Program of the NVPHBV Spring Meeting 2016

20th May 2016, Erasmus MC, Rotterdam

09:00 Doors open, coffee

09:30 **Welcome address**
Veronika Cheplygina, Erasmus MC

09:40 **Session 1**, chaired by *Erik Bekkers*, Eindhoven University of Technology

09:40 **A Stochastic Quasi-Newton Method for Non-rigid Image Registration**
Yuchuan Qiao, Zhuo Sun, Boudewijn P.F. Lelieveldt, Marius Staring, LUMC

10:00 **Hands-Free Interactive Segmentation of Medical Volumes**
Florian Dubost, Erasmus MC & *Lois Peter*, Technical University Munich

10:20 **Anomaly detection in radiological images using deep learning**
Ioannis Katramados, COSMONiO

10:40 Coffee

11:00 **Keynote**: **Surgical data science for the operating room of the future**
Lena Maier-Hein, German Cancer Research Center (DKFZ)

12:00 Lunch

13:00 Members meeting

13:30 **Keynote: Structured prediction with weak feedback from interaction logs**

Thorsten Joachims, Cornell University

14:30 Break

14:50 **Session 2, ,** chaired by *Mitko Veta*, Eindhoven University of Technology

14:50 **Target Contrastive Estimator for Robust Domain Adaptation**
Wouter Kouw & Marco Loog, Delft University of Technology

15:10 **Active Learning Using Uncertainty Information**
Yazhou Yang & Marco Loog, Delft University of Technology

15:30 Coffee

15:50 **Session 3,** chaired by *Carolyn Langen*, Erasmus MC

15:50 **Geometric connectivity analysis in curvilinear images based on the data-driven edge co-occurrences**
S. Abbasi-Sureshjani, J. Zhang, G. Sanguinetti, R. Duits, B. ter Haar Romeny, Eindhoven University of Technology

16:10 **Augmented Reality: from ARToolKit to HoloLens**
Ir. Lex van der Sluijs & Dr. John Schavemaker, TWNKLS

16:30 **Panel discussion** with COSMONiO, Quantib, De Tijdelijke Expert and TWNKLS

17:15 Drinks (until 19:00)

Keynote 11:00 – 12:00

Surgical data science for the operating room of the future

Lena Maier-Hein, German Cancer Research Center (DKFZ)

Despite spectacular advances in the field of medical imaging in the past decades, early cancer diagnosis and precise tumor therapy remain major healthcare challenges with high socioeconomic importance. Surgical data science is an emerging scientific discipline with the objective of improving the safety, quality, effectiveness, and efficiency of surgical care by means of data acquisition, modeling, and analysis. One key goal is to support physicians throughout the entire process of disease diagnosis, therapy and follow-up with the right information at the right time. To achieve this, we propose integration and advancement of methods from the research fields of machine learning, semantic modelling, medical image processing and biophotonics. The talk will highlight some of our recent contributions in this context with a particular focus on multispectral image analysis for cancer detection and crowdsourcing-based large-scale image annotation for context-aware guidance in tumor therapy.

Lena Maier-Hein received the diploma degree in computer science from Karlsruhe Institute of Technolgy (KIT) in 2005 (with distinction) and from Imperial College London in 2003 (International Diploma). She earned the PhD degree from KIT with distinction in 2009 and conducted her postdoctoral research in the Division of Medical and Biological Informatics at the German Cancer Research Center (DKFZ) and at the Hamlyn Centre for Robotics Surgery at Imperial College London. As an independent junior group leader at the DKFZ, she is now working in the field of computer-assisted medical interventions with a focus on multi-modal image processing, knowledge-based systems and computational biophotonics. She is/ has been (co-) principal investigator on a number of national and international grants including the Collaborative Research Centre 125 "Cognition guided surgery" (2012-2024) funded by the German Research Foundation (DFG) and a European Research Council (ERC) starting grant 2015. For her scientific contributions, she earned numerous scientific awards, including the Waltraud-Lewenz Prize 2008, the Ingrid-zu-Solms Prize for Natural Sciences 2009/2010 and the Heinz Maier-Leibnitz Prize 2013, which is known as the most important German award for young researchers of all disciplines.

