

Lecture 15: RANSAC and 3D-3D Correspondences

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 covers:

- RANSAC (which we will use to estimate the relative pose from pixel correspondences, possibly including wrong correspondences),
- estimation of the relative pose from 3D-3D correspondences (e.g., the ones given by an RGB-D camera).

For both topics, we follow the presentation of the original papers (ref. [2] for RANSAC, and ref. [1] for pose estimation from 3D-3D correspondences).

15.1 Recap on Essential matrix Estimation

In the last lecture we saw how to estimate the essential matrix from 8 point correspondences. In particular, we observed that, in the absence of noise, (a vectorized version of) the essential matrix can be obtained by solving the linear system:

$$\mathbf{A}\mathbf{e} = 0 \quad (15.1)$$

where \mathbf{e} is assumed to be nonzero (e.g., $\|\mathbf{e}\| = 1$) and \mathbf{A} is a suitable 8×9 matrix computed from the pixel correspondences. By solving the linear system (15.1) and re-arranging the entries of \mathbf{e} into a 3×3 matrix, we obtain the desired essential matrix \mathbf{E} . Note that since $\mathbf{A}\mathbf{e} = -\mathbf{A}\mathbf{e} = 0$ both \mathbf{E} and $-\mathbf{E}$ are valid solutions to the linear system (15.1), so we need to consider both as potential essential matrices (we discussed how to resolve this ambiguity in the previous lecture).

Noisy pixel measurements. Since the pixels measurements are typically affected by noise, the solution of the linear system may not be an essential matrix. Therefore, it is common to project the solution onto the essential space using eq. (13.15) in Lecture 13.

More than 8 points. If we have more than $N > 8$ point correspondences and in absence of noise, we can still expect to compute the right essential matrix via (15.1) (in general, $\mathbf{A} \in \mathbb{R}^{N \times 9}$). However, in presence of noise, the linear equations will not be satisfied exactly and we might look for a solution that is as close as possible to satisfying the linear equalities:

$$\min_{\mathbf{E} \in \mathcal{S}_E} \|\mathbf{A} \text{vec}(\mathbf{E})\|^2 \quad (15.2)$$

where we remarked that the minimization has to be restricted to valid essential matrices and $\text{vec}(\mathbf{E})$ denotes a vector of 9 elements stacking the entries of \mathbf{E} . We can follow the same logic of the 8-point algorithm and approximate the solution of (15.2) by the following steps:

- solve the relaxed problem: $\arg \min_{\|\mathbf{e}\|=1} \|\mathbf{A}\mathbf{e}\|^2$ (whose solution can be computed in closed form: it is simply the eigenvector corresponding to the smallest eigenvalue of $\mathbf{A}^\top \mathbf{A}$, see https://en.wikipedia.org/wiki/Rayleigh_quotient.)

- project the result to the Essential space using eq. (13.15) in Lecture 13.

What if we have wrong correspondences? If some of the point correspondences are wrong, the minimization (15.2) will not produce a meaningful result. Therefore, we introduce RANSAC, a method that will allow computing the essential matrix even when a subset of the correspondences are wrong.

15.2 RANSAC

From the previous section:

- we know how to compute the essential matrix from 8 “good” point correspondences.
- issue: we are given $N \gg 8$ point correspondences and some of them may be wrong.

RANSAC (**R**andom **S**ample **C**onsensus) is an approach to perform estimation (in our case: estimate the essential matrix) in presence of outliers (in our case: wrong pixel correspondences).

The basic idea behind RANSAC is straightforward: say that in general we can compute a model from n points (in the case of the 8-point method, clearly $n = 8$). Then the idea is to sample a subset of n elements from the N available correspondences till we get a “good” set. What is a “good” set? A set for which the corresponding model (i.e., the essential matrix) explains as many remaining correspondences as possible.

