



**OPPA European Social Fund
Prague & EU: We invest in your future.**



► The Full Problem of Matching and Fundamental Matrix Estimation

Problem: Given two sets of image points $X = \{x_i\}_{i=1}^m$ and $Y = \{y_j\}_{j=1}^n$ and their descriptors D , find the most probable

1. inliers $S_X \subseteq X$, $S_Y \subseteq Y$
2. one-to-one perfect matching $M: S_X \rightarrow S_Y$
3. fundamental matrix \mathbf{F} such that $\text{rank } \mathbf{F} = 2$
4. such that for each $x_i \in S_X$ and $y_j = M(x_i)$ it is probable that
 - a. the image descriptor $D(x_i)$ is similar to $D(y_j)$, and
 - b. the total geometric error $\sum_{i,j} e_{ij}^2(\mathbf{F})$ is small
5. inlier-outlier and outlier-outlier matches are improbable

perfect matching: 1-factor of the bipartite graph

note a slight change in notation: e_{ij}



$$(M^*, \mathbf{F}^*) = \arg \max_{M, \mathbf{F}} p(M, \mathbf{F} \mid X, Y, D) \quad (17)$$

- probabilistic model: an efficient language for task formulation
- the (17) is a p.d.f. for all the involved variables (there is a constant number of variables!)
- binary matching table $M_{ij} \in \{0, 1\}$ of fixed size $m \times n$
 - each row/column contains at most one unity
 - zero rows/columns correspond to unmatched point x_i/y_j

Deriving A Robust Matching Model by Marginalization

For algorithmic efficiency, instead of $(M^*, \mathbf{F}^*) = \arg \max_{M, \mathbf{F}} p(M, \mathbf{F} | X, Y, D)$ we will solve

$$\sum_a p(a, b) = p(b)$$

$$\mathbf{F}^* = \arg \max_{\mathbf{F}} p(\mathbf{F} | X, Y, D)$$

$$p_{\mathbf{F}|X,Y,D}(\mathbf{F}|X,Y,D) \mathbf{F} \quad (18)$$

$$p_{M,\mathbf{F}|X,Y,D}(M, \mathbf{F}|X,Y,D)$$

by marginalization of $p(M, \mathbf{F} | X, Y, D)$ over M

this simplification changes the problem!

$$p(M, \mathbf{F} | X, Y, D) \simeq p(M, \mathbf{F}, X, Y, D) = p(X, Y, D, M | \mathbf{F}) \cdot p(\mathbf{F})$$

assuming correspondence-wise independence:

$$p(X, Y, D, M | \mathbf{F}) = \prod_{i=1}^m \prod_{j=1}^n p(x_i, y_j, D, m_{ij} | \mathbf{F}) \stackrel{\text{def}}{=} \prod_{i=1}^m \prod_{j=1}^n p_e(e_{ij}, d_{ij}, m_{ij} | \mathbf{F})$$

- e_{ij} represents geometric error for match $x_i \leftrightarrow y_j$: $e_{ij}(x_i, y_j | \mathbf{F})$
- d_{ij} represents descriptor similarity for match $x_i \leftrightarrow y_j$: $d_{ij} = \|\mathbf{d}(x_i) - \mathbf{d}(y_j)\|$

SIFT etc

Marginalization:

$$\sum_{m_{11} \in \{0,1\}} \sum_{m_{12}} \cdots \sum_{m_{mn}} p(X, Y, D, M | \mathbf{F}) = \sum_{m_{11}} \sum_{m_{12}} \cdots \sum_{m_{mn}} \prod_{i=1}^m \prod_{j=1}^n p_e(e_{ij}, d_{ij}, m_{ij} | \mathbf{F}) =$$

$$\sum_{m_1} \sum_{m_2} (p_1(m_1) p_2(m_2))$$

$$\left(\sum_{m_1} p_1(m_1) \right) \left(\sum_{m_2} p_2(m_2) \right)$$

$$= \cdots = \prod_{i=1}^m \prod_{j=1}^n \sum_{m_{ij} \in \{0,1\}} p_e(e_{ij}, d_{ij}, m_{ij} | \mathbf{F}) = p(X, Y, D | \mathbf{F})$$

we will continue with this term

Robust Matching Model (cont'd)

