

Lecture 16: Noise Models, Optimization, and Estimation

Lecturer: Luca Carlone

Scribes: -

Disclaimer: *These notes have not been subjected to the usual scrutiny reserved for formal publications. They may be distributed outside this class only with the permission of the Instructor(s).*

This lecture discusses:

- how to model uncertainty and noise on positions, rotations, and poses.
- the use of optimization to mitigate the presence of measurement noise and the fact that one can choose different objective functions in these optimization problems, and
- introduces Maximum Likelihood and Maximum a Posteriori estimation which provide a probabilistically grounded way to select objective functions given assumptions on the measurement noise.

An introduction to (nonlinear) estimation is given in [1, Chapter 4]. A good introduction to uncertainty modeling for poses and rotations is [1, p. 255-283].

16.1 Representing uncertainty on positions, rotations, and poses

16.1.1 Representing uncertainty on positions

As we saw, positions and translations live in a vector space, hence we can use any multivariate distribution from traditional statistics to model an uncertain vector.

A particularly popular choice is to use a multivariate Gaussian or Normal distribution to model an uncertain vector in \mathbb{R}^d :

$$\mathbf{t} \sim \mathcal{N}(\bar{\mathbf{t}}, \Sigma) \doteq \frac{1}{\sqrt{(2\pi)^d \det(\Sigma)}} \exp\left(-\frac{1}{2}(\mathbf{t} - \bar{\mathbf{t}})^\top \Sigma^{-1}(\mathbf{t} - \bar{\mathbf{t}})\right) \quad (16.1)$$

where $\bar{\mathbf{t}}$ is the *mean* of the distribution and $\Sigma \in \mathbb{R}^{d \times d}$ is the *covariance* matrix. The inverse of Σ , is often called the *information matrix* $\Omega \doteq \Sigma^{-1}$.

Theorem 1 (Sum of uncertain vectors). *Given two Normally distributed vectors $\mathbf{t}_1 \sim \mathcal{N}(\bar{\mathbf{t}}_1, \Sigma_1)$ and $\mathbf{t}_2 \sim \mathcal{N}(\bar{\mathbf{t}}_2, \Sigma_2)$, it holds:*

$$\mathbf{t}_1 + \mathbf{t}_2 \sim \mathcal{N}(\bar{\mathbf{t}}_1 + \bar{\mathbf{t}}_2, \Sigma_1 + \Sigma_2) \quad (16.2)$$

As a consequence of Theorem 1, for a fixed $\bar{\mathbf{t}}$, we can write $\mathbf{t} \sim \mathcal{N}(\bar{\mathbf{t}}, \Sigma)$ equivalently as:

$$\mathbf{t} = \bar{\mathbf{t}} + \boldsymbol{\epsilon}_t \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}_d, \Sigma_t) \quad (16.3)$$

16.1.2 Representing uncertainty on rotations

The Gaussian distribution is a popular choice to represent uncertainty on positions. Unfortunately, the Gaussian is defined on a vector space and we know that rotations (and poses) are not vector spaces, but

Lie groups. Moreover, there are multiple potential choices of distributions that “behave” similarly to the Gaussian distribution. We review two of them in the following.

16.1.2.1 Wrapped Gaussian distribution

A fairly natural way to define the equivalent of a Gaussian distribution on a rotation is:

$$\mathbf{R} = \bar{\mathbf{R}} \exp(\epsilon_r^\wedge) \quad \epsilon_r \sim \mathcal{N}(\mathbf{0}_d, \Sigma_r) \quad (16.4)$$

The previous relation implies a distribution over the rotation \mathbf{R} :

$$\mathbf{R} \sim \mathcal{W}(\bar{\mathbf{R}}, \Sigma_r) \doteq \frac{1}{\mathbf{J}_r(\log(\bar{\mathbf{R}}^\top \mathbf{R})^\vee)} \frac{1}{\sqrt{(2\pi)^d \det(\Sigma_r)}} \exp\left(-\frac{1}{2}(\log(\bar{\mathbf{R}}^\top \mathbf{R})^\vee)^\top \Sigma_r^{-1} (\log(\bar{\mathbf{R}}^\top \mathbf{R})^\vee)\right) \quad (16.5)$$

where $\bar{\mathbf{R}}$ is the *mean* rotation and Σ_r is the covariance of the distribution.

