Shape Classification Through Structured Learning of Matching Measures

Longbin Chen[1], Julian J. McAuley[2], Rogerio S. Feris[3], Tibério S. Caetano[2], Matthew Turk[1]

1: University of California, Santa Barbara
2: Statistical Machine Learning group, NICTA, and the Australian National University
3: IBM T. J. Watson Research Center

Abstract

Many traditional methods for **shape classification** use **shape matching** scores as similarity measures. Previously, learning has only been applied to this process after the matching scores have been obtained. In our paper, instead of simply taking the matching scores for granted and then learning a classifier, we **learn the matching scores** themselves so as to produce shape similarity scores that **minimize the classification loss**.

Our approach, compared to existing approaches

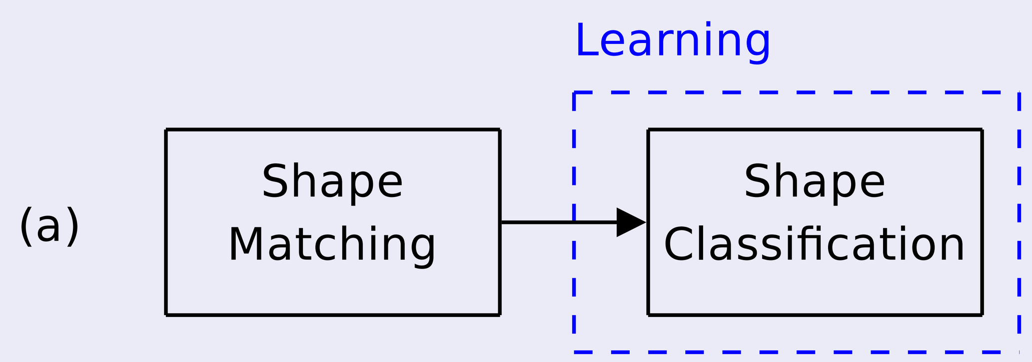

Conventional methods only apply learning **after** matching scores have been obtained.

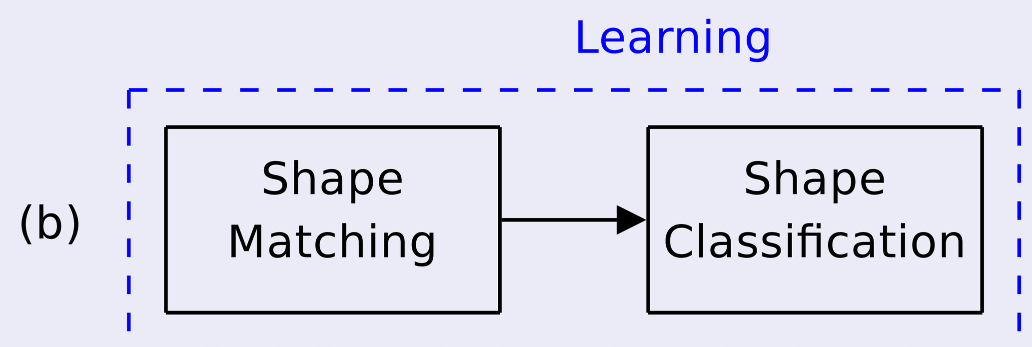

Instead, we apply learning to matching itself, so as to minimize the classification loss.