Example: Essential matrix estimation RANSAC proceeds as follows, when applied to the estimation of the Essential matrix. Given (i) N point correspondences, possibly including wrong matches, (ii) a minimal solver that can estimate the Essential matrix from n points (e.g., $n = 8$), RANSAC computes an estimate of \mathbf{E} as follows:

1. Randomly select a subset \mathcal{S} of n point correspondences out of the N available correspondences.
2. Estimate the Essential matrix \mathbf{E} from the correspondences in \mathcal{S}
3. Compute the set \mathcal{S}^* of correspondences (i, j) that are such that:

$$\mathbf{y}_j^\top \mathbf{E} \mathbf{y}_i \approx 0 \quad (15.3)$$

or more formally $\|\mathbf{y}_j^\top \mathbf{E} \mathbf{y}_i\| < \epsilon$ for some small threshold ϵ . The set \mathcal{S}^* is called the *consensus set* of \mathcal{S} .

4. Case (a): if the cardinality of \mathcal{S}^* (number of elements of \mathcal{S}^* , denoted as $|\mathcal{S}^*|$) is larger than a given threshold T , accept \mathcal{S}^* as the set of inliers and use it to recompute a better estimate of the essential matrix \mathbf{E} via (15.2).

Case (b): if $|\mathcal{S}^*| < T$ repeat from Step 1. Stop after a maximum number of iterations k_{\max} .

General algorithm We can generalize the example above as follows. Given (i) N data points, possibly corrupted by outliers, and (ii) a minimal solver that can estimate the unknown model \mathcal{P} from n data points, RANSAC proceeds as follows:

1. Randomly select a subset \mathcal{S} of n data points out of the N available points.
2. Estimate the model \mathcal{P} from the data points in \mathcal{S}

3. Compute the set \mathcal{S}^* of correspondences (i, j) that are in agreement with the model \mathcal{P} up to some given tolerance ϵ .
4. Case (a): if $|\mathcal{S}^*| \geq T$, accept \mathcal{S}^* as the set of inliers and use it to recompute a better estimate of the model \mathcal{P} .

Case (b): if $|\mathcal{S}^*| < T$ repeat from Step 1. Stop after a maximum number of iterations k_{\max} .

15.2.1 Parameter tuning

15.2.1.1 Error Tolerance ϵ

How to set the tolerance ϵ that is used to decide whether a given correspondence belongs to the consensus set? Let us consider our Essential matrix estimation example, where the condition to satisfy is:

$$\|\mathbf{y}_j^T \mathbf{E} \mathbf{y}_i\| < \epsilon \quad (15.4)$$

In this case one may select ϵ as a function of the pixel noise, but that would not account for the fact that the matrix \mathbf{E} estimated from the 8 points in the set \mathcal{S} is also noisy. In some cases one may do some error propagation, but it is more common to set this tolerance empirically by trial and error.

15.2.1.2 Maximum number of iterations k_{\max}

Assume we know the probability w that a given correspondence is an inlier¹. Moreover, call k the number of trials required to select an outlier-free set (note: k is a random variable since we sample at random). Then the expected value of k is:

$$\mathbb{E}[k] = \sum_{i=1}^{\infty} i \mathbb{P}(k = i) = b + 2ab + 3a^2b + \dots = b[1 + 2a + 3a^2 + \dots] \quad (15.5)$$

where $b = w^n$ and $a = (1 - b)$. Intuitively: $\mathbb{P}(k = 1)$ is the probability of extracting all n “good points” at the first iteration, hence $\mathbb{P}(k = 1) = w^n$; $\mathbb{P}(k = 2)$ is the probability of **not** extracting all good points at the first iteration (probability: $1 - w^n$) and extracting all good points at the second (probability: w^n), and so on. The infinite series in the expression above is equal to:

$$1 + 2a + 3a^2 + \dots = \frac{1}{(1 - a)^2} \quad (15.6)$$

Recalling that $a = (1 - b)$:

$$\mathbb{E}[k] = b \frac{1}{(1 - a)^2} = b \frac{1}{b^2} = \frac{1}{b} = w^{-n} \quad (15.7)$$

Typically, one would set the maximum number of iterations k_{\max} equal to 2 – 3 times the expected value above. Ref. [2] provides a nice justification for that in terms of the variance of k .

15.2.1.3 Admissible size of an acceptable consensus set T

Calling q the probability that a given data point is within the tolerance of a wrong model, then q^{T-n} is the probability that $T - n$ points agree with the incorrect model. Assuming $q < 0.5$ (i.e., it’s somehow unlikely

¹More precisely, this is the probability that a correspondence falls in the consensus set built from a set of inliers.

that a point agrees with an incorrect model by chance), then for $T = n + 5$, it holds $q^{T-n} < 0.5^5 < 0.03$, i.e., it is unlikely to have a consensus set larger than $T = n + 5$ for an incorrect model.