Keynote 13:30 – 14:30

Structured prediction with weak feedback from interaction logs

Thorsten Joachims, Cornell University

Log data is one of the most ubiquitous forms of data available, as it can be recorded from a variety of systems (e.g., search engine, ad placement engine, image segmenter) at little cost. The interaction logs of such systems typically contain a record of the input to the system (e.g., image), the prediction made by the system (e.g., segmentation of image) and the feedback provided by the user (e.g., user accepted/rejected this segmentation). This feedback, however, provides only partial-information feedback — aka "bandit feedback" — limited to the particular prediction shown by the system. The loss of all other possible predictions is not observed. This is fundamentally different from conventional supervised learning, where "correct" predictions (e.g., correct segmentation) together with a loss function provide full-information feedback.

In this talk, I will explore approaches and methods for batch learning from logged bandit feedback (BLBF). Unlike the well-explored problem of online learning with bandit feedback, batch learning with bandit feedback does not require interactive experimental control of the underlying system, but merely exploits logged interaction data collected in the past. The talk presents a new inductive principle for BLBF, new counterfactual risk estimators, and a new method for structured output prediction with BLBF. Joint work with Adith Swaminathan.

Thorsten Joachims is a Professor in the Department of Computer Science and the Department of Information Science at Cornell University. His research interests center on a synthesis of theory and system building in machine learning, with applications in information access, language technology, and recommendation. His past research focused on support vector machines, text classification, structured output prediction, convex optimization, learning to rank, learning with preferences, and learning from implicit feedback. In 2001, he finished his dissertation advised by Prof. Katharina Morik at the University of Dortmund. From 1994 to 1996 he was a visiting scholar with Prof. Tom Mitchell at Carnegie Mellon University. He is an ACM Fellow, AAAI Fellow, and Humboldt Fellow.

09:40 – 10:00

A Stochastic Quasi-Newton Method for Non-rigid Image Registration

Yuchuan Qiao, Zhuo Sun, Boudewijn P.F. Lelieveldt, Marius Staring
Leiden University Medical Center
Y.Qiao@lumc.nl

Image registration is important in the field of medical image analysis. However, this process is often very slow because of the large number of voxels in the images and the complexity of the registration algorithm. A powerful optimization method is needed to shorten the time consumption during the registration process, which would benefit time-critical intra-operative procedures relying on image guidance.

We propose a stochastic quasi-Newton method (s-LBFGS) specifically for non-rigid image registration, inspired by Byrd et al. [1]. Different from Byrd's method, the proposed method employs only gradients and avoids computing second order derivatives of the cost function to capture the curvature. Secondly, we employ an automatic and adaptive scheme for optimization step size estimation instead of a fixed manual scheme. Finally, we propose a restarting mechanism where the optimal step size is recomputed when a new Hessian approximation becomes available, i.e. every L iterations. The proposed method and some variations are compared to a deterministic quasi-Newton method (LBFGS) and adaptive stochastic gradient descent (ASGD). Validation is performed using 3D lung CT follow-up data and manually annotated corresponding points.

10:00 – 10:20

Hands-Free Interactive Segmentation of Medical Volumes

Florian Dubost [1], Loic Peter [2]
[1] Erasmus MC, [2] Technical University Munich
floriandubost1@gmail.com

We propose a novel hands-free method to interactively segment 3D medical volumes with binary user inputs only. In our scenario, a human user progressively segments an organ by answering a series of questions of the form "Is this voxel inside the object to segment?". At each iteration, the chosen question is defined as the one halving the set of candidate segmentations given the answers already provided. To handle the increased complexity of the task inherent to the large size of 3D data, our question selection strategy builds on a pre-trained boosting classifier encoding available prior knowledge about the organ to be segmented. Experiments on a dataset of 10 MRI volumes for left-atrium segmentation and on a dataset of 15 MRI volumes for prostate segmentation demonstrate the potential of our strategy. Future work directions focus on the reduction of the waiting time between two consecutive questions, and include online updates of the classifier.

10:20 – 10:40

Anomaly detection in radiological images using deep learning

Ioannis Katramados, COSMONiO
io@cosmonio.com

Deep learning is usually linked to Big Data. However, in a medical context radiologists often need help in diagnosing rare conditions with limited historical data. Can Deep Learning prove a useful tool in such cases? COSMONiO is designing NOUS, a deep neural network platform that aims to make high-accuracy predictions using significantly smaller training datasets generated from X-ray, CT, MRI and PET scanners. We will discuss the main challenges and how we address them.

