Robust Near-Isometric Matching via Structured Learning of Graphical Models

Julian J. McAuley, Tibério S. Caetano, Alexander J. Smola

Statistical Machine Learning program, NICTA, and the Research School of Information Sciences and Engineering, ANU

Abstract

In our paper [1], we present a model for matching shapes in images that are related by **near-isometric transformations**. We do so by combining two previous approaches: we use the graphical model of [2], which quickly solves isometric matching problems in 2-d point sets. We use a **structured learning** approach similar to that in [3], which parametrises the matching scores used by linear and quadratic assignment. The result is an **efficient** and **exact** shape-matching algorithm, which uses machine learning to parametrise first, second, and **third-order** features.

Common approaches to the matching problem

Commonly, the problem of matching objects in images is expressed as a problem of matching graphs, which are attributed by high-dimensional local image descriptors. That is, we want to find a mapping y from nodes in a template graph $\mathcal{G}$ to nodes in a target graph $\mathcal{G}'$:

$$\underset{y}{\operatorname{argmin}} \sum_{i=1}^{|\mathcal{G}|} \underbrace{\Phi_1(g_i, y(g_i))}_{\text{node features}} + \sum_{i=1}^{|\mathcal{G}|} \sum_{j=1}^{|\mathcal{G}|} \underbrace{\Phi_2(g_i, g_j, y(g_i), y(g_j))}_{\text{edge features}}.$$

Unfortunately, this corresponds to the **quadratic assignment problem**, which is in general **NP-hard**.

Our approach

Instead, we use the recently proposed graphical model from [2], which can solve **near-isometric** matching problems **exactly** in $O(|\mathcal{G}||\mathcal{G}'|^3)$ time:

