Overview of the 2014 Workshop on Medical Computer Vision — Algorithms for Big Data (MCV 2014)

Henning Müller[1,2], Bjoern Menze[3,4], Georg Langs[5,6], Albert Montillo[7], Michael Kelm[8], Shaoting Zhang[9], Weidong Cai[10,11], and Dimitris Metaxas[12]

[1]University of Applied Sciences Western Switzerland (HES–SO), Switzerland
[2]University Hospitals and University of Geneva, Switzerland
[3]Technical University of Munich, Germany
[4]INRIA, Sophia–Antipolis, France
[5]Medical University of Vienna, Austria
[6]MIT, Cambridge, MA, USA
[7]GE Global Research, USA
[8]Siemens Healthcare, Erlangen, Germany
[9] UNC Charlotte, USA
[10] University of Sydney, Australia
[11] Harvard Medical School, USA
[12] Rutgers University, USA
`henning.mueller@hevs.ch`

Abstract. The 2014 workshop on medical computer vision (MCV): algorithms for big data took place in Cambridge, MA, USA in connection with MICCAI (Medical Image Computing for Computer Assisted Intervention). It is the fourth MICCAI MCV workshop after those held in 2010, 2012 and 2013 with another edition held at CVPR 2012. This workshop aims at exploring the use of modern computer vision technology in tasks such as automatic segmentation and registration, localisation of anatomical features and extraction of meaningful visual features. It emphasises questions of harvesting, organising and learning from large–scale medical imaging data sets and general–purpose automatic understanding of medical images. The workshop is especially interested in modern, scalable and efficient algorithms which generalise well to previously unseen images. The strong participation in the workshop of over 80 persons shows the importance of and interest in Medical Computer Vision. This overview article describes the papers presented in the workshop as either oral presentations or short presentations and posters. It also describes the invited talks and the results of the VISCERAL session in the workshop on the use of big data in medical imaging.

Keywords: medical image analysis, medical computer vision, segmentation, detection

1 Introduction

The Medical Computer Vision workshop (MCV) took place in conjunction with MICCAI (Medical Image Computing for Computer–Assisted Interventions) on September 18, 2014 in Cambridge, MA in the USA. This fifth workshop on medical computer vision was organised in connection with MICCAI after the workshops in 2010 [12], 2012 [10] and 2013 [11] and with CVPR in 2012. The workshop received 29 paper submissions of which five were submitted to the VISCERAL session. All papers were reviewed by at least three external reviewers of the scientific committee of the workshop. Then, all borderline papers were reviewed in addition by at least one member of the workshop organisers. The 13 best papers were presented as oral presentations and authors had the possibility to also present a poster on their techniques for discussions during the lunch break.

With the increasing importance of large datasets (and the addition of big data in the workshop title) it was also decided to again add a session on an evaluation campaign called VISCERAL[1] (VISual Concept ExtRaction challenge in RAdioLogy) in 2014. The VISCERAL project [5] is creating large amounts of manually annotated 3D medical data, and is making them available to the research community in four benchmark challenges. The first two benchmarks were focussed on the automatic detection of organs in the body and include annotations of over 20 organs and 50 landmarks in multiple modalities. The third benchmark is on lesion detection and the fourth on the retrieval of similar cases in very large data sets.

This text also gives an overview of the most important discussions that took place during the medical computer vision workshop and the challenges that were identified in the field. Participants gave very good feedback and all agreed to again organize the workshop during future MICCAI conferences.

2 Papers Presented at the Workshop

The oral presentations were separated into four topic areas, papers on segmentation, feature extraction, multi–atlas techniques and the last session on translational medical computer vision.

2.1 Segmentation of Big Medical Data

Wu et al. [21] addressed the problem of segmentation and registration of infant brains from subjects at different ages. They estimated tissue probability maps separately using only training at the respective age and used the probability maps as a good initialization to guide the level set segmentation.

Then, Wang et al. [20] presented a random forest based approach for infant brain image segmentation that fuses multi–contrast MRI and tissue probability

[1] `http://visceral.eu/`

maps. Next, Harmouche et al. [3] proposed a method to segment the pectoralis muscle in CT. Their approach constructs a likelihood using a multivariate distribution of pairwise registered similar training subjects while the posterior tissue map probability is used to drive a graph cuts segmentation.

2.2 Advanced Feature Extraction

In the first paper in this section, van Tulder and de Bruijne [19] adapted a convolutional classification restricted Boltzmann machine to learn features well suited for discriminative feature learning and apply it for texture–based tissue classification on two lung CT problems. Then, Stühler [17] argued that for large scale longitudinal key point tracking in brain MRI of dementia studies, time–consuming non–rigid registration could be avoided by employing local invariant features that are independent of image scale and orientation. Maraci et al. [9] then showed how they combined techniques from the computer vision and medical imaging communities to increase the degree of automation in ultrasound acquisition. They introduced new symmetric SIFT features and used them to represent the acquired image for classification of fetal image anatomical structures. Next Schlegel et al. [14] addressed the need to learn from data collected across multiple hospitals with heterogeneous medical imaging equipment. Using unsupervised pre–training of convolutional neural networks they inject information from hospitals or image classes for which no annotations are available and they show how this can leads to improved classification accuracy in the classification of lung tissue.

