AKMI: Final Report

The project Advanced Kernel Methods for Medical Imaging (AKMI) has developed and analysed machine learning methods for solving relevant medical imaging problems. The AKMI project considered the diagnosis of osteoarthritis (OA) as well as breast cancer risk prediction, two of the world's most burdening diseases. Methodologically, the project focussed on learning features—that is, alternative representations of the input images—for kernel-based learning algorithms and deep neural networks.

Deep learning

Learning classifiers based on raw medical imaging data is challenging. Proper features have to be extracted from the images to achieve good generalization, that is, good performance on new patients. Deep learning methods process inputs in a multi-layered, hierarchical fashion, where each layer corresponds to a different feature representation of the input—a principle inspired by the human visual system. In the ALMI project, different feature learning architectures have been researched and implemented, namely hierarchical kernel methods (Diao, 2012), restricted Boltzmann machines (RBMs), as well as convolutional neural networks (CNNs).

Our insights into RBMs, which have found application in a wide range of pattern recognition tasks and are the building blocks of deep belief networks, were summarized in an article (Fischer & Igel, Pattern Recognition 47(1), 2014). The well-received survey rigorously introduces learning RBMs from the viewpoint of Markov random fields and Markov chain Monte Carlo (MCMC) methods.

Kernel methods

Kernel-based based machine learning algorithms such as support vector machines (SVMs) are well understood theoretically and give excellent results in practice. Here learning proper features amounts to learning the kernel and can be viewed as part of model selection (i.e., adapting parameters of the learning algorithm itself).

We implemented and analyzed model selection strategies for kernel SVMs. Stochastic derivative-free as well as gradient-based methods were considered. Model selection bears the risk of overfitting to the available data, which leads to overoptimistic conclusions. In the AKMI project, we published an article to raise the awareness of the different types of overfitting that can occur including guidelines for avoiding them (Igel, IEEE TEC 17(3), 2013).

We derived a way to proof new generalization bounds for multi-class SVMs, that is, a way to obtain guarantees on how well a classifier will perform on unseen data (Dogan et al., LNCS 7523, 2012). Then we addressed the fundamental question about the computational complexity of determining whether there is a hypothesis class containing a classifier such that the upper bound on the generalization error is below a certain value. Results of this type are important for model comparison and selection. In a first step, we proved that minimizing a basic margin-bound is NP-hard when considering linear classifiers (e.g., SVMs) and the ramp loss (which is used in robust SVMs). This result directly implies the hardness of ramp loss minimization (Frejstrup Maibing & Igel, AISTATS, 2015).

In addition, we worked on hashing for speeding up kernel methods for large-scale data. Hashing algorithms have been successfully applied to scaling up the computation of set and bit-string similarities, and we evaluated different hashing strategies in the context of SVMs and nearest neighbor search (Dahlgaard et al., IEEE Big Data, 2013).

Diagnosis of osteoarthritis

We applied machine learning for segmenting cartilage from knee MRI scans. This is needed for the quantitative analysis of the deterioration of articular cartilage, which causes osteoarthritis. Osteoarthritis is one of main causes of work disability throughout the world, especially for the elderly, and non-invasive assessment of articular cartilage is commonly based on MRI scans. Using more than one classification stage and exploiting class population imbalance allowed us to employ complex classifiers, which would otherwise require long training times. Our approach reached a higher accuracy in comparison to the state-of-the-art method. It also achieved better inter-scan segmentation reproducibility when compared to a radiologist as well as the current state-of-the-art method (Prasoon et al., IEEE EMBC, 2012). In the next step, we developed a novel CNN architecture for efficient 3D image analysis, which further increased the accuracy (Prasoon et al., LNCS 8150, 2013). This result confirms that by employing a deep learning architecture that autonomously learns the features from images we can indeed improve medical image analysis, which is the main claim of the AKMI project.

Breast cancer risk analysis

We have shown that computerized techniques are needed to fully overcome the impact of subjectivity when evaluating mammographic density and parenchymal patterns, two important breast cancer risk factors (Winkel et al., BMC Cancer 15, 2015). Mammographic risk scoring has commonly been automated by extracting a set of handcrafted features from mammograms, and relating the responses to breast cancer risk. We developed a method that learns a feature hierarchy from unlabeled data, which serves as the basis for breast density segmentation and scoring of mammographic texture. The learned breast density scores have a very strong positive relationship with manual ones, and the learned texture scores are predictive of breast cancer (Kallenberg et al., IEEE TMI 35(5), 2016).