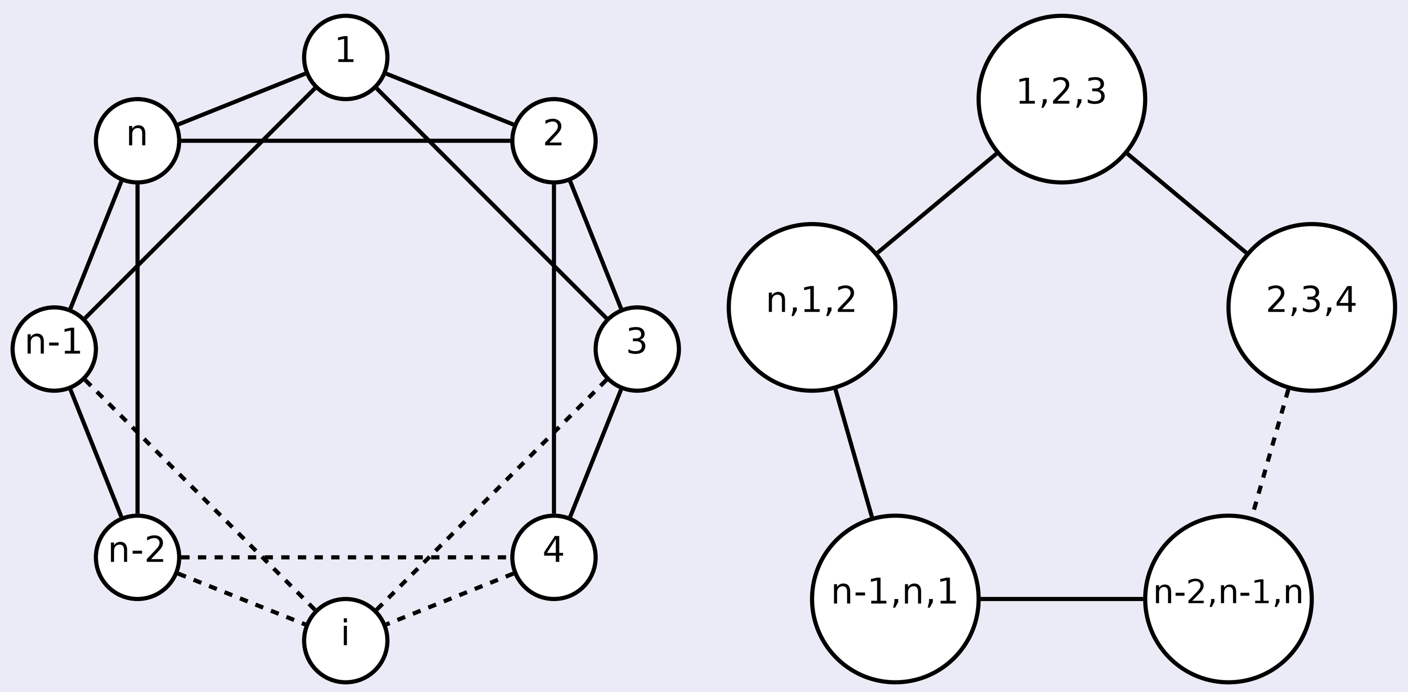

the graphical model from [2] (left), and its clique-graph (right)

Nodes in this model correspond to points in the template graph, and their assignments correspond to points in the target graph. Because the clique-graph contains only a single loop, **loopy belief-propagation** in this graph will converge to the optimal solution.

Our energy function

The energy minimisation problem we now want to solve is

$$\underset{y}{\operatorname{argmin}} \sum_{i=1}^{|\mathcal{G}|} \underbrace{\Phi_3(g_i, g_{i+1}, g_{i+2}, y(g_i), y(g_{i+1}), y(g_{i+2}))}_{\text{third-order features}}.$$

Φ_3 allows us to include first-order features (such as SIFT, Shape Context), second-order features (such as adjacency information, distances), and **third-order** features (such as angles, triangle similarity).

Matching 'shapes' and 'point-patterns'

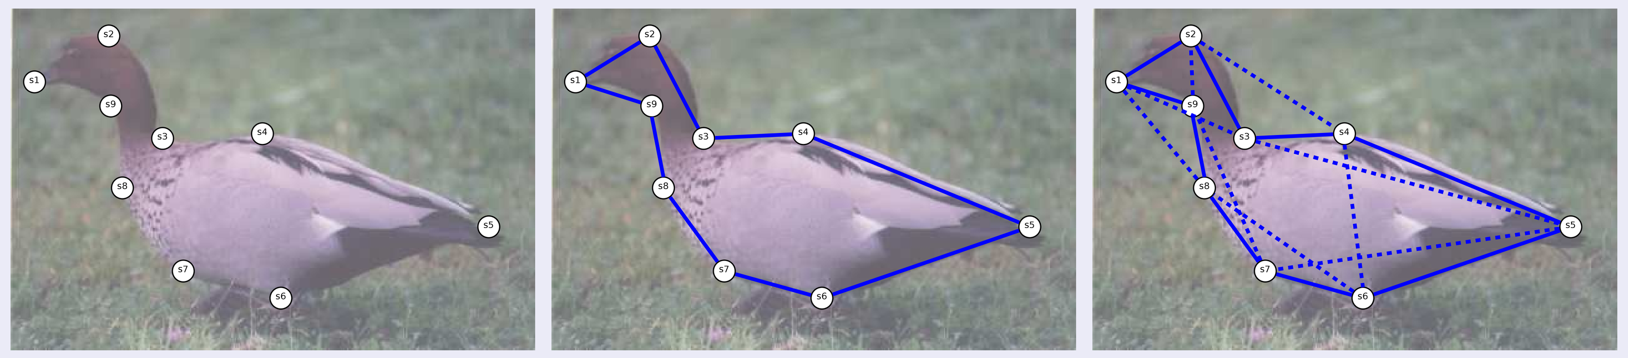

When the objects being matched are characterised by their shape, our model seems to capture exactly the desired dependencies.

