Outlier-tolerant parameter estimation

Bayesian Methods in physics, statistics, machine learning, and signal processing (SS 2003)

Friedrich Fraundorfer

fraunfri@icg.tu-graz.ac.at

Computer Graphics and Vision

TUG

Graz University of Technology

Outline

- Epipolar Geometry
- A Bayesian Formulation
- MAP estimation using RANSAC
- Example

Epipolar Geometry

A Bayesian Formulation

$$\hat{\Theta} = \arg\max_{\Theta} \Pr(\Theta | D, M, I) =$$

$$= \arg\max_{\Theta} \frac{\Pr(D | \Theta, M, I)\Pr(\Theta | M, I)}{\Pr(D | M, I)}$$

$$\Theta = \{\alpha, \beta, \gamma\}$$

- Θ .. set of parameters
- $\hat{\Theta}$ .. MAP estimate
- D .. data (image correspondences)
- M .. hypothesized model (motion model)
- I .. conditioning information

Parameter set Θ

$$\Theta = \{\alpha, \beta, \gamma\}$$

- α .. 9 parameters of fundamental matrix

$$\mathbf{F} = \begin{bmatrix} \alpha_1 & \alpha_2 & \alpha_3 \\ \alpha_4 & \alpha_5 & \alpha_6 \\ \alpha_7 & \alpha_8 & \alpha_9 \end{bmatrix} \qquad \mathbf{x}^{2T}\mathbf{F}\mathbf{x}^1 = 0$$

- β .. parameters of point correspondences

$$m_i = (x_i^1, y_i^1, x_i^2, y_i^2), i = 1..n$$

- γ .. set of indicator variables, γ_i, such that γ_i = 1 for inlier, γ_i = 0 for outlier $i = 1..n$

What to estimate (CV)?

- MAP estimate of α, β, γ .
- Recovers camera viewpoint and 3D structure.
- Parameters of the relation α
- Which matches are inliers γ
- 3D coordinates of matches β

Likelihood [1|3]

$$\Pr(D|\Theta, M, I) = \Pr(D|\alpha, \beta, \gamma, M, I) = \Pr(D|\beta, \gamma, M, I)$$

- Depends only on β the parameters of the matches and γ which determines whether or not the matches are outliers or inliers.
- We now assume that the likelihood is a mixture model of distributions for inliers and outliers.

Likelihood [2|3]

- We assume that we only know the first two moments of the inlier distribution and only the range of the outliers, then the maximally non informative distributions are Gaussian distribution for inliers and uniform distribution for outliers.

$$\Pr(D|\beta, M, I) = \left(\frac{1}{\sqrt{2\pi\sigma^2}}\right)^{nD}$$

$$\times \prod_{i=1..n} \exp\left(-\frac{\sum_{j=1,2} (\hat{x}_i^j - x_i^j)^2 + (\hat{y}_i^j - y_i^j)^2}{2\sigma^2}\right)$$

$$e_i^2 = \sum_{j=1,2} (\hat{x}_i^j - x_i^j)^2 + (\hat{y}_i^j - y_i^j)^2$$

Likelihood [3|3]

- Outlier distribution

$$\Pr(D|\gamma, M, I) = \frac{1}{v}$$

$$v = L \times L \times S \times S$$

$$L \times L \text{...Bildgroesse}$$

$$S \times S \text{...Searchwindow}$$

$$\Pr(D|\beta, \gamma, M, I) = \prod_{i=1..n} \psi(e_i)$$

$$= \prod_{i=1..n} \underbrace{(\gamma_i (\frac{1}{\sqrt{2\pi\sigma^2}})^D \exp(-\frac{e_i^2}{2\sigma^2})}_{\text{inlier (distr.)}} + \underbrace{(1-\gamma_i)\frac{1}{v}}_{\text{outlier (distr).}})$$

Prior [1|2]

$$\Pr(\Theta|M,I) = \Pr(\alpha,\beta,\gamma|M,I)$$

$$= \underbrace{\Pr(\beta,\gamma|\alpha,M,I)}_{\text{prior of matches}}\underbrace{\Pr(\alpha|M,I)}_{\text{prior of relation}}$$