Problems with existing approaches

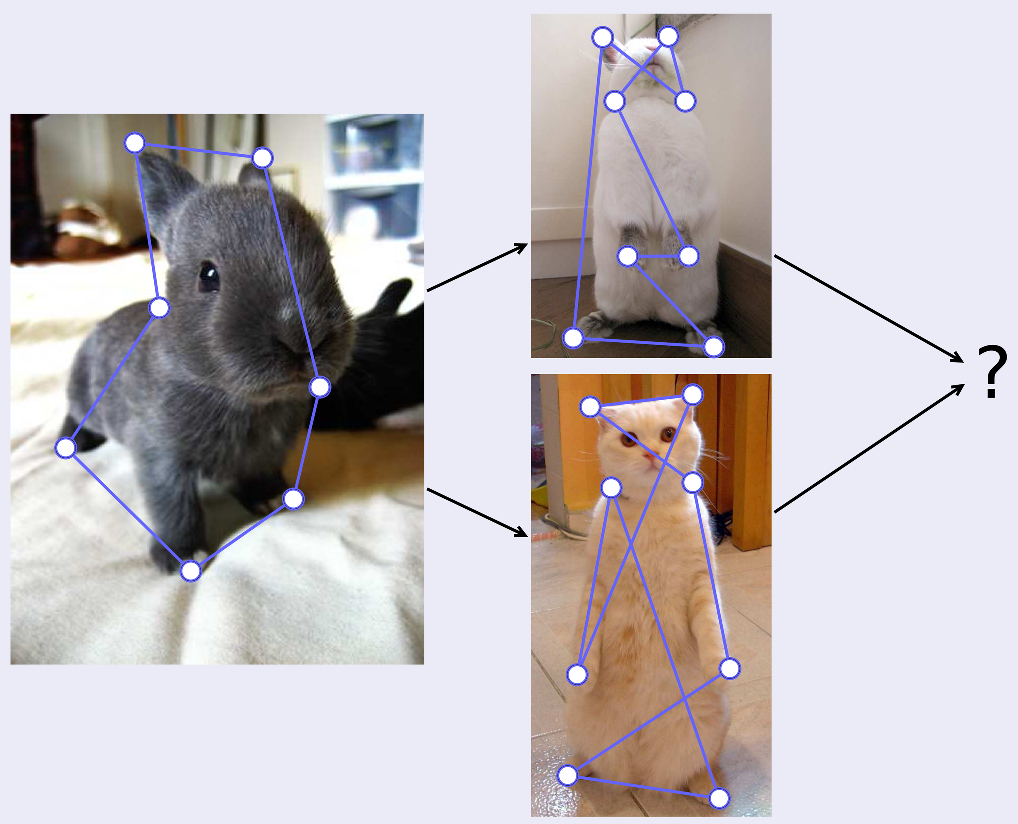

Performing learning *after* matching scores have been obtained makes little sense if the matches themselves are poor. Alternately, even if matches between objects of the *same* class are good, the method may not be good at **discriminating between objects** of different classes.

Our risk function, compared to existing risk functions

$$\sum_i \Delta\Big[\overbrace{\mathrm{class}\Big(\operatorname*{argmin}_{\mathcal{G}} \theta_{\mathrm{class}(\mathcal{G})} \underbrace{\min_y \sum_{i=1}^{|\mathcal{G}|} \|\phi(g_i) - \phi(y(g_i))\|_2^2}_{\text{linear assignment objective}}\Big)}^{\text{class of the graph with the lowest linear assignment score}}, \overbrace{\mathrm{class}(y^i)}^{\text{training label}}\Big]$$

In the conventional setting, θ assigns a weight to each class. Δ is a $0/1$ loss function indicating whether the chosen class is correct.

$$\sum_i \Delta\Big[\mathrm{class}\Big(\operatorname*{argmin}_{\mathcal{G}} \underbrace{\min_y \sum_{i=1}^{|\mathcal{G}|} \langle \phi(g_i) - \phi(y(g_i)), \theta\rangle^2}_{\textbf{parametrized}\text{ linear assignment objective}}\Big), \mathrm{class}(y^i)\Big]$$

We use the same loss, but **we parametrize the linear assignment objective itself**.

Structured learning

Whereas conventional learning scenarios only require *class labels* for training, our training data consists of **class labels and matches between graphs** of the same class.

This is an example of **structured learning**, and learning is done in the framework of [TJHA05].

Our experiments

MPEG-7: 70 Shape categories, 20 samples in each category [LLE00].

MNIST: 70,000 hand-written digits [LBH98].

The main difficulty we face with these datasets is to produce labeled correspondences between shapes of the same class for our training set. However, we show that even when using semi-automatic methods to produce correspondences, we still yield substantial benefits from learning.

Results

We demonstrate that both the conventional and the proposed formulation improve over non-learning results. However, the improvement of learning over non-learning is greater when the matching criterion is parametrized.

The framework presented in our paper potentially **applies to any classification method based on graph matching**.

