Learning Near-Isometric Matching Models

Julian McAuley, Tibério Caetano
julian.mcauley@nicta.com.au tiberio.caetano@nicta.com.au

Abstract

Shape matching problems are often described in terms of structural properties, encoding the 'elasticity' of certain joints in a rigid object. Encoding such constraints often results in a hard optimization problem, presenting a problem for many learning approaches that require exact inference schemes as a subroutine. We note that certain types of constraints, such as isometries, result in 'easier' optimization problems that can be solved with **low tree-width graphical models**. This allows us to apply **learning in near-isometric matching scenarios**, encoding rich first-order properties, as well as isometric and topological information.

Matching objectives

The well-known 'quadratic assignment' problem takes the form

$$\hat{f} = \operatorname*{argmin}_{f:A\to B} \sum_{a,b\in A} D_{a,b,f(a),f(b)}.$$

Many structural matching problems can be expressed in this form, such as *isometric* matching:

$$\hat{f} = \operatorname*{argmin}_{f:\mathcal{S}\to\mathcal{T}} \sum_{s_i,s_j\in\mathcal{S}} \left|d(s_i,s_j) - d(f(s_i),f(s_j))\right|,$$

where $\mathcal{S}$ and $\mathcal{T}$ represent *shapes*, f maps points in $\mathcal{S}$ to points in $\mathcal{T}$, and d is a distance function. In cases where this equation has a zero-cost solution, we note that far more efficient solutions to this problem can be found: we need not consider all edges between pairs of points in $\mathcal{S}$, but rather a subset of edges that define a 'globally rigid' graph.

Some of our models with globally rigid embeddings

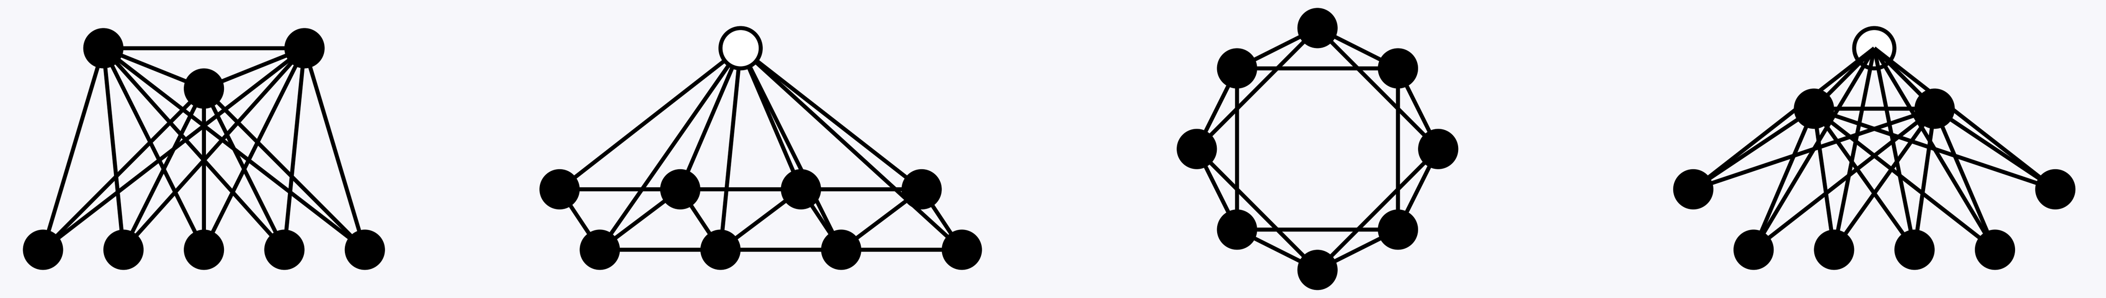

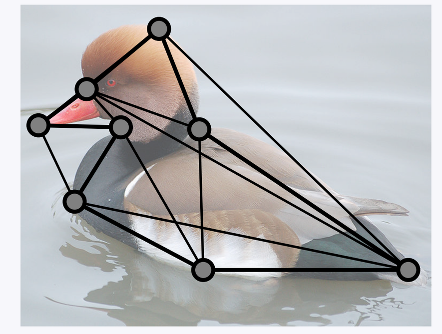

model:				
complexity:	$O(MN^4)$	$O(MN^3)$	$O(MN^2\sqrt{N})$	$O(MN^2 \log N)$
reference:	[CC06]	[MdCC10]	[MCS08]	

Some of our graphical models and their associated running times (for a template of size $|\mathcal{S}| = M$ in a scene of size $|\mathcal{T}| = N$); white nodes denote boolean variables. The nodes of this graph are 'embedded' in the plane, as shown for the model from [MCS08] at right. Our best methods are sub-cubic in N, meaning that we improve upon the running time implied by the tree-widths of the graphs.

Structured matching objectives

We can augment our potentials to encode structural constraints other than distance preservation:

$$\hat{f} = \operatorname*{argmin}_{f:\mathcal{S}\to\mathcal{T}} \sum_{i,j\in\mathcal{G}} \langle \Phi_{i,j}(s_i,s_j,f(s_i),f(s_j)),\theta\rangle.$$

Here Φ is a *feature vector*, which could include topological information, and first-order properties such as Shape Contexts or SIFT features; we can also create higher-order features $\Phi_{i,j,k}$ encoding angle and scale information. θ is a *parameter vector*, chosen using the *structured learning* framework of [THJA04]. Our objective function for choosing the best $\hat{\theta}$ is

$$\hat{\theta} = \operatorname*{argmin}_{\theta} \underbrace{\frac{1}{K}\sum_{i=1}^{K}\Delta(\hat{f}^i, f^i)}_{\text{empirical risk}} + \underbrace{\frac{\lambda}{2}\|\theta\|_2^2}_{L_2\text{ regularizer}},$$

where $f^1 \dots f^K$ is our training set, $\Delta(\hat{f}^i, f^i)$ is a *loss function*, and λ is a regularization constant.

Our findings

We find that models based on rigid graphs – in which inference can be done exactly – significantly outperform approximate methods based on quadratic assignment once learning is applied. This confirms the need for efficient, exact inference procedures in structured learning settings. Our models also outperform first-order models based on Shape Contexts and SIFT features, confirming the need for high-order *structural* objectives.

[CC06] T. S. Caetano and T. Caelli, *A unified formulation of invariant point pattern matching.*, ICPR, 2006.

[MCS08] J. J. McAuley, T. S. Caetano, and A. J. Smola, *Robust near-isometric matching via structured learning of graphical models*, NIPS, 2008.

[MdCC10] J. J. McAuley, T. de Campos, and T. S. Caetano, *Unified graph matching in euclidean spaces*, CVPR, 2010.

[THJA04] I. Tsochantaridis, T. Hofmann, T. Joachims, and Y. Altun, *Support vector machine learning for interdependent and structured output spaces*, ICML, 2004.