Joint Probabilistic Matching Using m-Best Solutions

Seyed Hamid Rezatofighi[1] Anton Milan[1] Zhen Zhang[2] Qinfeng Shi[1] Anthony Dick[1] Ian Reid[1]

[1]Australian Centre for Visual Technologies (ACVT), School of Computer Science, The University of Adelaide, Australia

[2]School of Computer Science and Technology, Northwestern Polytechnical University, Xian, China

WELCOME TO Fabulous CVPR 2016 LAS VEGAS

Motivation & Contribution

Graph matching is typically approached by solving a MAP problem, e.g. maximizing a joint matching score.

We argue that

- Globally optimal solution may or may not be easily achieved,
- Even the optimal solution does not necessarily yield the correct matching assignment

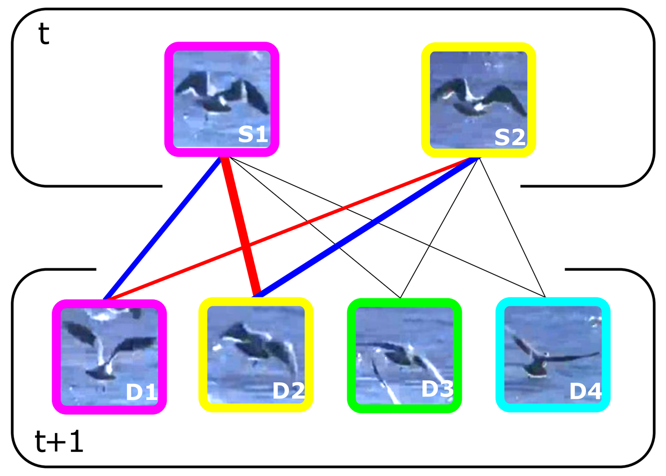

To improve matching results, we propose to use...

Approx. marginals instead of MAP:
m-best soutions instead of a single solution.

One-to-One Graph Matching

Bipartite, K-Partite and Multipartite Graphs

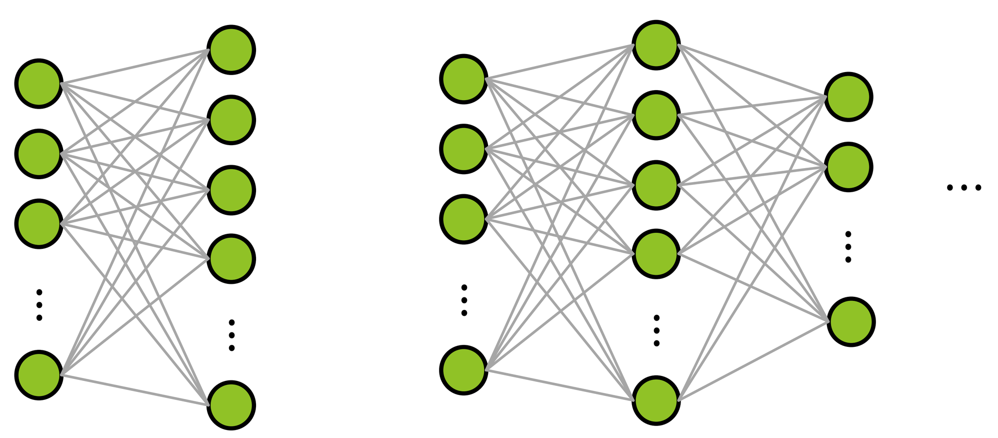

A constrained binary program

Maximizing (or minimising) a joint matching probability $p(\cdot)$ (or objective cost $f(\cdot)$)

$$X^* = \underset{X\in\mathcal{X}}{\text{argmax}}\; p(X), \quad \underset{X\in\mathcal{X}}{\text{argmin}}\; f(X), \qquad (1)$$

where $\mathcal{X}$ is the one-to-one matching space

$$\mathcal{X} = \left\{X = \left(x_i^j\right)_{\forall i,j} \;\middle|\; x_i^j \in \{0,1\}, \forall j: \sum_i x_i^j \leqslant 1, \forall i: \sum_j x_i^j = 1\right\}$$

A linear inequality constraint
$AX \leqslant B$

Marginalization vs. MAP

MAP estimate ingnores underlying distribution and picks only one solution.