14:50 – 15:10

Target Contrastive Estimator for Robust Domain Adaptation

Wouter Kouw & Marco Loog
Delft University of Technology
W.M.Kouw@tudelft.nl

In domain adaptation, a classifier that is built to adapt to samples from a different distribution (target domain) can perform worse than the original unadapted classifier (source domain). However, one can construct an estimator by assuming a worst-case labeling and contrasting it with the unadapted model that will only produce a different estimate when it is sure that it can improve. Experiments show that this estimator never performs worse than the unadapted classifier in terms of the log-likelihood of the classification model, is competitive with popular other robust estimators and can even outperform non-robust estimators consistently.

15:10 – 15:30

Active Learning Using Uncertainty Information

Yazhou Yang & Marco Loog
Delft University of Technology
yazhouy@gmail.com

Many active learning methods belong to the retraining-based approaches, which select one unlabeled instance, add it to the training set with its possible labels, retrain the classification model, and evaluate the criteria that we base our selection on. However, since the true label of the selected instance is unknown, these methods resort to calculating the average-case or worse-case performance with respect to the unknown label. In this paper, we propose a different method to solve this problem. In particular, our method aims to make use of the uncertainty information to enhance the performance of retraining-based models. We apply our method to two state-of-the-art algorithms and carry out extensive experiments on a wide variety of real-world datasets. The results clearly demonstrate the effectiveness of the proposed method and indicate it can reduce human labeling efforts in many real-life applications.

15:50 – 16:10

Geometric connectivity analysis in curvilinear images based on the data-driven edge co-occurrences

S. (Samaneh) Abbasi-Sureshjani, J. Zhang, G. Sanguinetti, R. Duits, B. ter Haar Romeny
Eindhoven University of Technology
s.abbasi@tue.nl

Natural images contain several curvilinear structures, which might be disconnected, partly occluded or interrupted. Recovering the connectivity information of disconnected structures is an open issue and needs appropriate geometric reasoning. The human brain is capable of grouping the local edge element to global contours perfectly. The contour perception is closely related to the statistics of edge co-occurrences in natural images, and it is well explained by the Gestalt law of good continuation. This contextual information which models the relation of the information at two different positions, provides perfect similarity measurements that can be used in an spectral clustering step for retrieving the connectivity information in natural images. In this work, after an initial segmentation of blood vessels in retinal images and their orientation extraction, the edge co-occurrences have been calculated based on a cross-correlation histogram (called the connectivity kernel) on a training set. Further comparisons between the fundamental solution of the time independent Fokker-Planck operator and this data-driven kernel bridges between the geometric model and the edge statistics. To validate the data-driven kernel, it is used to calculate the affinity matrix for a set of test image patches. For each patch, a self-tuning spectral clustering is applied, which is capable of grouping the data so that the pixels belonging to individual blood vessels (called as perceptual units) are detected immediately. These analysis can be further used not only for junction resolution and tracking algorithms, but also for tuning the computer vision algorithms for perceptual completion problems.

16:10 – 16:30

Augmented Reality: from ARToolKit to HoloLens

Ir. Lex van der Sluijs, Dr. John Schavemaker
TWNKLS | Augmented Reality
lex@twnkls.com

Augmented Reality, or AR in short, is the visual overlay of virtual elements on the real world. It differs from virtual reality in the sense that the view of the real world is combined with a virtual one. Because that needs to be done very quickly and accurately it is a big challenge for the underlying computer vision algorithms that determine the 6 degrees-of-freedom pose of the user's viewpoint with respect to (objects in) the real world. In this presentation we give a short historical overview of augmented reality. We will touch upon the hardware and software challenges: most notably the field-of-view of the AR display and the real-time aspects of the required computer-vision software to enable jitter less tracking. We show how reality can be augmented by placing artificial markers in the world (the well-known black-and-white markers from the original ARToolKit are a good example) or using the real word as is by applying so-called natural features. At the end, we show some recent hardware and software advances (e.g. Microsoft HoloLens, Meta Glasses, 3D SLAM) in the field.