When Σ_r is “small”, $\epsilon_r = \log(\bar{\mathbf{R}}^\top \mathbf{R})^\vee$ is also small (with high probability) and $\mathbf{J}_r(\epsilon_r) \approx \mathbf{I}_d$ and the distribution (16.5) is approximated as:

$$\mathcal{W}(\bar{\mathbf{R}}, \Sigma_r) \approx \frac{1}{\sqrt{(2\pi)^d \det(\Sigma_r)}} \exp\left(-\frac{1}{2}(\log(\bar{\mathbf{R}}^\top \mathbf{R})^\vee)^\top \Sigma_r^{-1} (\log(\bar{\mathbf{R}}^\top \mathbf{R})^\vee)\right) \quad (16.6)$$

16.1.2.2 Isotropic Langevin distribution

The Isotropic Langevin distribution is defined as [3]:

$$\mathcal{L}(\bar{\mathbf{R}}, \kappa) = \frac{1}{c_d(\kappa)} \exp(\kappa \operatorname{tr}(\bar{\mathbf{R}}^\top \mathbf{R})) \quad (16.7)$$

where $\bar{\mathbf{R}}$ and κ are called the mode and the concentration parameter of the distribution, and $c_d(\kappa)$ is a normalization constant.

In the 2D case, the Langevin distribution is known as the Von Mises distribution:

$$\mathcal{V}(\bar{\theta}, \kappa) = \frac{1}{c_2(\kappa)} \exp(\kappa \cos(\theta - \bar{\theta})) \quad (16.8)$$

where the equivalence can be seen from $\operatorname{tr}(\bar{\mathbf{R}}^\top \mathbf{R}) = \operatorname{tr}(\mathbf{R}(\bar{\theta})^\top \mathbf{R}(\theta)) = 2 \cos(\theta - \bar{\theta})$, where the factor “2” is included in the normalization constant $c_2(\kappa)$.

16.1.3 Representing uncertainty on poses

A fairly natural way to define the equivalent of a Gaussian distribution on a pose is:

$$\mathbf{T} = \bar{\mathbf{T}} \exp(\epsilon_T^\wedge) \quad \epsilon_T \sim \mathcal{N}(\mathbf{0}, \Sigma_T) \quad (16.9)$$

where $\bar{\mathbf{T}}$ is the *mean* pose and Σ_T is the covariance of the distribution.

In alternative, one can consider a pose at a pair of rotation and translation and assume (uncorrelated) distributions on each, e.g., use a Langevin distribution for the rotation and a Gaussian distribution on the translation. A longer discussion on how to represent uncertainty on poses is given by [2].

16.2 Estimating unknown quantities via optimization

Similarly to what we often suggested in the previous lectures, here we discuss how to use optimization to mitigate the presence of noise in the measurements, when estimating some unknown quantity. In particular, we consider 2 examples: point triangulation, and bundle adjustment.

16.2.1 Example 1: Point triangulation

Consider the case where we have two cameras at known poses $(\mathbf{R}_w^{c_1}, \mathbf{t}_w^{c_1})$ and $(\mathbf{R}_w^{c_2}, \mathbf{t}_w^{c_2})$ and known calibration matrices \mathbf{K}_1 and \mathbf{K}_2 . Assume that we measure the pixel projection $\tilde{\mathbf{x}}_1$ and $\tilde{\mathbf{x}}_2$ of an unknown 3D point \mathbf{p}^w in both cameras. The triangulation problem (or “structure reconstruction”) is: estimate the 3D position of the point given the pixel measurements.

