Deterministic Sample Consensus with Multiple Match Hypotheses

Paul McIlroy
http://mi.eng.cam.ac.uk/~pmm33

Ed Rosten
http://mi.eng.cam.ac.uk/~er258

Simon Taylor
http://mi.eng.cam.ac.uk/~sjt59

Tom Drummond
http://mi.eng.cam.ac.uk/~twd20

Machine Intelligence Laboratory
Department of Engineering
University of Cambridge
Cambridge, UK

In computer vision, a common task is to estimate model parameters from a set of feature matches containing outliers using a hypothesize-and-verify approach [2]. We present a deterministic scheme for selecting samples that combines matching scores, ambiguity and past performance of hypotheses generated by the matches. At every stage the best matches are chosen to generate a hypothesis. The result is a system that is able to operate very efficiently on ambiguous data.

Multiple matches belonging to a single feature are mutually exclusive events, so two matches of equal quality score each have a probability ≤ 0.5 and are much less likely to be valid than a unique match with the same score. We exploit the additional information provided by the alternative match scores in order to compute a probability for each putative correspondence.
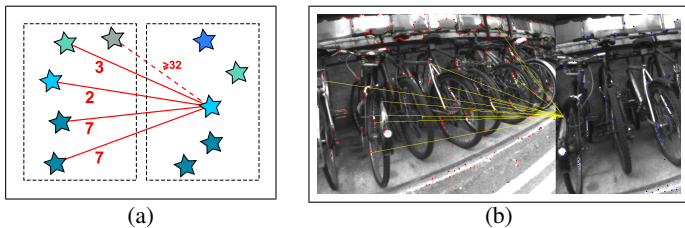

Figure 1: Multiple matches: (a) a discrete error score is calculated for each correspondence; (b) multiple matches from a single feature.

The frame-to-frame matching process is illustrated by Figure 1. Each feature in the current frame is compared to all n features in the previous frame. Each feature will generate at most one correct match. If one match is correct, the remaining $(n-1)$ match scores are generated by incorrect correspondences. Alternatively, all n match scores may be generated by incorrect correspondences due to occlusion.

Each false match generates a match score drawn from a probability distribution. The match scores are generated by each putative correspondence as independent trials and the probability that n false matches will generate a particular set of match scores is given by the multinomial distribution,

$$\text{Mult}(m_1, m_2, \ldots, m_K | \mu, n) = \frac{n!}{m_1! m_2! \cdots m_K!} \prod_{k=1}^{K} \mu_k^{m_k}, \qquad (1)$$

where m_k is the number of false matches that generated a match with score k. The parameters $\mu = (\mu_1, \ldots, \mu_K)^\top$ give the probability of a single false match generating match score k, subject to $0 \leq \mu_k \leq 1$ and $\sum_k \mu_k = 1$.
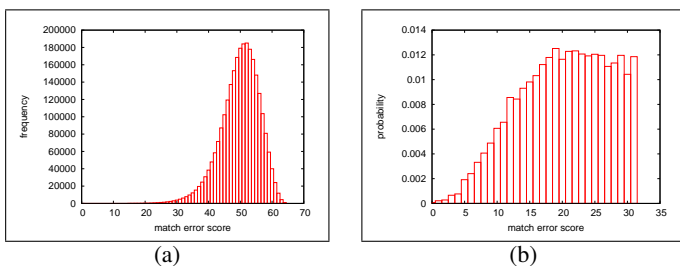

Figure 2: Probability distributions for error scores generated by (a) false matches; (b) correct matches.

Several competing explanations are compatible with the set of observed match scores for each feature. Let $P(c_k)$ be the probability that a feature has a correct correspondence with match score k and $P(c_\emptyset)$ be the probability that the feature has no correct correspondence in the previous image. The probability of the event E_j, one correct match with score j and $(n-1)$ false matches, is given by

$$P(E_j) = P(c_j)\text{Mult}(m_1, m_2, \ldots, m_j - 1, \ldots, m_K | \mu, n-1). \qquad (2)$$

$$P(E_\emptyset) = P(c_\emptyset)\text{Mult}(m_1, m_2, \ldots, m_K | \mu, n) \qquad (3)$$

is the probability of event $E_\emptyset$, no correct match and n false matches. The posterior probability $P(c_i | m_1, \ldots, m_K)$ that a correct match exists with error i given the observed set of match scores is therefore

$$P(c_i | m_1, \ldots, m_K) = \frac{P(E_i)}{\sum_{j=0}^{K} P(E_j) + P(E_\emptyset)}. \qquad (4)$$

The parameters μ are obtained by considering *intra*frame matches. Matches between two distinct features in the same image are guaranteed to be false. False matches are easily harvested from video to obtain the false match distribution in Figure 2.
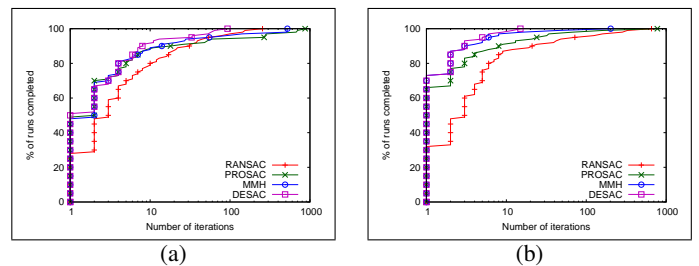

Figure 3: Time-to-solution comparison: (a) indoor; (b) hedge.

PROSAC [1] draws each hypothesis from a gradually expanding pool of the best samples ordered by match score. The predetermined rate at which this pool is expanded is designed to provide a balance between reliance on the initial sorting and the RANSAC approach which treats all correspondences as equally likely. PROSAC only exploits the ordering of the samples by match score. If more accurate probabilities are available (by taking into account alternative match hypotheses) these can be used to govern the rate at which new samples are tried.

If a hypothesis generated from a set of samples fails to generate consensus, then it is likely that at least one of the samples used is a false match. This can be used to update (by reducing) the probability that each match being considered is correct. Consider a minimal set consisting of four matches A, B, C and D. The prior probability of the minimal set containing only correct matches is given by $P(A \cap B \cap C \cap D)$. The outcome of a failed test removes this event from the joint distribution and the posterior probability that match A is correct becomes

$$P(A | \bar{A} \cup \bar{B} \cup \bar{C} \cup \bar{D}) = \frac{P(A) - P(A \cap B \cap C \cap D)}{1 - P(A \cap B \cap C \cap D)}. \qquad (5)$$

The updated matches are removed from and reinserted into the list so that it remains sorted by match probability. At each iteration we always choose the most likely matches at the top of the list.

[1] O. Chum and J. Matas. Matching with PROSAC-progressive sample consensus. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2005.

[2] M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981.