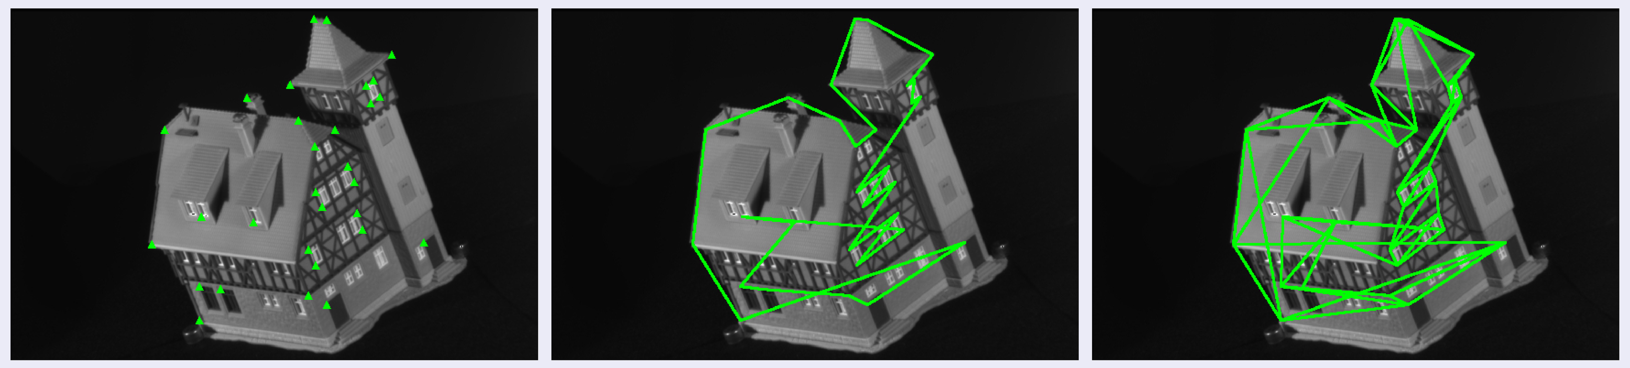

In other cases, we capture only a fraction of the desired dependencies, but we benefit from the addition of third-order features, and from the fact that our model is exact.

Structured learning

Given a collection of labeled training matches, we use a **structured learning** approach similar to that of [3] to parametrise Φ_3. This allows us to determine the degree to which changes in appearance play a role, versus changes in rotation and scale.

An example match...

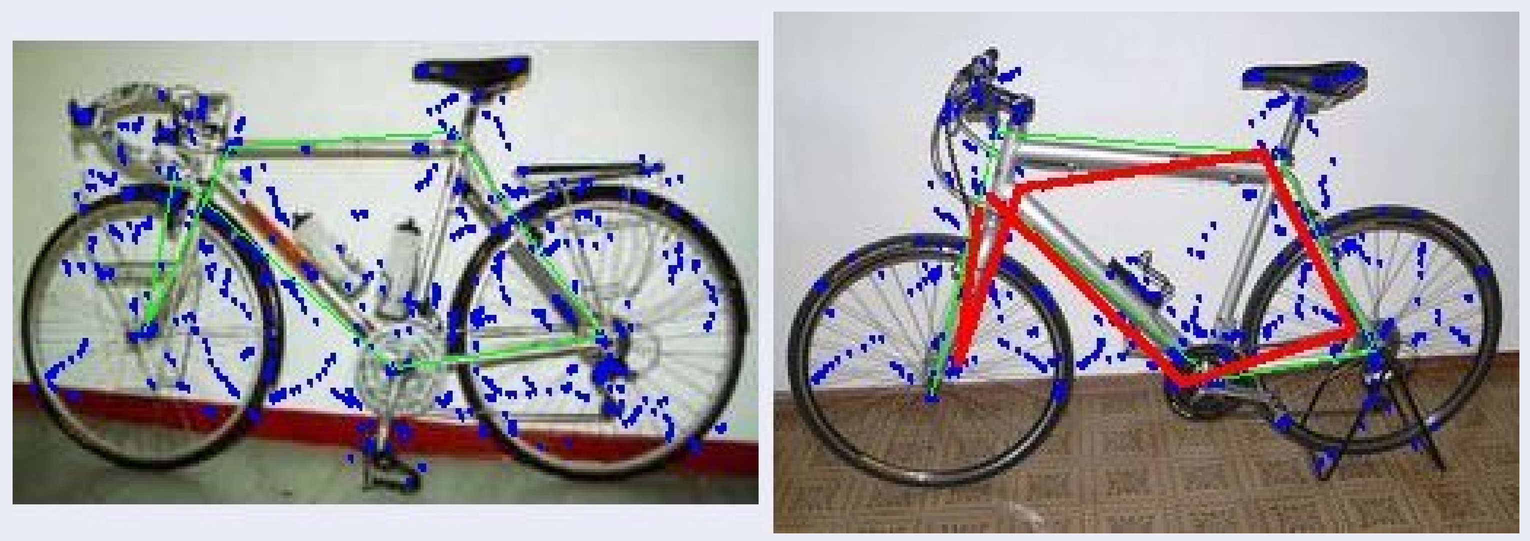

A labeled shape (left), and its match found in another scene. The correct label is shown in green, the inferred label using our model is shown in red.