From basic perspective projection, we know that:

$$\lambda_1 \tilde{\mathbf{x}}_1 = \mathbf{K}_1 [\mathbf{R}_w^{c_1} \ \mathbf{t}_w^{c_1}] \tilde{\mathbf{p}}^w \quad (16.10)$$

$$\lambda_2 \tilde{\mathbf{x}}_2 = \mathbf{K}_2 [\mathbf{R}_w^{c_2} \ \mathbf{t}_w^{c_2}] \tilde{\mathbf{p}}^w \quad (16.11)$$

Using a more compact notation:

$$\lambda_1 \tilde{\mathbf{x}}_1 = \mathbf{\Pi}_1 \tilde{\mathbf{p}}^w \quad (16.12)$$

$$\lambda_2 \tilde{\mathbf{x}}_2 = \mathbf{\Pi}_2 \tilde{\mathbf{p}}^w \quad (16.13)$$

where $\mathbf{\Pi}_1 = \mathbf{K}_1 [\mathbf{R}_w^{c_1} \ \mathbf{t}_w^{c_1}]$ and $\mathbf{\Pi}_2 = \mathbf{K}_2 [\mathbf{R}_w^{c_2} \ \mathbf{t}_w^{c_2}]$.

To get rid of the scale λ_1 and λ_2 , let us multiply the first equation by $[\tilde{\mathbf{x}}_1]_{\times}$ and the second equation by $[\tilde{\mathbf{x}}_2]_{\times}$, from which we get:

$$[\tilde{\mathbf{x}}_1]_{\times} \mathbf{\Pi}_1 \tilde{\mathbf{p}}^w = \mathbf{0} \quad (16.14)$$

$$[\tilde{\mathbf{x}}_2]_{\times} \mathbf{\Pi}_2 \tilde{\mathbf{p}}^w = \mathbf{0} \quad (16.15)$$

or more succinctly:

$$\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} \tilde{\mathbf{p}}^w = \mathbf{0} \quad (16.16)$$

where $\mathbf{A}_1 = [\tilde{\mathbf{x}}_1]_{\times} \mathbf{\Pi}_1$ and $\mathbf{A}_2 = [\tilde{\mathbf{x}}_2]_{\times} \mathbf{\Pi}_2$.

In the presence of noise the equation above (6 equalities in 3 variables) will have no solution, hence we rather look for a least square solution:

$$\min_{\|\tilde{\mathbf{p}}^w\|=1} \left\| \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} \tilde{\mathbf{p}}^w \right\|^2 = \min_{\|\tilde{\mathbf{p}}^w\|=1} \|\mathbf{A} \tilde{\mathbf{p}}^w\|^2 \quad \text{with } \mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} \quad (16.17)$$

which can be solved in closed form (solution is the eigenvector corresponding to the smallest eigenvalue of $\mathbf{A}^T \mathbf{A}$). The objective in (16.17) is said to minimize the “algebraic” error.

An alternative objective. Recall that the pixel measurements can be written as:

$$\mathbf{x}_1 = \begin{bmatrix} u_1 \\ v_1 \end{bmatrix} = \begin{bmatrix} \frac{[\mathbf{\Pi}_1 \tilde{\mathbf{p}}^w]_1}{[\mathbf{\Pi}_1 \tilde{\mathbf{p}}^w]_3} \\ \frac{[\mathbf{\Pi}_1 \tilde{\mathbf{p}}^w]_2}{[\mathbf{\Pi}_1 \tilde{\mathbf{p}}^w]_3} \end{bmatrix} = \pi(\mathbf{R}_{c_1}^w, \mathbf{t}_{c_1}^w, \mathbf{p}^w) \quad (16.18)$$

$$\mathbf{x}_2 = \begin{bmatrix} u_2 \\ v_2 \end{bmatrix} = \begin{bmatrix} \frac{[\mathbf{\Pi}_2 \tilde{\mathbf{p}}^w]_1}{[\mathbf{\Pi}_2 \tilde{\mathbf{p}}^w]_3} \\ \frac{[\mathbf{\Pi}_2 \tilde{\mathbf{p}}^w]_2}{[\mathbf{\Pi}_2 \tilde{\mathbf{p}}^w]_3} \end{bmatrix} = \pi(\mathbf{R}_{c_2}^w, \mathbf{t}_{c_2}^w, \mathbf{p}^w) \quad (16.19)$$

where on the right-hand-side we introduced a more compact notation for the perspective projection function $\pi(\cdot)$. Therefore, an alternative objective can be obtained as follows:

$$\min_{\mathbf{p}^w} \|\mathbf{x}_1 - \pi(\mathbf{R}_{c_1}^w, \mathbf{t}_{c_1}^w, \mathbf{p}^w)\|^2 + \|\mathbf{x}_2 - \pi(\mathbf{R}_{c_2}^w, \mathbf{t}_{c_2}^w, \mathbf{p}^w)\|^2 \quad (16.20)$$