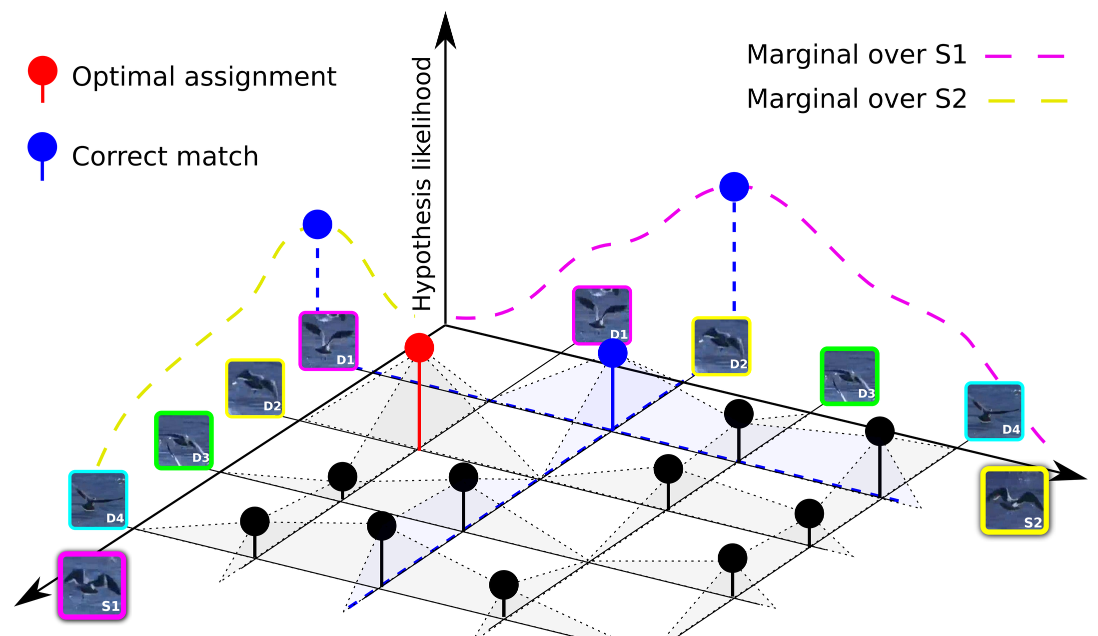

Marginalization, a safer choice

- Encodes the entire distribution to untangle potential ambiguities,
- Improves matching ranking due to averaging / smoothing property

$$p(x_i^j = 1) = \sum_{\{X\in\mathcal{X}|x_i^j=1\}} p(X),$$
$$c_i^j = -\log \sum_{\{X\in\mathcal{X}|x_i^j=1\}} e^{-f(X)}. \qquad (2)$$

Exact marginalization is NP-hard:
It requires all feasible solutions to build the distribution.