Examples

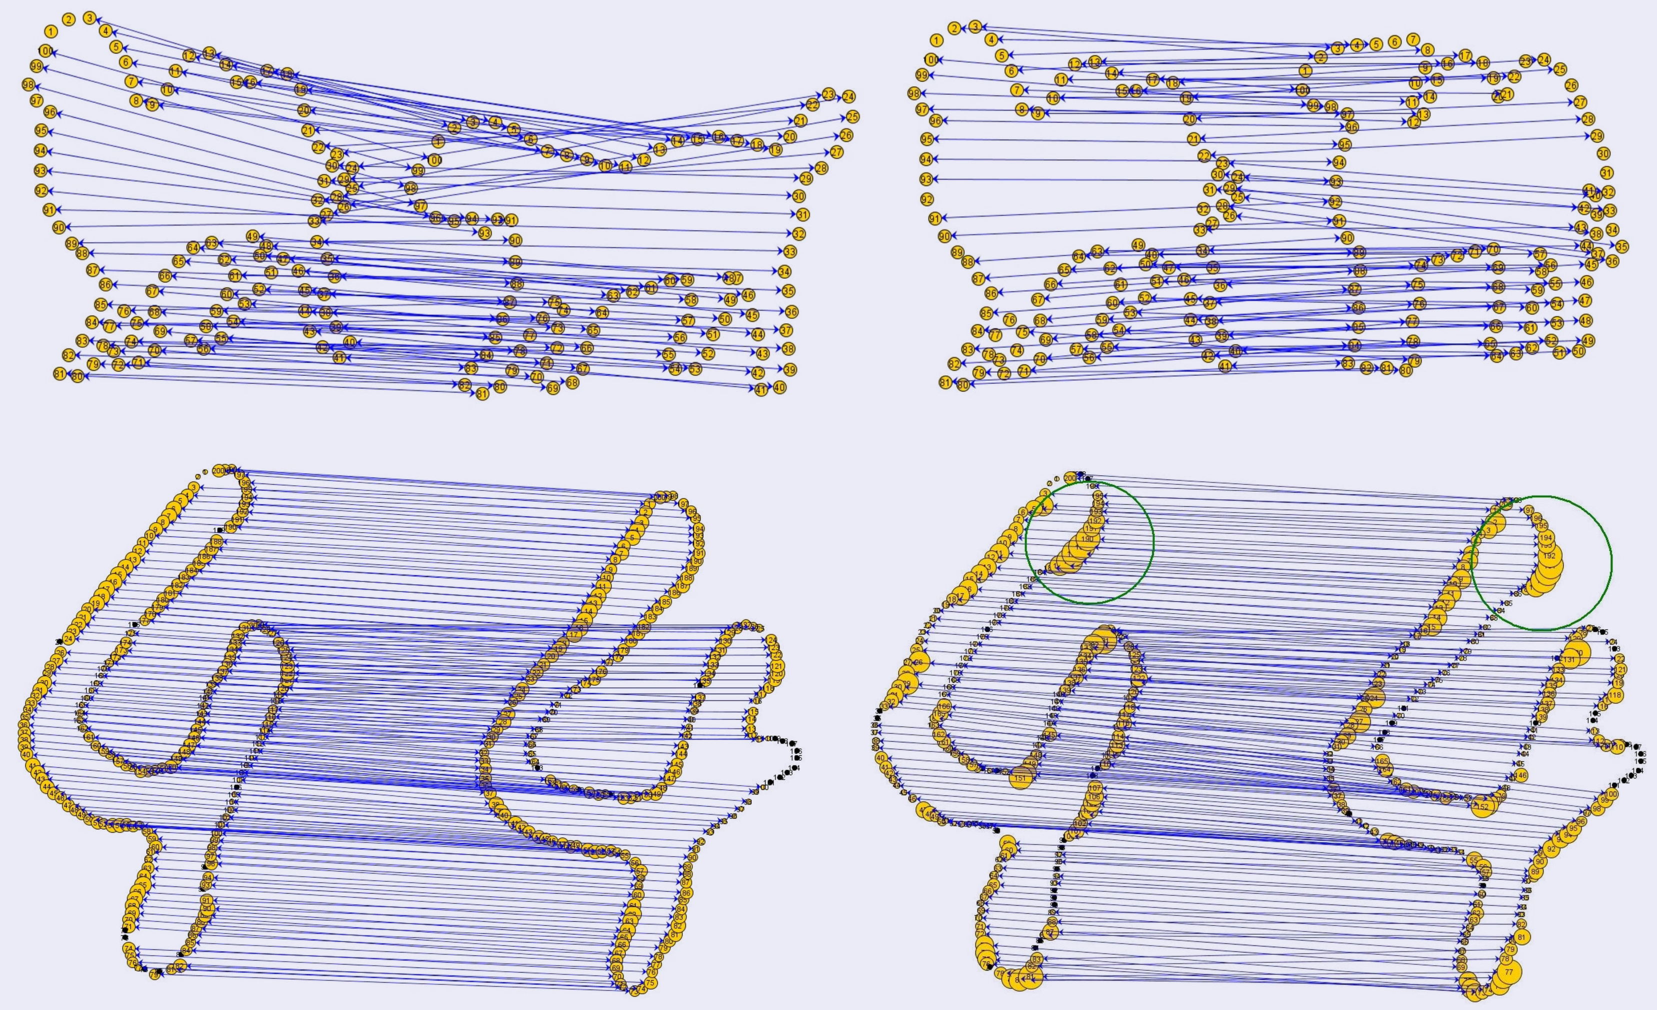

Top: classification without learning (left), and with learning (right). Without learning, the incorrect class is chosen, possibly due to the poor quality of the match. Bottom: an illustration of the weights learned by our method, before and after learning. High weight has been given to those features that are useful in distinguishing '4' from '9' (for example).

Bibliography

[CMF+09] L. Chen, J. J. McAuley, R. S. Feris, T. S. Caetano, and M. Turk. Shape classification through structured learning of matching measures. *CVPR*, 2009.

[LBH98] Yann Lecun, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.

[LLE00] L. Latecki, R. Lakamper, and U. Eckhardt. Shape descriptors for non-rigid shapes with a single closed contour. In *CVPR*, 2000.

[TJHA05] I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *JMLR*, 6:1453–1484, 2005.

lbchen@cs.ucsb.edu, julian.mcauley@nicta.com.au, rsferis@us.ibm.com, tiberio.caetano@nicta.com.au, mturk@cs.ucsb.edu