In general, the solution of (16.20) will be different from the solution of (16.17). The objective in (16.20) is said to minimize the “geometric” error.

Invariance. The formulation (16.20) is invariant to a rigid-body transformation of all the two cameras. Numerically, one can show that the formulation (16.17), instead, is not invariant to rigid-body transformations.

16.2.2 Example 2: Bundle Adjustment

In the previous lecture we have seen how to estimate the relative pose between two cameras given N pixel correspondences (let us assume for now that there are no outliers). Now let’s assume that we can track the same N features across 3 consecutive images, i.e., we can detect and match features in image 1 and 2, and then we can match features in image 2 against features extracted in image 3.

As done in the previous lecture we can compute the motion between camera 1 and 2 and between 2 and 3 by estimating the corresponding essential matrices:

$$\mathbf{E}_{12} = \arg \min_{\mathbf{E}_{12} \in \mathcal{S}_E} \sum_{k=1}^N |\tilde{\mathbf{y}}_{k,2}^\top \mathbf{E}_{12} \tilde{\mathbf{y}}_{k,1}|^2 \quad (16.21)$$

$$\mathbf{E}_{23} = \arg \min_{\mathbf{E}_{23} \in \mathcal{S}_E} \sum_{k=1}^N |\tilde{\mathbf{y}}_{k,3}^\top \mathbf{E}_{23} \tilde{\mathbf{y}}_{k,2}|^2 \quad (16.22)$$

from which one can compute the relative poses (up to scale).

An alternative objective. An alternative approach to estimate the poses between the three cameras is to formulate a single optimization problem:

$$\min_{\substack{(\mathbf{R}_{c_i}^w, \mathbf{t}_{c_i}^w), i=1,2,3 \\ \mathbf{p}_k^w, k=1, \dots, N}} \sum_{k=1}^N \sum_{i=1}^3 \|\mathbf{x}_{k,i} - \pi(\mathbf{R}_{c_i}^w, \mathbf{t}_{c_i}^w, \mathbf{p}_k^w)\|^2 \quad (16.23)$$

where we minimize the *reprojection error* of the N points in each camera. Note that in (16.23), besides optimizing for the camera poses, we also optimize for the unknown positions of the 3D points \mathbf{p}_k^w , $k = 1, \dots, N$.

The previous formulation, which easily extends to the case of K cameras is known as *bundle adjustment*. Since it does not admit a closed-form solution in general, in the next lecture we will discuss numerical optimization methods to obtain a (local) minimizer of (16.23).

Invariance. The formulation (16.23) is invariant to a rigid-body transformation of all the cameras.

16.3 Maximum Likelihood and Maximum a Posteriori estimation

The discussion in the previous section highlights a basic question: what is the “right” objective function to minimize in order to estimate a variable of interest? Each minimization provides a different solution so is

there a grounded way to pick an objective function?

This question finds a satisfactory answer in estimation theory, where there are two frameworks (*Maximum Likelihood* and *Maximum a Posteriori* estimation, discussed below) that discuss how the assumption we make on the measurement noise dictates a meaningful objective function to minimize.