Approximation using m-best solutions

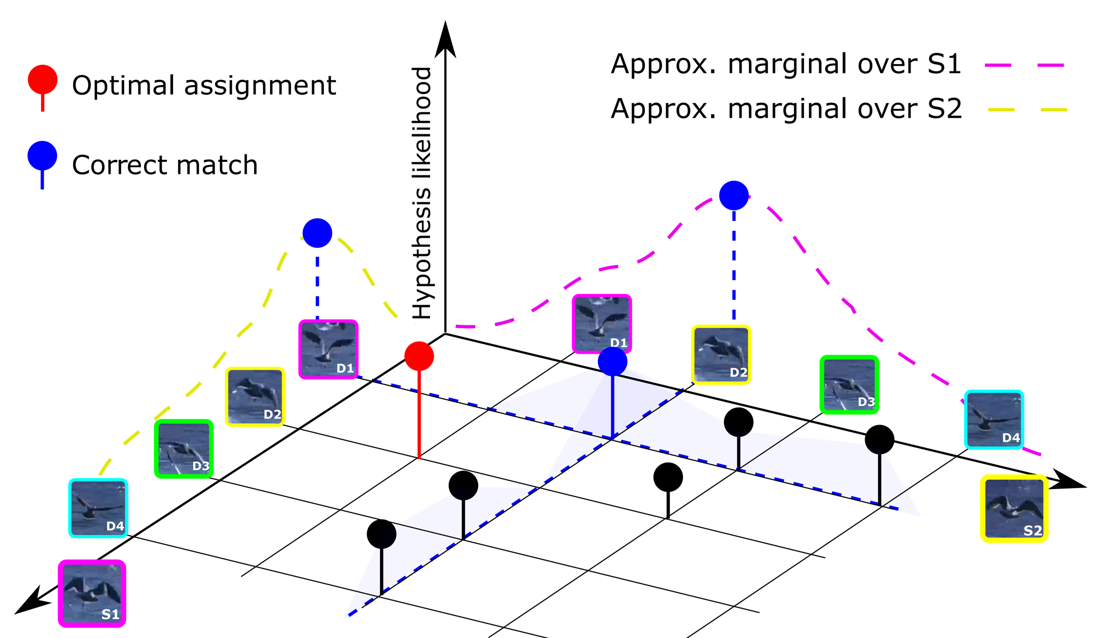

Computing m-Best Solutions

Naive exclusion strategy

- General approach,
- Impractical for large values of m

$$X_m^* = \underset{}{\text{argmin}}\; f(X) \quad AX \leqslant B \quad \hat{A}X \leqslant \hat{B}$$

Binary tree partitioning [5]

Partition the space into a set of disjoint subspaces

- Efficient approach,
- Not a good strategy for weak solvers

Experimental Results

Applications with linear objectives $X^* = \underset{X\in\mathcal{X}}{\text{argmin}}\; C^\top X$

Person Re-Identification

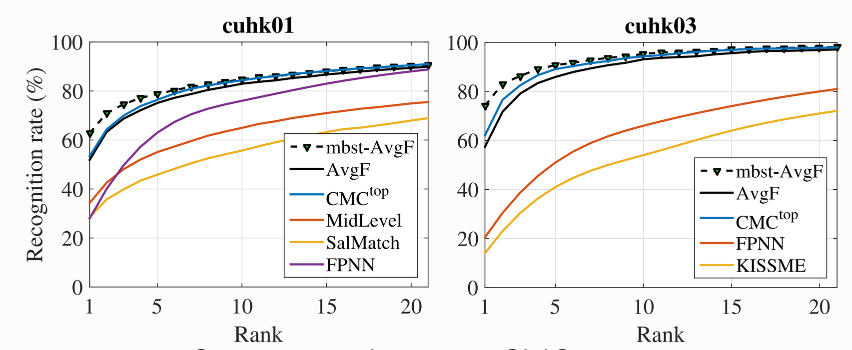

Comparison between CMC curves

Dataset (size)	Method	Recognition rate % Rank-1	Rank-2	Rank-5	Time (Sec.)
RAiD (20 × 20)	FT [1]	74.0	82.0	96.0	
	mbst-FT	**85.0**	**99.0**	**100.0**	1.6
WARD (35 × 35)	FT [1]	50.3	70.9	88.0	
	mbst-FT	**72.0**	**81.1**	**92.6**	4.2
iLIDS (59 × 59)	AvgF [4]	51.9	60.7	72.4	
	mbst-AvgF	**54.7**	**63.6**	**75.4**	15.4
3DPeS (96 × 96)	AvgF [4]	53.6	64.1	76.9	
	mbst-AvgF	**57.5**	**67.9**	**79.5**	31.8
VIPeR (316 × 316)	AvgF [4]	44.9	58.3	76.3	
	mbst-AvgF	**50.5**	**63.0**	**78.0**	201.9
CUHK01 (485 × 485)	AvgF [4]	51.9	63.3	75.1	
	mbst-AvgF	**62.8**	**70.9**	**78.8**	485.6
CUHK03 (100 × 100)	AvgF [4]	57.4	71.7	85.9	
	mbst-AvgF	**74.2**	**83.1**	**90.7**	33.5