Then we have applied our method to an even more challenging task, namely to breast cancer classification of mammographic images that were mis-classified by radiologists in a mammographic screening. The results show that our CNN model performs on par with radiologists in estimating breast density and that it carries an approximately 10% stronger signal for predicting cancer compared to radiologists' texture scores.

Project homepage:

http://www.diku.dk/english/research/imagesection/machine-learning/akmi/

Contact information:

Christian Igel, Professor, Dr. habil.
University of Copenhagen, Department of Computer Science
Universitetsparken 1, 2100 København Ø
Email: igel@diku.dk

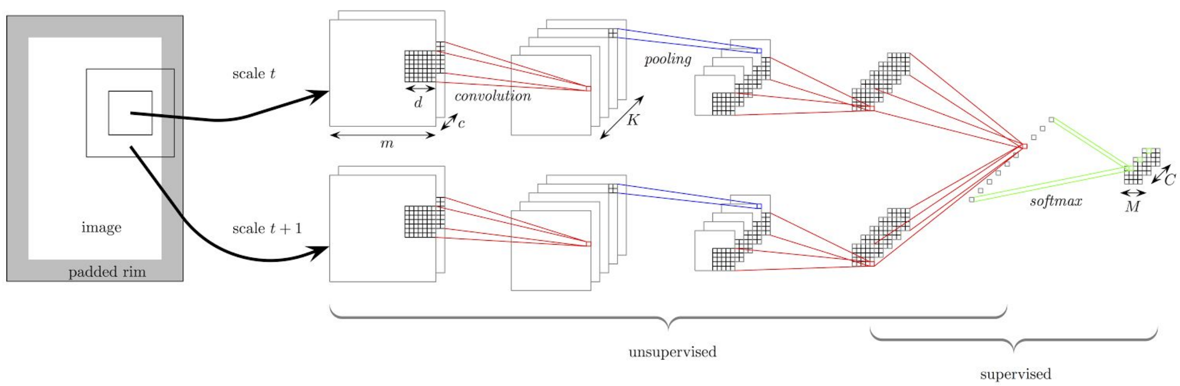

Fig. 1: New convolutional neural network (CNN) architecture for breast density segmentation and scoring of mammographic texture (Kallenberg et al., IEEE TMI 35(5), 2016).

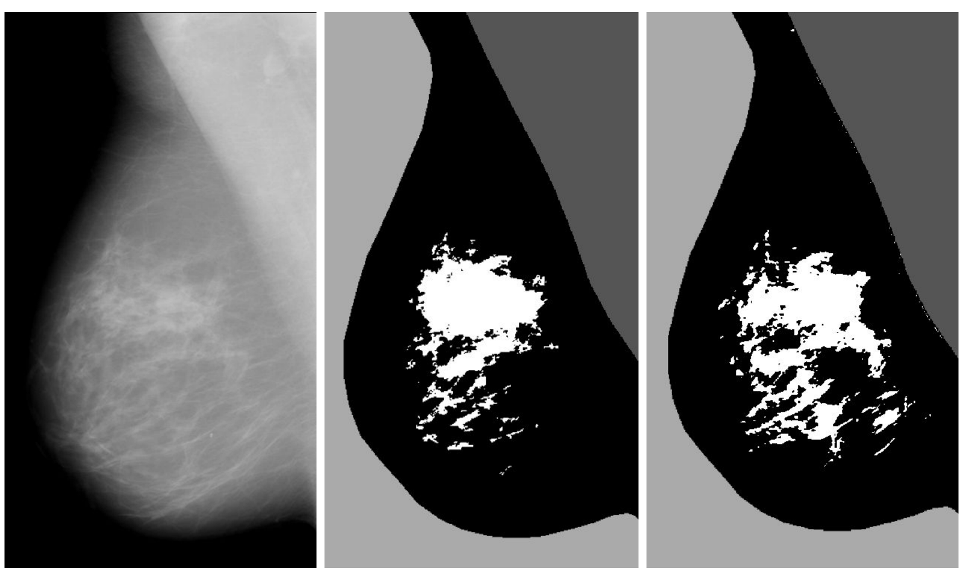

Fig. 2: Example of dense tissue segmentation (Kallenberg et al., IEEE TMI 35(5), 2016), original image (left), expert segmentation (middle), CNN segmentation (right).