$$\begin{aligned}
 \sum_{m_{ij} \in \{0,1\}} p_e(e_{ij}, d_{ij}, m_{ij} | \mathbf{F}) &= \sum_{m_{ij} \in \{0,1\}} p_e(e_{ij}, d_{ij} | m_{ij}, \mathbf{F}) \cdot p(m_{ij} | \mathbf{F}) = \\
 &= \underbrace{p_e(e_{ij}, d_{ij} | m_{ij} = 1, \mathbf{F})}_{p_1(e_{ij}, d_{ij} | \mathbf{F})} \cdot \underbrace{p(m_{ij} = 1 | \mathbf{F})}_{1 - \alpha_0} + \underbrace{p_e(e_{ij}, d_{ij} | m_{ij} = 0, \mathbf{F})}_{p_0(e_{ij}, d_{ij} | \mathbf{F})} \cdot \underbrace{p(m_{ij} = 0 | \mathbf{F})}_{\alpha_0} = \\
 &= (1 - \alpha_0) \left[p_1(e_{ij}, d_{ij} | \mathbf{F}) + \frac{\alpha_0}{1 - \alpha_0} p_0(e_{ij}, d_{ij} | \mathbf{F}) \right] \quad (19)
 \end{aligned}$$

- the $p_0(e_{ij}, d_{ij} | \mathbf{F}) \approx \text{const}$ is a penalty for 'missing a correspondence' but it should be a p.d.f. (cannot be a constant) (see Slide 108 for a simplification)

$$\alpha_0 \rightarrow 1, \quad p_0 \rightarrow 0 \quad \text{so that} \quad \frac{\alpha_0}{1 - \alpha_0} p_0 \approx \text{const}$$

- the $p_1(e_{ij}, d_{ij} | \mathbf{F})$ is typically an easy-to-design component: assuming independence of geometric error and descriptor similarity:

$$p_1(e_{ij}, d_{ij} | \mathbf{F}) = p_1(e_{ij} | \mathbf{F}) \cdot p_1(d_{ij})$$

- we choose, eg.

$$p_1(e_{ij} | \mathbf{F}) = \frac{1}{T_e(\sigma_1, \mathbf{F})} e^{-\frac{e_{ij}^2(\mathbf{F})}{2\sigma_1^2}}, \quad p_1(d_{ij}) = \frac{1}{T_d(\sigma_d, \text{dim } \mathbf{d})} e^{-\frac{\|\mathbf{d}(x_i) - \mathbf{d}(y_j)\|^2}{2\sigma_d^2}} \quad (20)$$

- $\sigma_1, \sigma_d, \alpha_0$ are 'hyper-parameters'
- the form of $T(\sigma, \mathbf{F})$ depends on error definition
- we will continue with the result from (19)

$$= \sqrt{2\sigma_d^2} \sigma_d^{\text{dim } \mathbf{d}}$$

$$S = \sigma_d^2 \mathbf{I}_{\text{dim } \mathbf{d}}$$

► Simplified Robust Energy (Error) Function

- assuming the choice of p_1 as in (20), we are simplifying

$$p(X, Y, D | \mathbf{F}) = \prod_{i=1}^m \prod_{j=1}^n \left[(1 - \alpha_0) p_1(e_{ij}, d_{ij} | \mathbf{F}) \oplus \alpha_0 p_0(e_{ij}, d_{ij} | \mathbf{F}) \right] \quad (21)$$

- we define 'energy' as: $V(x) = -\log p(x)$ $H = -\int p(x) \log p(x) dx$ this helps simplify the formulas
entropy = mean energy
- for simplicity, we omit d_{ij}
- we choose $\sigma_0 \gg \sigma_1$ and the missed-correspondence penalty function as

$$p_0(e_{ij} | \mathbf{F}) = \frac{1}{T_e(\sigma_0, \mathbf{F})} e^{-\frac{e_{ij}^2(\mathbf{F})}{2\sigma_0^2}} \quad \sigma_0 \rightarrow \infty$$

- then

$$V(X, Y, D | \mathbf{F}) = \sum_{i=1}^m \sum_{j=1}^n \left[\underbrace{-\log \frac{1 - \alpha_0}{T_e(\sigma_1, \mathbf{F})}}_{\Delta(\mathbf{F})} - \log \left(e^{-\frac{e_{ij}^2(\mathbf{F})}{2\sigma_1^2}} \oplus \frac{\alpha_0}{1 - \alpha_0} \frac{T_e(\sigma_1, \mathbf{F})}{T_e(\sigma_0, \mathbf{F})} e^{-\frac{e_{ij}^2(\mathbf{F})}{2\sigma_0^2}} \right) \right] \quad t \approx \text{const}$$