Sequential Re-Identification

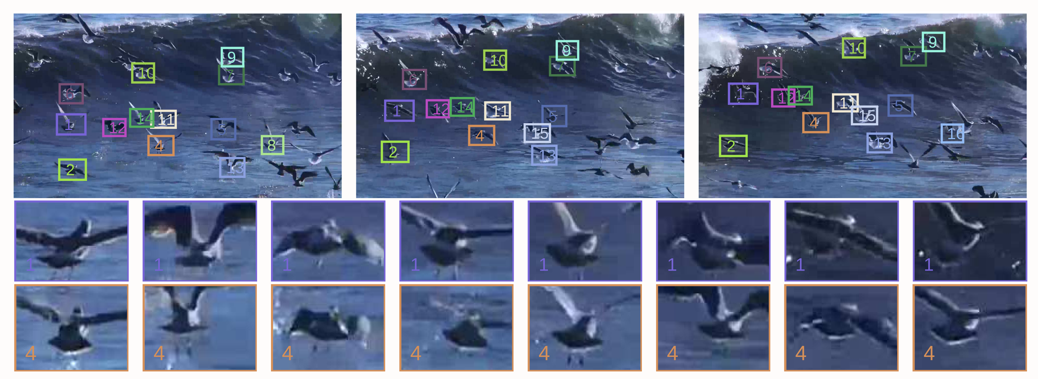

Method	Dicle *et al.* [2]	Our SeqReID
MOTA % ↑	92.46	**97.16**

Experimental Results

Application with a quadratic objective $X^* = \underset{X\in\mathcal{X}}{\text{argmax}}\; X^\top K X$

Feature Matching

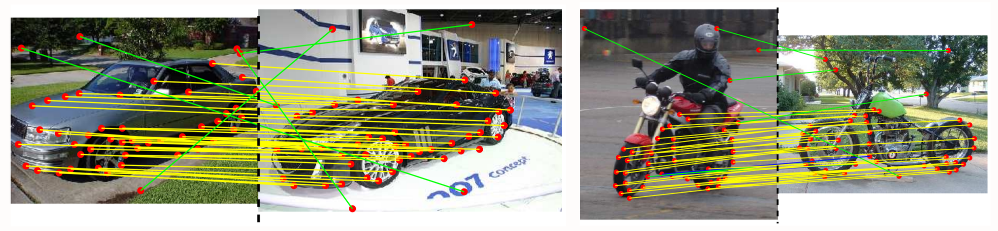

We used two different solvers (IPFP [3] and BP [6]) for this application.