- We assume that β and γ are independent (reasonable for this CV problem)

$$= \Pr(\beta|\alpha,M,I)\Pr(\gamma|\alpha,M,I)\Pr(\alpha|M,I)$$

- β_i uniformly distributed (the parameters β_i associated with each match are independent)

$$\Pr(\beta|\alpha,M,I) = \frac{1}{c}$$

$c \approx L \times L \times S$

$L \times L$...Bildgroesse

S...Searchrange

Prior [2|2]

- We assume that the γ_i are independent and that the prior probability of a given point being an inlier is constant over all i

$$\{\Pr(\gamma_i) = 1\} = \gamma; \forall i$$

- γ is the prior expectation of seeing an inlier (expected percentage of inliers in the data)

- $\Pr(\alpha | M, I)$ will typically have little effect as $n \to \infty$ and won‘t be included in the posterior model

Posterior of the model

$$\Pr(\Theta \mid D, M, I) \propto$$

$$\Pr(D \mid \alpha, \beta, \gamma, M, I)\Pr(\alpha, \beta, \gamma \mid M, I) =$$

$$\Pr(D \mid \beta, \gamma, M, I)\Pr(\beta \mid \alpha, M, I)\Pr(\gamma \mid \alpha, M, I)\Pr(\alpha \mid M, I) \propto$$

$$\prod_{i=1..n}(\gamma_i(\frac{1}{\sqrt{2\pi\sigma^2}})^D \exp(-\frac{e_i^2}{2\sigma^2}) + (1-\gamma_i)\frac{1}{v})\Pr(\beta_i \mid \alpha, M, I)$$

$$\times \Pr(\gamma \mid \alpha, M, I)$$

Marginalization over $\beta = (\beta_1 \ldots \beta_n)$ [1|2]

$$\int_{\beta} \Pr(\Theta | D, M, I) \partial \beta \propto$$

$$\Big(\prod_{i=1..n} \int_{\beta_i} \big(\gamma_i \big(\frac{1}{\sqrt{2\pi\sigma^2}}\big)^D \exp\big(-\frac{e_i^2}{2\sigma^2}\big) + (1-\gamma_i)\frac{1}{v}\big) \Pr(\beta_i | \alpha, M, I) \partial \beta_i \Big)$$

$$\times \Pr(\gamma | \alpha, M, I) \Pr(\alpha | M, I)$$

Integrate the uniform distribution

$$\int_{\beta_i} (1-\gamma_i)\frac{1}{v}) \frac{1}{c} \partial \beta_i = (1-\gamma_i)\frac{1}{v}$$

$$\Pr(\beta_i | \alpha, M, I) = \frac{1}{c}$$

Marginalization over $\beta = (\beta_1 \ldots \beta_n)$ [2|2]

Integrate the Gaussian part

$$\int_{\beta_i} \gamma_i \left(\frac{1}{\sqrt{2\pi\sigma^2}} \right)^D \exp\left(-\frac{e_i^2}{2\sigma^2} \right) \partial\beta_i \approx \gamma_i \frac{\sqrt{2\pi\sigma^2}^{\,d-D}}{c} \exp\left(-\frac{e_\perp^2}{2\sigma^2} \right)$$

The MAP estimate

$$\int_{\beta_i} (\ldots) = \gamma_i \left(\frac{\sqrt{2\pi\sigma^2}^{\,d-D}}{c} \exp\left(-\frac{e_\perp^2}{2\sigma^2} \right) \right) + (1-\gamma_i)\frac{1}{v}$$

Determining indicator variables γ

Point correspondences with an re-projection error < T are considered to be inliers

$$\hat{\gamma}_i = \begin{cases} 1, & \dfrac{e^2}{\sigma^2} \le T \\ 0, & \dfrac{e^2}{\sigma^2} > T \end{cases} \qquad T = -2\log\left(\frac{1-\gamma}{\gamma}\frac{c}{v}\sqrt{2\pi\sigma^2}^{\,D-d}\right)$$

The MAP estimate marginalizing over β and γ

$$\int_{\beta_i}(\ldots) = \hat{\gamma}_i\left(\frac{\sqrt{2\pi\sigma^2}^{\,d-D}}{c}\exp\left(-\frac{e_\perp^2}{2\sigma^2}\right)\right) + (1-\hat{\gamma}_i)\frac{1}{v}$$