Results on a video sequence

We compare our method to that of [3] (which parametrises features for linear and quadratic assignment) on frames in a video sequence. The proportion of incorrect labels is reported as the baseline (separation between frames) varies.

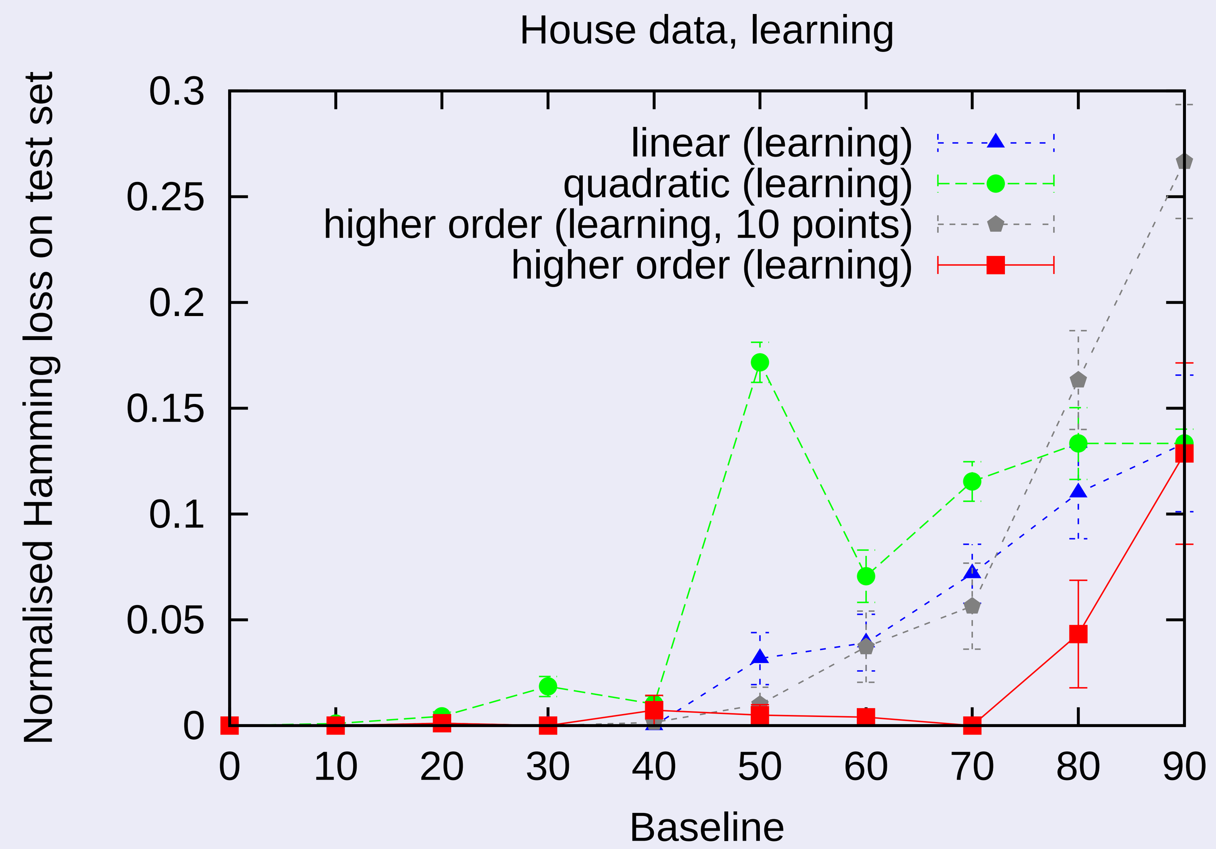

By considering only a '**likely subset**' of p points in the target graph, we reduce the running time from $O(|\mathcal{G}||\mathcal{G}'|^3)$ to $O(|\mathcal{G}|p^3)$. With $p = 10$, our algorithm is favourable in terms of performance and running time.