16.3.1 Maximum Likelihood Estimation (MLE)

Assume we are given N measurements $\mathbf{z}_1, \dots, \mathbf{z}_N$ (e.g., pixel measurements) that are function of a variable we want to estimate \mathbf{x} (e.g., camera poses, points). Assume that we are also given the conditional distributions:

$$\mathbb{P}(\mathbf{z}_j|\mathbf{x}) \quad (16.24)$$

Then the *maximum likelihood* estimator (MLE) is defined as:

$$\mathbf{x}_{\text{MLE}} = \arg \max_{\mathbf{x}} \mathbb{P}(\mathbf{z}_1, \dots, \mathbf{z}_N|\mathbf{x}) \quad (16.25)$$

where $\mathbb{P}(\mathbf{z}_1, \dots, \mathbf{z}_N|\mathbf{x})$ is also called the *measurement likelihood*. Equivalently:

$$\mathbf{x}_{\text{MLE}} = \arg \min_{\mathbf{x}} -\log \mathbb{P}(\mathbf{z}_1, \dots, \mathbf{z}_N|\mathbf{x}) \quad (16.26)$$

Assuming conditional independence between the measurements, the MLE estimator becomes:

$$\mathbf{x}_{\text{MLE}} = \arg \max_{\mathbf{x}} \prod_{j=1}^N \mathbb{P}(\mathbf{z}_j|\mathbf{x}) \quad (16.27)$$

or equivalently:

$$\mathbf{x}_{\text{MLE}} = \arg \min_{\mathbf{x}} -\log \prod_{j=1}^N \mathbb{P}(\mathbf{z}_j|\mathbf{x}) = \arg \min_{\mathbf{x}} -\sum_{j=1}^N \log \mathbb{P}(\mathbf{z}_j|\mathbf{x}) \quad (16.28)$$

Example: Linear measurements with Gaussian noise. Assume we have N random measurements of an unknown variable \mathbf{x} we want to estimate:

$$\mathbf{z}_j = \mathbf{A}_j \mathbf{x} + \boldsymbol{\epsilon}_j \quad (16.29)$$

where \mathbf{A}_j are known matrices of suitable dimensions and $\boldsymbol{\epsilon}_j \in \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_j)$.

Now we note that if $\boldsymbol{\epsilon}_j \in \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_j)$, then $\mathbf{z}_j \in \mathcal{N}(\mathbf{A}_j \mathbf{x}, \boldsymbol{\Sigma}_j)$ or equivalently:

$$\mathbb{P}(\mathbf{z}_j|\mathbf{x}) = \frac{1}{\kappa_j} \exp -\frac{1}{2}(\mathbf{z}_j - \mathbf{A}_j \mathbf{x})^\top \boldsymbol{\Sigma}_j^{-1}(\mathbf{z}_j - \mathbf{A}_j \mathbf{x}) \quad (16.30)$$

where κ_j is the normalization constant of the Normal distribution. Then the MLE is:

$$\mathbf{x}_{\text{MLE}} = \arg \min_{\mathbf{x}} -\sum_{j=1}^N \log \mathbb{P}(\mathbf{z}_j|\mathbf{x}) = \arg \min_{\mathbf{x}} -\sum_{j=1}^N \log \frac{1}{\kappa_j} \exp -\frac{1}{2}(\mathbf{z}_j - \mathbf{A}_j \mathbf{x})^\top \boldsymbol{\Sigma}_j^{-1}(\mathbf{z}_j - \mathbf{A}_j \mathbf{x}) = \quad (16.31)$$

$$\arg \min_{\mathbf{x}} -\sum_{j=1}^N \log \frac{1}{\kappa_j} - \sum_{j=1}^N \log \exp -\frac{1}{2}(\mathbf{z}_j - \mathbf{A}_j \mathbf{x})^\top \boldsymbol{\Sigma}_j^{-1}(\mathbf{z}_j - \mathbf{A}_j \mathbf{x}) = \quad (16.32)$$

$$\arg \min_{\mathbf{x}} \sum_{j=1}^N \frac{1}{2}(\mathbf{z}_j - \mathbf{A}_j \mathbf{x})^\top \boldsymbol{\Sigma}_j^{-1}(\mathbf{z}_j - \mathbf{A}_j \mathbf{x}) \quad (16.33)$$

when $\Sigma_j = \mathbf{I}$ for all j , we obtain standard least squares:

$$\mathbf{x}_{\text{MLE}} = \arg \min \sum_{j=1}^N \frac{1}{2} \|\mathbf{z}_j - \mathbf{A}_j \mathbf{x}\|^2 \quad (16.34)$$