MAPSAC: An algorithm for the MAP Estimator

- We want an algorithm for efficiently finding modes in the posterior.
- Idea is to extend the RANSAC-Algorithm
- MAPSAC – Maximum a posteriori sample consensus

RANSAC [1|2]

- *Ran*dom *Sa*mple *C*onsensus

- Divide data points into inliers and outliers
- Use only inliers to calculate the model
- Randomly sample subsets of the data in the hope to get one subset with no outliers

RANSAC [2|2]

1. From the data set D select randomly a sample of d data points. Estimate a model from this subset.
2. Determine the set of data points d_i which are within a distance threshold T of the model. This is the consensus set of the sample and contains the inliers for the current model.
3. If the number of inliers (from d_i) is greater than some threshold T terminate. Re-estimate the model using only the inliers.
4. If the number of inliers is less than T go back to step 1.
5. After N trials the largest consensus set is selected and the model is re-estimated using only the inliers, if the run hasn‘t been terminated before.

Using RANSAC for MAP Estimation

RANSAC finds the miniumum of a cost function defined as

$$C = \sum_i \rho\left(\frac{e_i^2}{\sigma^2}\right) \qquad \rho\left(\frac{e^2}{\sigma^2}\right) = \begin{cases} 0, & \dfrac{e^2}{\sigma^2} < T \\ pos.const., & \dfrac{e^2}{\sigma^2} \geq T \end{cases}$$

Approximate the posterior by minimizing

$$C_2 = \sum_i \rho_2\left(\frac{e_i^2}{\sigma^2}\right) \qquad \rho_2\left(\frac{e^2}{\sigma^2}\right) = \begin{cases} \dfrac{e^2}{\sigma^2}, & \dfrac{e^2}{\sigma^2} < T \\ T, & \dfrac{e^2}{\sigma^2} \geq T \end{cases}$$

MAPSAC-Algorithm

Table 2. A brief summary of the stages of the Bayesian random sampling estimation algorithm MAPSAC.

1. Obtain data $\mathbf{m}_i$, $i = 1 \ldots n$ from some source, and determine the manifold to be fit.
2. Repeat (a)–(c) until a large number of samples have been taken or "jump out" occurs as described in Torr and Zisserman (1997).
   (a) Select a random sample of the minimum number of points (matches) $S_m = \{\mathbf{m}_i\}$ to estimate θ.
   (b) Estimate θ consistent with this minimal set S_m.
   (c) Calculate the posterior for θ: $\Pr(\theta \mid \mathcal{D}, \mathcal{M}, \mathcal{I}) \propto \prod_{i=1 \ldots n} (\psi(e_i)) \Pr(\beta \mid \alpha, \mathcal{M}, \mathcal{I}) \Pr(\gamma \mid \alpha, \mathcal{M}, \mathcal{I})$ $\Pr(\alpha \mid \mathcal{M}, \mathcal{I})$.. This can either be approximated by maximizing γ, or by marginalizing γ via EM
      (i) *Either*: Approximate posterior by minimizing $C_2 = \sum_i \rho_2(\frac{e^2}{2\sigma_i^2})$, with $\rho_2(\ )$ defined in (23), and T defined in (18). (*Lower computational cost*).
      (ii) *Or*: use EM to marginalize γ and maximize the posterior $\Pr(\theta \mid \mathcal{D}, \mathcal{M}, \mathcal{I})$ directly. (*Higher computational cost*).
3. Select the best solution $\hat{\theta}$ over all the samples i.e. that which maximized the posterior (or minimizes C_2).
4. Maximize the posterior (minimize C_2) using non-linear optimization as described in Torr and Zisserman (1997).

[Torr2002]

Results on Epipolar Geometry estimation

- Set of 20 point correspondences
- minimum number of 7 required for Fundamental Matrix estimation
- 14 true inliers
- 6 true outliers

Using inliers only

Using inliers&outliers

Robust estimation

References

- Torr, P.H.S.: Bayesian Model Estimation and Selection for Epipolar Geometry and Generic Manifold Fitting
- Fischler, M.A. and Bolles, R.C.: Random sample consensus: A paradigm for model fitting with application to image analysis and automated cartography