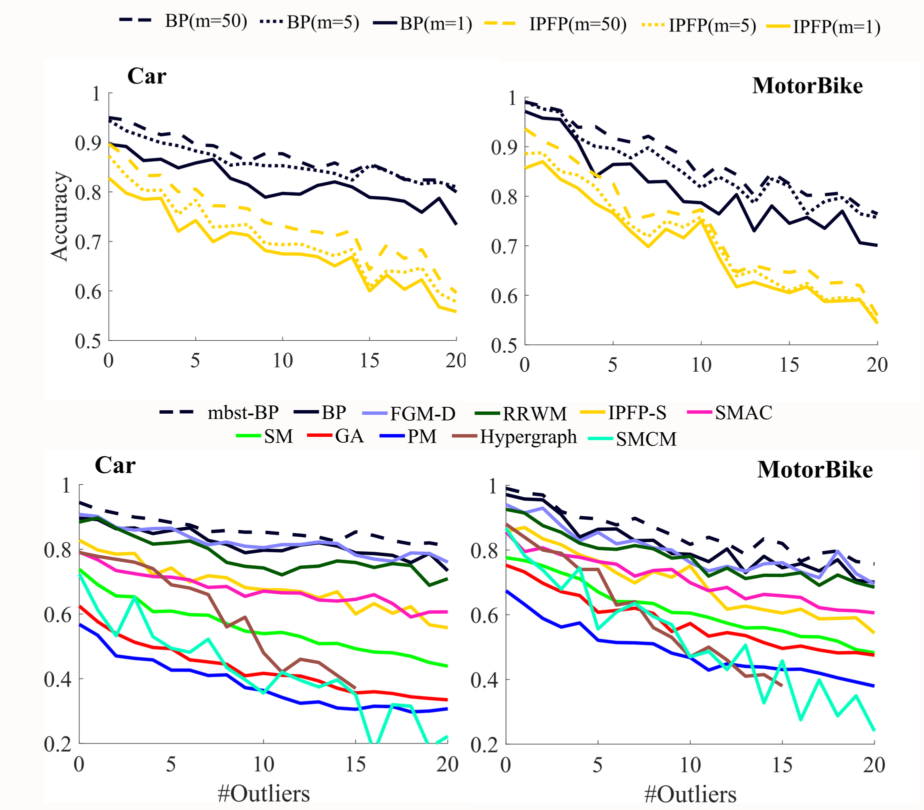

Discussion

- No apparent correlation between similarity of solutions and their contribution toward accuracy,
- Finding a solution with a better objective

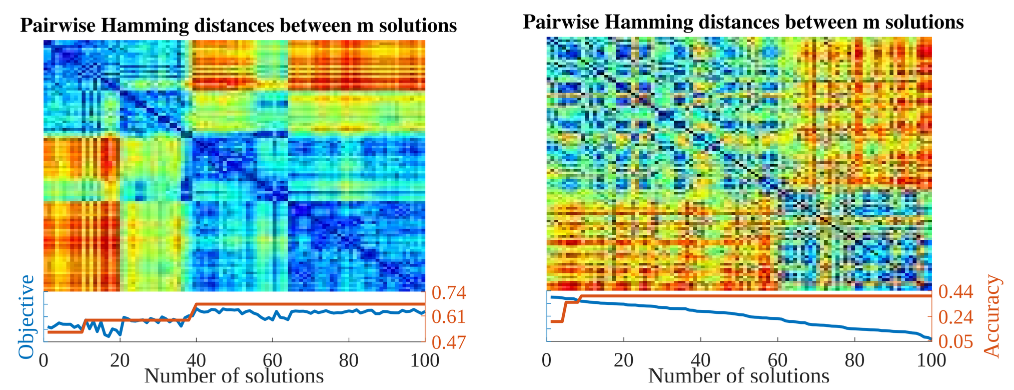

References

[1] A. Das, A. Chakraborty, and A. K. Roy-Chowdhury. Consistent re-identification in a camera network. In *ECCV 2014*.

[2] C. Dicle, O. I. Camps, and M. Sznaier. The way they move: Tracking multiple targets with similar appearance. In *ICCV*, 2013.

[3] M. Leordeanu, R. Sukthankar, and M. Hebert. Unsupervised learning for graph matching. *IJCV*, 96(1):28–45, Apr. 2011.

[4] S. Paisitkriangkrai, C. Shen, and A. v. d. Hengel. Learning to rank in person re-identification with metric ensembles. In *CVPR*, 2015.

[5] H. S. Rezatofighi, A. Milan, Z. Zhang, Q. Shi, A. Dick, and I. Reid. Joint probabilistic data association revisited. In *ICCV 2015*.

[6] Z. Zhang, Q. Shi, J. McAuley, W. Wei, Y. Zhang, and A. van den Hengel. Pairwise matching through Max-Weight bipartite belief propagation. In *CVPR 2016*.

We gratefully acknowledge the financial support of the Australian Research Council Linkage Project LP130100154R and ARC Laureate Fellowship FL130100102.