Example: Nonlinear measurements with additive Gaussian noise. Assume we have N random measurements of an unknown variable \mathbf{x} we want to estimate:

$$\mathbf{z}_j = f_j(\mathbf{x}) + \epsilon_j \quad (16.35)$$

Repeating the same derivation of the linear case, we get that the MLE is:

$$\mathbf{x}_{\text{MLE}} = \arg \min \sum_{j=1}^N \frac{1}{2} (\mathbf{z}_j - f_j(\mathbf{x}))^\top \Sigma_j^{-1} (\mathbf{z}_j - f_j(\mathbf{x})) \quad (16.36)$$

which is a *nonlinear least squares* problem.

Example: Translation averaging. Let us assume we are given N measurements of an unknown translation \mathbf{t} . Assume that each measurement is affected by zero mean Gaussian noise with identity covariance:

$$\bar{\mathbf{t}}_i = \mathbf{t} + \epsilon_i, \quad i = 1, \dots, N \quad \text{with } \epsilon_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (16.37)$$

Following the same derivation of the linear case with Gaussian noise in the previous example, we obtain the MLE estimator:

$$\mathbf{t}_{\text{MLE}} = \arg \min_{\mathbf{t}} \sum_{i=1}^N \|\mathbf{t} - \bar{\mathbf{t}}_i\|^2 \quad (16.38)$$

whose solution can be computed in closed form as:

$$\mathbf{t}_{\text{MLE}} = \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{t}}_i \quad (16.39)$$

Example: Rotation averaging. Let us assume we are given N measurements of an unknown rotation \mathbf{R} . Assume that each measurement $\bar{\mathbf{R}}_i$ is modeled as:

$$\bar{\mathbf{R}}_i = \mathbf{R} \mathbf{R}_\epsilon, \quad i = 1, \dots, N \quad \text{with } \mathbf{R}_\epsilon \sim \mathcal{L}(\mathbf{I}, \kappa) \quad (16.40)$$

where $\mathcal{L}(\mathbf{I}, \kappa)$ is the Langevin distribution (introduced in Section 16.1.2.2), with the identity mode and concentration parameter κ . Let us assume $\kappa = 1$.

Applying Maximum Likelihood Estimation, we obtain the following estimator (you will prove this as part of Lab 7):

$$\mathbf{R}_{\text{MLE}} = \arg \min_{\mathbf{R} \in \text{SO}(d)} \sum_{i=1}^N \|\mathbf{R} - \bar{\mathbf{R}}_i\|_F^2 \quad (16.41)$$

One would be tempted to compute \mathbf{R}_{MLE} as the average of $\bar{\mathbf{R}}_i$ as we did in the previous examples, but unfortunately, that would not be a valid rotation, and hence violate the constraint $\mathbf{R} \in \text{SO}(d)$.¹ Therefore, we need to compute the correct “average” rotation by solving (16.41) (including the constraint $\mathbf{R} \in \text{SO}(d)$).

¹Such approach would be incorrect even in the case of 2D rotations, where each rotation $\bar{\mathbf{R}}_i$ can be identified by a single angle $\bar{\theta}_i$. In this case, one would suggest $\theta_{\text{MLE}} = \frac{1}{N} \sum_{i=1}^N \bar{\theta}_i$, but it is easy to see that this does not lead to a meaningful average (e.g., pick $\bar{\theta}_1 = \frac{3}{4}\pi$ and $\bar{\theta}_2 = -\frac{3}{4}\pi$: the arithmetic average is zero, while one would expect an average rotation of π).