- by choosing representative of \mathbf{F} such that $\Delta(\mathbf{F}) = \text{const}$, we get

$$V(X, Y, D | \mathbf{F}) = m n \Delta + \sum_{i=1}^m \sum_{j=1}^n \underbrace{-\log \left(e^{-\frac{e_{ij}^2(\mathbf{F})}{2\sigma_1^2}} + t \right)}_{\hat{V}(e_{ij})} \quad (22)$$

note that m, n are fixed

► The Action of the Robust Matching Model on Data

Example for $\hat{V}(e)$ from (22): $e_{i,j}^2(\mathbf{F})$



red – the usual (non-robust) error when $t = 0$
 blue – the rejected correspondence penalty t
 green – ‘robust energy’ (22)

- if the error of a correspondence exceeds a limit, it is ignored
- then $\hat{V}(e) = \text{const}$ and we essentially count outliers in (22)
- t controls the ‘turn-off’ point
- the inlier/outlier threshold is e_T is the error for which $(1 - \alpha_0) p_1(e_T) = \alpha_0 p_0(e_T)$: note that $t \approx 0$

$$e_T = \sigma_1 \sqrt{-\log t^2} \quad (23)$$


The full optimization problem is (18):

$$\mathbf{F}^* = \arg \max_{\mathbf{F}} p(\mathbf{F} | X, Y, D) = \arg \max_{\mathbf{F}} \frac{\overbrace{p(X, Y, D | \mathbf{F})}^{\text{likelihood}} \cdot \overbrace{p(\mathbf{F})}^{\text{prior}}}{\underbrace{p(X, Y, D)}_{\text{evidence}}} = \arg \min_{\mathbf{F}} \{V(X, Y, D | \mathbf{F}) + V(\mathbf{F})\}$$


- typically we take $V(\mathbf{F}) = 0$ unless we need to stabilize a computation, e.g. when video camera moves smoothly (on a high-mass vehicle) and we have a prediction for \mathbf{F}
- evidence is not needed unless we want to compare different models

Discussion: On The Art of Probabilistic Model Design...

- a few models for fitting zero-centered circle C of radius r to points in \mathbb{R}^2



How To Find the Global Maxima (Modes) of a PDF?



- consider the function $p(x)$ at left p.d.f. on $[0, 1]$, mode at 0.1
- consider several methods:

1. exhaustive search

$$\underset{x}{\operatorname{argmax}} p(x)$$

```

step = 1/(iterations-1);
for x = 0:step:1
    if p(x) > bestp
        bestx = x; bestp = p(x);
    end
end
    
```

- slow algorithm (definite quantization); faster variants exist
- fast to implement

2. randomized search with uniform sampling

```

x = rand(1); x = prnd;
if p(x) > bestp
    bestx = x; bestp = p(x);
end
    
```


- slow algorithm but better convergence
- fast to implement
- how to stop it?

3. random sampling from $p(x)$ (Gibbs sampler)

- faster algorithm
- fast to implement but often infeasible (e.g. when $p(x)$ is data dependent (our case))


4. Metropolis-Hastings sampling

- almost as fast (with care)
- not so fast to implement
- rarely infeasible
- RANSAC belongs here



- averaged over 10^4 trials
- number of proposals before $|x - x_{\text{true}}| \leq \text{step}$
- uniform and Gibbs give the theoretical result

How To Generate Random Samples from a Complex Distribution?



- red: probability density function $p(x)$ of a toy distribution on the unit interval **target distribution**

$$p(x) = \sum_{i=1}^4 \alpha_i \text{Be}(x; \alpha_i, \beta_i), \quad \sum_{i=1}^4 \alpha_i = 1, \quad \alpha_i \geq 0$$

$$\text{Be}(x; \alpha, \beta) = \frac{1}{\text{B}(\alpha, \beta)} \cdot x^{\alpha-1} (1-x)^{\beta-1}$$


$(\alpha_1, \alpha_2, \alpha_3, \alpha_4) \quad (0.1, 0.2, 0.5, 0.2) \quad | \quad (0.1, 0.3, 0.8, 1.0)$


- note we can generate samples from this $p(x)$ **how?**
- suppose we cannot sample from $p(x)$ but we can sample from some 'simple' distribution, given the last sample x_0 (blue) **proposal distribution**


$$q(x | x_0) = \begin{cases} \text{U}_{0,1}(x) & \text{(independent) uniform sampling} \\ \text{Be}(x; \frac{x_0}{T} + 1, \frac{1-x_0}{T} + 1) & \text{'beta' diffusion (crawler) } T - \text{temperature} \\ p(x) & \text{(independent) Gibbs sampler} \end{cases}$$

- note we have unified all the random sampling methods on the previous slide
- how to transform proposal samples $q(x | x_0)$ to target distribution $p(x)$ samples?

Thank You









**OPPA European Social Fund
Prague & EU: We invest in your future.**