Again, since these theoretical bounds are based on (often unrealistic) simplifying assumptions, it is more common to set T to a larger number, e.g., $T > N/2$.

15.3 Relative Pose Estimation from 3D-3D Correspondences

Assume that we are given N (calibrated) pixel correspondences $(\tilde{\mathbf{y}}_{1,k}, \tilde{\mathbf{y}}_{2,k})$ for $k = 1, \dots, N$ picturing corresponding points in 2 cameras. Moreover, assume that we have an RGB-D camera which is also able to measure the depth at each pixel. How can we estimate the relative pose between the cameras?

Since we have the depth $\lambda_{1,k}$ for each pixel $\tilde{\mathbf{y}}_{1,k}$, we can reproject the corresponding 3D point as $\mathbf{p}_{1,k} = \lambda_{1,k} \tilde{\mathbf{y}}_{1,k}$ (and similarly for the points in the second camera). Therefore the problem can be understood as the estimation of the relative pose between 2 sets of 3D points observed by 2 cameras, given correspondences $(\mathbf{p}_{1,k}, \mathbf{p}_{2,k})$, for $k = 1, \dots, N$.

In the absence of noise, each point k must satisfy:

$$\mathbf{p}_{1,k} = \mathbf{R}_{c_2}^{c_1} \mathbf{p}_{2,k} + \mathbf{t}_{c_2}^{c_1} \quad k = 1, \dots, N \quad (15.8)$$

i.e., the points seen by the two cameras only differ by a rigid-body transformation $(\mathbf{R}_{c_2}^{c_1}, \mathbf{t}_{c_2}^{c_1})$ corresponding to the relative pose between the cameras. For simplicity of notation, in the following we simply use \mathbf{R} (instead of $\mathbf{R}_{c_2}^{c_1}$) and \mathbf{t} (instead of $\mathbf{t}_{c_2}^{c_1}$) since this is the only pose of interest in the rest of these notes. Therefore, in absence of noise we can compute \mathbf{R} and \mathbf{t} by solving the set of linear equations (15.8).

In presence of noise, there is no pose (\mathbf{R}, \mathbf{t}) that exactly satisfies all the equations in (15.8) and we can only look for a pose that makes $\mathbf{p}_{1,k} - (\mathbf{R}\mathbf{p}_{2,k} + \mathbf{t})$ close to zero. This suggests estimating the pose (\mathbf{R}, \mathbf{t}) by solving the following minimization problem:

$$\min_{\substack{\mathbf{R} \in \text{SO}(3) \\ \mathbf{t} \in \mathbb{R}^3}} \sum_{k=1}^N \|\mathbf{p}_{1,k} - (\mathbf{R}\mathbf{p}_{2,k} + \mathbf{t})\|^2 = \sum_{k=1}^N \|\mathbf{p}_{1,k} - \mathbf{R}\mathbf{p}_{2,k} - \mathbf{t}\|^2 \quad (15.9)$$

For any choice of \mathbf{R} , the minimum over \mathbf{t} is attained at

$$\mathbf{t}^* = \frac{1}{N} \sum_{k=1}^N (\mathbf{p}_{1,k} - \mathbf{R}\mathbf{p}_{2,k}) \quad (15.10)$$

[hint: from “the point of view” of \mathbf{t} , problem (15.9) is the same as $\min_{\mathbf{t}} \sum_{k=1}^N \|\mathbf{t}_k - \mathbf{t}\|^2$, whose solution is the average: $\mathbf{t}^* = \frac{1}{N} \sum_{k=1}^N \mathbf{t}_k$].