It turns out this optimization problem admits a closed-form solution:

$$\arg \min_{\mathbf{R} \in \text{SO}(d)} \sum_{i=1}^N \|\mathbf{R} - \bar{\mathbf{R}}_i\|_F^2 = \arg \min_{\mathbf{R} \in \text{SO}(d)} \sum_{i=1}^N \|\mathbf{R}\|_F^2 + \|\bar{\mathbf{R}}_i\|_F^2 - 2\text{tr}(\bar{\mathbf{R}}_i^\top \mathbf{R}) = \arg \min_{\mathbf{R} \in \text{SO}(d)} \sum_{i=1}^N 2d - 2\text{tr}(\bar{\mathbf{R}}_i^\top \mathbf{R}) \quad (16.42)$$

Dropping constants (irrelevant for optimization) and calling $\mathbf{N} \doteq \sum_{i=1}^N \bar{\mathbf{R}}_i$:

$$\arg \min_{\mathbf{R} \in \text{SO}(d)} \sum_{i=1}^N 2d - 2\text{tr}(\bar{\mathbf{R}}_i^\top \mathbf{R}) = \arg \min_{\mathbf{R} \in \text{SO}(d)} \sum_{i=1}^N -\text{tr}(\bar{\mathbf{R}}_i^\top \mathbf{R}) = \arg \min_{\mathbf{R} \in \text{SO}(d)} -\text{tr}\left(\left(\sum_{i=1}^N \bar{\mathbf{R}}_i\right)^\top \mathbf{R}\right) = \arg \max_{\mathbf{R} \in \text{SO}(d)} \text{tr}(\mathbf{N}^\top \mathbf{R})$$

Now, once again note that:

$$\text{tr}(\mathbf{N}^\top \mathbf{R}) = \frac{1}{2} \left(\|\mathbf{N}\|_F^2 + \|\mathbf{R}\|_F^2 - \|\mathbf{N} - \mathbf{R}\|_F^2 \right) = \frac{1}{2} \left(\|\mathbf{N}\|_F^2 + d - \|\mathbf{N} - \mathbf{R}\|_F^2 \right). \quad (16.43)$$

Since \mathbf{N} is constant we have:

$$\arg \max_{\mathbf{R} \in \text{SO}(d)} \text{tr}(\mathbf{N}^\top \mathbf{R}) = \arg \min_{\mathbf{R} \in \text{SO}(d)} \|\mathbf{N} - \mathbf{R}\|_F \quad (16.44)$$

which is simply a projection of \mathbf{N} onto $\text{SO}(3)$. Such a projection can be computed as follows [4] (we also saw this in Lecture 15):

$$\arg \min_{\mathbf{R} \in \text{SO}(d)} \|\mathbf{N} - \mathbf{R}\|_F = \begin{cases} \mathbf{U}\mathbf{V}^\top & \text{if } \det(\mathbf{U}\mathbf{V}^\top) \geq 0 \\ \mathbf{U} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{bmatrix} \mathbf{V}^\top & \text{otherwise} \end{cases} \quad (16.45)$$

where $\mathbf{N} = \mathbf{U}\mathbf{D}\mathbf{V}^\top$ is a singular value decomposition of the matrix \mathbf{N} .

Therefore, the optimal solution to the rotation ‘‘averaging’’ problem (16.41) is to compute the Euclidean average $\mathbf{N} \doteq \sum_{i=1}^N \bar{\mathbf{R}}_i$ and then project this matrix to $\text{SO}(3)$.

16.3.2 Maximum a Posteriori Estimation (MAP)

Assume we are given N measurements $\mathbf{z}_1, \dots, \mathbf{z}_N$ (e.g., pixel measurements) that are function of a variable we want to estimate \mathbf{x} (e.g., camera poses, points). *Maximum a Posteriori Estimation* (MAP) is a generalization of MLE. The MAP estimator is:

$$\mathbf{x}_{\text{MAP}} = \arg \max_{\mathbf{x}} \mathbb{P}(\mathbf{x} | \mathbf{z}_1, \dots, \mathbf{z}_N) \quad (16.46)$$

