol{\omega}_i$ is distributed according to $\mathcal{N}(\boldsymbol{\omega}_i^*, H_i)$, where $\boldsymbol{\omega}_i^*$ is the ground truth axis-angle vector and $H_i = \sum_{j:\{ij\} \text{ was observed}} H_{ij}$, assuming that all other cameras are fixed. Let $\Delta\boldsymbol{\omega}_i^- = \boldsymbol{\omega}_i - \boldsymbol{\omega}_i^*$ and

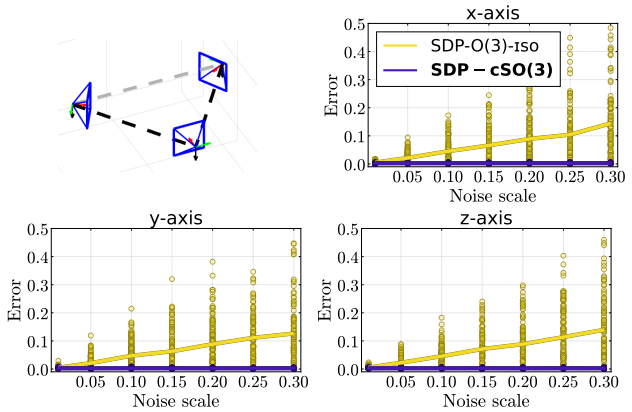

Figure D.1. The configuration of the cameras (top left). Rotation error wrt. ground truth $\sqrt{\sum_i \|R_i - R_i^*\|_F^2}$ (top right, bottom) for the increasing noise scale around one of the axes.

$\Delta\boldsymbol{\omega}_i^+ = \boldsymbol{\omega}_i + \boldsymbol{\omega}_i^*$, then the Mahalanobis error is

$$\sqrt{\sum_i \min\{\Delta\boldsymbol{\omega}_i^{-\top} H_i \Delta\boldsymbol{\omega}_i^-, \Delta\boldsymbol{\omega}_i^{+\top} H_i \Delta\boldsymbol{\omega}_i^+\}}, \quad \text{(D.17)}$$

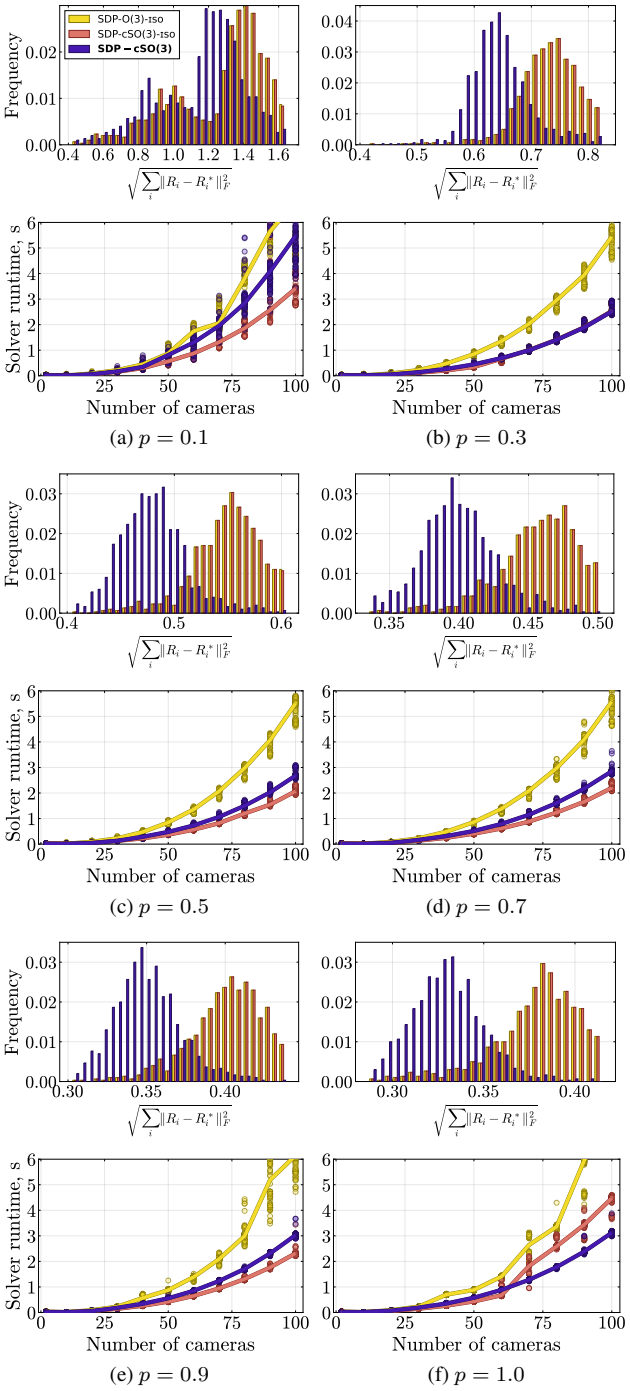

Figure D.2. Histograms of rotation errors wrt. ground truth (rows 1,3,5). Corresponding solver runtime (s) wrt. the increasing number of cameras (rows 2,4,6). Results are shown for different fractions p of observed relative rotations.

where the minimization is done to account for the sign ambiguity in the axis-angle representation. We also present the RMS angular errors for a better geometric interpretation. As shown in Table D.1, in many cases, the proposed method leads to a much lower error.

Dataset	Method	Mahal. err.	Angl. err.
LU Sphinx	SDP-O(3)-iso	0.388	0.46
	Spectral	0.420	1.17
	SDP-cSO(3)	**0.207**	**0.36**
Round Church	SDP-O(3)-iso	0.631	0.59
	Spectral	0.437	1.20
	SDP-cSO(3)	**0.368**	**0.54**
UWO	SDP-O(3)-iso	1.481	1.19
	Spectral	6.125	7.07
	SDP-cSO(3)	**0.727**	**0.86**
Tsar Nikolai I	SDP-O(3)-iso	0.687	0.48
	Spectral	0.344	0.71
	SDP-cSO(3)	**0.188**	**0.22**
Vercingetorix	SDP-O(3)-iso	0.431	1.53
	Spectral	30.970	86.94
	SDP-cSO(3)	**0.423**	**1.42**
Eglise Du Dome	SDP-O(3)-iso	0.224	0.24
	Spectral	**0.119**	0.22
	SDP-cSO(3)	0.188	**0.21**
King's College	SDP-O(3)-iso	0.229	0.76
	Spectral	0.251	1.00
	SDP-cSO(3)	**0.130**	**0.37**
Kronan	SDP-O(3)-iso	**0.738**	**0.76**
	Spectral	2.622	4.36
	SDP-cSO(3)	1.111	1.38
Alcatraz	SDP-O(3)-iso	1.333	0.62
	Spectral	**0.667**	0.80
	SDP-cSO(3)	1.011	**0.45**
Museum Barcelona	SDP-O(3)-iso	2.710	0.79
	Spectral	16.588	7.35
	SDP-cSO(3)	**1.216**	**0.46**
Temple Singapore	SDP-O(3)-iso	2.420	0.86
	Spectral	**0.719**	**0.46**
	SDP-cSO(3)	1.076	0.55

Table D.1. Mahalanobis distance between the axis-angle vectors of the estimated and ground truth rotations and RMS angular errors (degrees) evaluated on the real datasets.