DD2434 Projects

Hedvig Kjellström

December 5, 2014

Abstract

The task of the project is to reproduce the results presented in a published scientific article, describe the article orally and in written form to your peer students, and argue for and against the method presented in the article. From this you will learn how to read scientific articles, how to implement and use a particular method, how to argue for and against a method, and how to adapt the presentation of a method to different target groups (i.e., adapt the presentation of the method in the article - targeted to active researchers in Machine Learning - so that it is understandable to first year Master students in Machine Learning).

The below 7 papers represent a range of different topics in Machine Learning, and have been selected by Hedvig, who will be the supervisor of these projects.

Some of the papers are more theoretical and while others are of a more practical nature. The requirements will change accordingly, so if you pick a more practical paper you will need to perform more experiments while a more theoretical paper requires you to show a more thorough analysis of the paper.

Detailed instructions about the project can be found on the course home page, Project in the menu to the left.

1 Graphical Models

Y. Boykov and M.-P. Jolly. "Interactive Graph Cuts for Optimal Boundary & Region Segmentation of Objects in ND Images". In: *International Conference on Computer Vision*. 2001

This paper is about image segmentation, and is extremely influential in Computer Vision. It treats an image as a Markov Random Field, optimizing segmentation boundaries using the Graph Cuts algorithm, also developed by Boykov. It is ok to use an existing implementation of the Graph Cuts algorithm, as an exception to the general demand of implementing from "scratch".

L. Breiman. "Random Forests". In: *Machine Learning* 45.1 (2001), pp. 5–32

The original Random Forests paper. Random Forests are used in diversity of applications; a recent application is the Microsoft person tracker. We have not studied them in the course, but AdaBoost, which we will meet in the Assignment 3, can be seen as a Random Forest.

D. M. Blei *et al.* "Latent Dirichlet allocation". In: *The Journal of Machine Learning Research* 3 (2003), pp. 993–1022

The original LDA paper. LDA was originally designed for text document representation and retrieval, and have after that been used for other kinds of data formulated as bags-of-words models; e.g., DNA sequences and video snippets.

2 Particle Filters

R. Van Der Merwe *et al.* "The unscented particle filter". In: *Advances in Neural Information Processing Systems* 12. 2000

The unscented particle filter is an improvement of the regular particle filter (which we meet in Assignment 3) that uses an unscented Kalman filter (UKF) to generate the importance proposal distribution.

3 Approximative Nearest Neighbor

A. Gionis *et al.* "Similarity search in high dimensions via hashing". In: *VLDB Conference*. 1999

Exakt kNN (as we study it in Assignment 3) becomes intractable for very large N. Locality Sensitive Hashing is a principled approximation to exact kNN, which pre-organizes the state space so as to restricting the neighbor search to a small subset of the space.

4 Reinforcement Learning

R. H. Crites and A. G. Barto. "Improving elevator performance using reinforcement learning". In: *Advances in Neural Information Processing Systems* 8. 1996

Reinforcement learning can be seen as a third learning paradigm, in addition to supervised and unsupervised learning. This paper should give you insights in how reinforcement works, without requiring the massive implementation effort that would be necessary with a more recent paper.

5 Kernels

K. Grauman and T. Darrell. "The pyramid match kernel: discriminative classification with sets of image features". In: *IEEE International Conference on Computer Vision*. 2005

Pyramid kernels are designed for very efficient comparison of sets of features. The algorithm is widely used in, e.g., large-scale retrieval applications, where computational efficiency is crucial.

References

- Y. Boykov and M.-P. Jolly. “Interactive Graph Cuts for Optimal Boundary & Region Segmentation of Objects in ND Images”. In: *International Conference on Computer Vision*. 2001.
- L. Breiman. “Random Forests”. In: *Machine Learning* 45.1 (2001), pp. 5–32.
- D. M. Blei *et al.* “Latent Dirichlet allocation”. In: *The Journal of Machine Learning Research* 3 (2003), pp. 993–1022.
- R. Van Der Merwe *et al.* “The unscented particle filter”. In: *Advances in Neural Information Processing Systems* 12. 2000.
- A. Gionis *et al.* “Similarity search in high dimensions via hashing”. In: *VLDB Conference*. 1999.
- R. H. Crites and A. G. Barto. “Improving elevator performance using reinforcement learning”. In: *Advances in Neural Information Processing Systems* 8. 1996.
- K. Grauman and T. Darrell. “The pyramid match kernel: discriminative classification with sets of image features”. In: *IEEE International Conference on Computer Vision*. 2005.