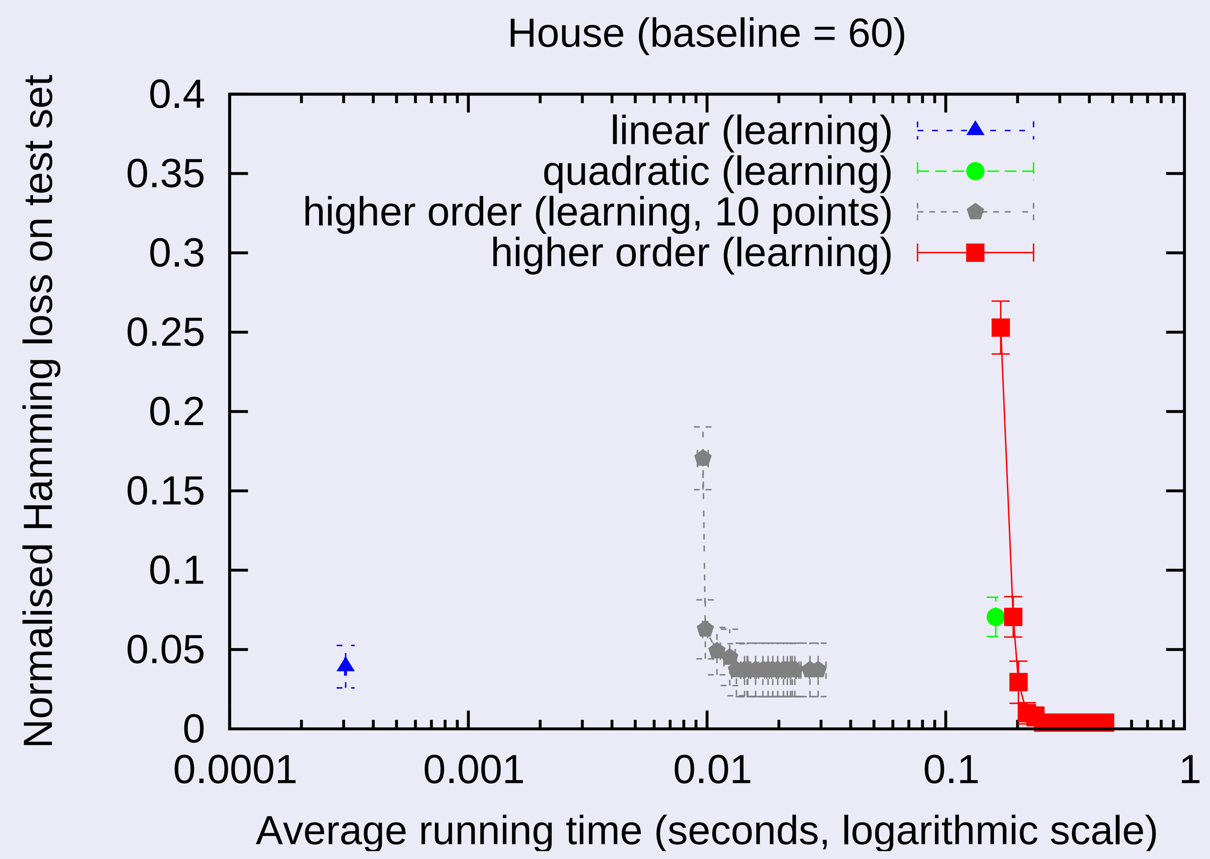

Bibliography

[1] J. J. McAuley, T. S. Caetano, and A. J. Smola.
Robust near-isometric matching via structured learning of graphical models.
NIPS, 2009.

[2] J. J. McAuley, T. S. Caetano, and M. S. Barbosa.
Graph rigidity, cyclic belief propagation and point pattern matching.
PAMI, 30(11):2047–2054, 2008.

[3] T.S. Caetano, L. Cheng, Q.V. Le, and A.J. Smola.
Learning graph matching.
In *ICCV*, pages 1–8, 2007.

http://users.rsise.anu.edu.au/~julianm/ julian.mcauley@nicta.com.au, tiberio.caetano@nicta.com.au, alex@smola.org