Using Bayes rule:

$$\mathbf{x}_{\text{MAP}} = \arg \max_{\mathbf{x}} \mathbb{P}(\mathbf{x} | \mathbf{z}_1, \dots, \mathbf{z}_N) = \quad (16.47)$$

$$\arg \max_{\mathbf{x}} \frac{\mathbb{P}(\mathbf{z}_1, \dots, \mathbf{z}_N | \mathbf{x}) \mathbb{P}(\mathbf{x})}{\mathbb{P}(\mathbf{z}_1, \dots, \mathbf{z}_N)} = \quad (16.48)$$

$$\arg \max_{\mathbf{x}} \mathbb{P}(\mathbf{z}_1, \dots, \mathbf{z}_N | \mathbf{x}) \mathbb{P}(\mathbf{x}) \quad (16.49)$$

where again $\mathbb{P}(\mathbf{z}_1, \dots, \mathbf{z}_N | \mathbf{x})$ is the *likelihood* of the measurements given \mathbf{x} , and $\mathbb{P}(\mathbf{x})$ is a *prior* probability over \mathbf{x} . It is easy to see that MAP reduces to MLE when the prior is uniform ($\mathbb{P}(\mathbf{x})$ is constant).

Assuming independence between the measurements, and repeating the same derivation of the previous section:

$$\mathbf{x}_{\text{MAP}} = \arg \min_{\mathbf{x}} - \sum_{j=1}^N \log \mathbb{P}(\mathbf{z}_j | \mathbf{x}) - \log \mathbb{P}(\mathbf{x}) \quad (16.50)$$

MLE Example: triangulation. Assume that a 3D point is observed by two known cameras and that pixel measurements are affected by Gaussian noise:

$$\mathbf{x}_1 = \pi(\mathbf{R}_{c_1}^w, \mathbf{t}_{c_1}^w, \mathbf{p}^w) + \boldsymbol{\epsilon}_1 \quad (16.51)$$

$$\mathbf{x}_2 = \pi(\mathbf{R}_{c_2}^w, \mathbf{t}_{c_2}^w, \mathbf{p}^w) + \boldsymbol{\epsilon}_2 \quad (16.52)$$

with $\boldsymbol{\epsilon}_1, \boldsymbol{\epsilon}_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$. Then the MLE is:

$$(\mathbf{p}^w)^* = \arg \min_{\mathbf{p}^w} \|\mathbf{x}_1 - \pi(\mathbf{R}_{c_1}^w, \mathbf{t}_{c_1}^w, \mathbf{p}^w)\|^2 + \|\mathbf{x}_2 - \pi(\mathbf{R}_{c_2}^w, \mathbf{t}_{c_2}^w, \mathbf{p}^w)\|^2 \quad (16.53)$$

MAP Example: bundle adjustment. Assume that N 3D points are observed by K cameras and that pixel measurements are affected by Gaussian noise. Moreover, assume we have a prior on the first camera being at the origin of the reference frame.

$$\min_{(\mathbf{R}_{c_i}^w, \mathbf{t}_{c_i}^w), i=1, \dots, K} \sum_{k=1}^N \sum_{i=1}^K \|u_{k,i} - \pi(\mathbf{R}_{c_i}^w, \mathbf{t}_{c_i}^w, \mathbf{p}_k)\|^2 + \rho \|\mathbf{t}_{c_1}^w\|^2 + \omega \|\mathbf{R}_{c_1}^w - \mathbf{I}_3\|_F^2 \quad (16.54)$$

where we assumed the prior:

$$\mathbb{P}(\mathbf{t}_{c_1}^w) = \mathcal{N}(\mathbf{0}, \frac{1}{\rho} \mathbf{I}_3) \quad (16.55)$$

$$\mathbb{P}(\mathbf{R}_{c_1}^w) = \text{Langevin}(\mathbf{I}_3, \omega) \quad (16.56)$$

References

- [1] T. Barfoot. *State Estimation for Robotics*. Cambridge University Press, 2017.
- [2] T. D. Barfoot and P. T. Furgale. Associating uncertainty with three-dimensional poses for use in estimation problems. *IEEE Trans. Robotics*, 30(3):679–693, 2014.
- [3] A. Chiuso, G. Picci, and S. Soatto. Wide-sense estimation on the special orthogonal group. *Commun. Inf. Syst.*, 8:185–200, 2008.
- [4] R. Hartley, J. Trunpf, Y. Dai, and H. Li. Rotation averaging. *IJCV*, 103(3):267–305, 2013.