Substituting \mathbf{t}^* back into (15.9) we get:

$$\min_{\mathbf{R} \in \text{SO}(3)} \sum_{k=1}^N \left\| \mathbf{p}_{1,k} - \mathbf{R}\mathbf{p}_{2,k} - \frac{1}{N} \sum_{k=1}^N (\mathbf{p}_{1,k} - \mathbf{R}\mathbf{p}_{2,k}) \right\|^2 = \quad (15.11)$$

$$\min_{\mathbf{R} \in \text{SO}(3)} \sum_{k=1}^N \left\| \left(\mathbf{p}_{1,k} - \frac{1}{N} \sum_{k=1}^N \mathbf{p}_{1,k} \right) - \mathbf{R} \left(\mathbf{p}_{2,k} - \frac{1}{N} \sum_{k=1}^N \mathbf{p}_{2,k} \right) \right\|^2 \quad (15.12)$$

Redefining $\mathbf{p}'_{1,k} = \mathbf{p}_{1,k} - \frac{1}{N} \sum_{k=1}^N \mathbf{p}_{1,k}$ and $\mathbf{p}'_{2,k} = \mathbf{p}_{2,k} - \frac{1}{N} \sum_{k=1}^N \mathbf{p}_{2,k}$:

$$\min_{\mathbf{R} \in \text{SO}(3)} \sum_{k=1}^N \|\mathbf{p}'_{1,k} - \mathbf{R}\mathbf{p}'_{2,k}\|^2 = \min_{\mathbf{R} \in \text{SO}(3)} \sum_{k=1}^N \|\mathbf{p}'_{1,k}\|^2 + \|\mathbf{R}\mathbf{p}'_{2,k}\|^2 - 2(\mathbf{p}'_{1,k})^\top \mathbf{R}\mathbf{p}'_{2,k} = \quad (15.13)$$

$$\min_{\mathbf{R} \in \text{SO}(3)} \sum_{k=1}^N -\text{tr}(\mathbf{p}'_{2,k}(\mathbf{p}'_{1,k})^\top \mathbf{R}) = \min_{\mathbf{R} \in \text{SO}(3)} -\text{tr}(\mathbf{M}^\top \mathbf{R}) = \max_{\mathbf{R} \in \text{SO}(3)} \text{tr}(\mathbf{M}^\top \mathbf{R}) \quad (15.14)$$

where $\mathbf{M} = \sum_{k=1}^N (\mathbf{p}'_{2,k}(\mathbf{p}'_{1,k})^\top)^\top$. Now we note that:

$$\arg \min_{\mathbf{R} \in \text{SO}(3)} \|\mathbf{R} - \mathbf{M}\|_F^2 = \arg \min_{\mathbf{R} \in \text{SO}(3)} \text{tr}((\mathbf{R} - \mathbf{M})(\mathbf{R} - \mathbf{M})^\top) \quad (15.15)$$

$$\begin{aligned} & \text{(recall } \mathbf{R}^\top \mathbf{R} = \mathbf{I} \text{ and that constants are irrelevant for the optimization)} \\ & = \arg \min_{\mathbf{R} \in \text{SO}(3)} \text{tr}(\mathbf{I} - 2\mathbf{M}^\top \mathbf{R} + \mathbf{M}\mathbf{M}^\top) = \arg \max_{\mathbf{R} \in \text{SO}(3)} \text{tr}(\mathbf{M}^\top \mathbf{R}) \end{aligned} \quad (15.16)$$

therefore maximizing $\text{tr}(\mathbf{M}^\top \mathbf{R})$ in eq. (15.14) is the same as solving $\arg \min_{\mathbf{R} \in \text{SO}(3)} \|\mathbf{R} - \mathbf{M}\|_F^2$, which is the definition of the *projection* of a matrix \mathbf{M} onto the set $\text{SO}(3)$. This projection has a closed-form solution (if you are interested, you can find the closed-form solution in the Appendix below).

The solution to relative pose estimation from 3D-3D correspondences outlined in this section is often called “Arun’s method” and was presented in [1]. One can also repeat a similar derivation by representing rotations via unit quaternions, and get the so called “Horn’s method”, which is described in the beautiful paper [3].

15.3.1 Appendix: Projection onto $\text{O}(n)$ and $\text{SO}(n)$

Given an arbitrary square matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, projecting \mathbf{M} onto $\text{O}(n)$ (resp. $\text{SO}(n)$) consists in finding the matrix $\mathbf{Q} \in \text{O}(n)$ (resp. $\mathbf{R} \in \text{SO}(n)$) that is *nearest* to \mathbf{M} :

$$\Pi_{\text{O}(n)}(\mathbf{M}) = \arg \min_{\mathbf{Q} \in \text{O}(n)} \|\mathbf{Q} - \mathbf{M}\|_F^2; \quad \Pi_{\text{SO}(n)}(\mathbf{M}) = \arg \min_{\mathbf{R} \in \text{SO}(n)} \|\mathbf{R} - \mathbf{M}\|_F^2. \quad (15.17)$$

This is a notable example of tractable nonconvex problems. In fact, the two projection problems in eq. (15.17) admit closed-form solutions, using *singular value decomposition* (SVD) [1, 4]. Formally, let $\mathbf{M} = \mathbf{U}\mathbf{S}\mathbf{V}^\top$ be the SVD of \mathbf{M} , where $\mathbf{U}, \mathbf{V} \in \text{O}(n)$, and \mathbf{S} contains the singular values of \mathbf{M} in *descending* order,² i.e., $S_{11} \geq S_{22} \geq \dots \geq S_{nn} \geq 0$, then the projections in (15.17) admit the following closed-form solutions:

$$\Pi_{\text{O}(n)}(\mathbf{M}) = \mathbf{U}\mathbf{V}^\top, \quad (15.18)$$

$$\Pi_{\text{SO}(n)}(\mathbf{M}) = \mathbf{U} \text{diag}([1, 1, \dots, \det(\mathbf{U}) \det(\mathbf{V})]) \mathbf{V}^\top. \quad (15.19)$$

Proof. We only prove the solution for projection onto $\text{O}(n)$ (eq. (15.18)), while we leave the proof for projection onto $\text{SO}(n)$ as an exercise. To prove eq. (15.18) is the solution for problem (15.17), we develop the cost function of problem (15.17):

$$\arg \min_{\mathbf{Q} \in \text{O}(n)} \|\mathbf{Q} - \mathbf{M}\|_F^2 = \arg \min_{\mathbf{Q} \in \text{O}(n)} \text{tr}((\mathbf{Q} - \mathbf{M})(\mathbf{Q} - \mathbf{M})^\top) \quad (15.20)$$

$$= \arg \min_{\mathbf{Q} \in \text{O}(n)} \text{tr}(\mathbf{Q}\mathbf{Q}^\top) + \text{tr}(\mathbf{M}\mathbf{M}^\top) - 2\text{tr}(\mathbf{Q}^\top \mathbf{M}) = \arg \max_{\mathbf{Q} \in \text{O}(n)} \text{tr}(\mathbf{Q}^\top \mathbf{M}) \quad (15.21)$$

$$= \arg \max_{\mathbf{Q} \in \text{O}(n)} \text{tr}(\mathbf{Q}^\top \mathbf{U}\mathbf{S}\mathbf{V}^\top) = \arg \max_{\mathbf{Q} \in \text{O}(n)} \text{tr}(\mathbf{S}\mathbf{V}^\top \mathbf{Q}^\top \mathbf{U}), \quad (15.22)$$

²For example, Matlab `svd` returns the singular values in descending order.

and observe that $\mathbf{V}^T \mathbf{Q}^T \mathbf{U}$ is an orthogonal matrix because $\mathbf{V}, \mathbf{Q}, \mathbf{U} \in O(n)$. Therefore, as \mathbf{S} is a diagonal matrix with positive entries, the maximum of $\text{tr}(\mathbf{S} \mathbf{V}^T \mathbf{Q}^T \mathbf{U})$ is attained if and only if $\mathbf{V}^T \mathbf{Q}^T \mathbf{U} = \mathbf{I}_n$, i.e., $\mathbf{Q} = \mathbf{U} \mathbf{V}^T$. \square

References

- [1] K.S. Arun, T.S. Huang, and S.D. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Trans. Pattern Anal. Machine Intell.*, 9(5):698–700, sept. 1987.
- [2] R. Bolles and M. Fischler. A RANSAC-based approach to model fitting and its application to finding cylinders in range data. In *Intl. Joint Conf. on AI (IJCAI)*, pages 637–643, Vancouver, BC, Canada, 1981.
- [3] Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *J. Opt. Soc. Amer.*, 4(4):629–642, Apr 1987.
- [4] P.H. Schonemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31:1–10, 1966.