2.3 Multi Atlas and Beyond

Ma et al. [8] kicked off our Multi–Atlas section by presenting a hybrid approach for brain anatomy segmentation that combines multi–atlas and learning based methods. Different from traditional learning–based labelling methods, their atlas–guided multi–channel forest learning method utilized information from both the target image and the aligned atlas for a voxel–wise labelling. Next the task of reducing registration cost for radiation therapy planning was addressed by Rivest–Henault et al. [13]. Their approach finds a proxy that can be used to hop from a given image A to a target image B with minimal distortion and they also defined both a clustering scheme and the transitivity error function. Last, Zikic et al. [22] adapted the Atlas Forest approach for the case when target and test brain images lack correspondences such as the case when there is a tumor in one. By training on only atlases similar to the test, they managed to overcome the inherent overtraining problem as shown in the results they presented on BraTS 2013 (Multimodal Brain Tumor Segmentation Challenge).

2.4 Translational Medical Computer Vision

The translation of concepts from computer vision to applications in the medical imaging domain was well represented in our workshop. Shao et al. [15] let off this section by describing a prostate boundary delineation method that forms an estimate of voxel boundary likelihood using votes cast by a regression forest and then form a discrete segmentation by fitting a deformable model. Next Lugauer et al. [7] proposed a model–guided segmentation approach to segment the lumen in coronary computed tomography angiography. Their method builds a Markov Random Field model with convex priors to ensure tubular solutions, which they optimize through a graph–cut based approach. Finally, a method for identifying local image characteristics capable of predicting the presence of local abnormal ventricular activities in the heart was proposed by Cabrera Lozoya et al. [1]. While determining the optimum intensity and texture–based local image features using a random forest, they developed an approach for integrating uncertainties due to errors in the training set and describe how this improves algorithm performance.

3 Invited Speakers

3.1 Xiang Sean Zhou

The first invited speaker was *Xiang Sean Zhou*, Head of Innovations at Siemens Medical Solutions in Malvern, PA, USA who presented his experiences in Medical Imaging research from a large company perspective. He described main lines of his research approach that has lead him to develop rapid and robust anatomy localization. This accomplishment has in part lead him to recently be awarded Inventor of the Year at Siemens. Emphasizing the principle of "robustness through redundancy", he argued that the three keys to achieve high robustness in medical image analysis are *"redundancy, redundancy, redundancy"*, (which reminded us of the saying that "the three key aspects in buying a house are 'location, location, location'"). The perspective on implementing or exploiting redundancy was also inspired by the space and aeronautics industries that are well known for pioneering work in fault tolerant design. Dr. Zhou described multiple ways in which he has successfully designed redundancy into his solutions. Redundancy through ensemble learning was an approach he adopted early on which increases reliability through aggregating multiple machine learning models, while redundancy through modality leverages multiple modalities when providing a clinical interpretation such as was prevalent in the Health-e-Child project. Redundancy through algorithm fusion entails using all of the best methods for image interpretation including detection, registration and segmentation. These as well as several other redundancies have formed the hallmarks of his research which lead to stimulating discussions about what academia might provide for industrial research. Dr. Zhou described how while many algorithms might achieve success on 80–90% of the cases, to find truly robust solutions that can work on 98-99% of the cases requires quantitative evaluation on large–scale standardized datasets, such as those represented in the large challenges now becoming popular at MICCAI or other events.

3.2 Eric G. Learned–Miller

The second invited speaker was *Eric G. Learned–Miller*, Associate Professor of Computer Science at the University of Massachusetts in Amherst, MA, USA who presented a talk entitled "Experience with Big Data: A Decade of Research in Face Recognition". In it he both presented an overview of computer vision techniques for face recognition and provided a historical perspective. A central theme in his research has been in the reduction of the face recognition problem to its salient components. Towards this aim he has created a widely used public resource for the community in the form of a curated face database, Labeled Faces in the Wild. This has over 13,000 faces collected from the web, each labeled with the name of the person pictured. Every face was identified with the Viola–Jones face detector making detection less of a concern. Additionally he provides subsets of the database in which the faces have been cropped, scaled and aligned to a standard reference frame, leaving only the core recognition task. As a nice complement to the discussion we had with Dr. Zhou, Dr Learned-Miller described how he sees a primary benefit coming from maintaining the database online, so that researchers can at any time benchmark their approaches to the state of the art. The system maintains a ranking of all methods submitted. This has enabled Dr. Learned-Miller to uncover trends in the approaches being applied. For example he has observed that, while at first computer vision researchers submitted the best methods, the recent trend has been that researchers from the machine learning community have attained top scores. Additionally, multiple ranking strategies are employed. In one, methods are allowed to train only from data in the database, while in another training data can come from any additional source, including Facebook which several methods employed successfully. This has enabled further insights including an understanding of the relative value in training on more data versus the development of new methodological models and approaches. This provided an excellent tutorial example to complement the MICCAI debate session "Signal Processing or Machine learning: What's right for MICCAI?".

4 VISCERAL Session

The VISCERAL session started with an overview of the challenges in multi–organ detection and the data that were annotated and made available in the VISCERAL project. An overview of the results was presented without detailing the various techniques of the participants. This included a description of the cloud–based evaluation infrastructure that avoids to physically distribute the data. The session also included five presentations of participants on the techniques employed in the benchmark.

Spanier et al. [16] presented their approach for multi–organ segmentation starting with the simplest organs and then going towards harder organs in the process. The process includes identification of the region of interest for each organ, thresholding, seed point identification and then slice growing. Gass et al. [2] used a multi-atlas approach for the segmentation of multiple organs and also the identification of landmarks. The goal was a data–driven and modality–idenpendent approach for multi–organ segmentation.

Only liver segmentation is done by Li et al. in [6]. The approach uses multiple prior knowledge models and an Adaboost classifier, reaching good results on the liver. Jimenez del Toro et al. [18] presented an approach to multi–organ segmentation that is entirely data driven and does not use any organ–specific optimizations. It uses first a global registration and then successive local registrations. For large organs with much contrast a single local registration is used and then for small organs with less contrast the registration of the larger organs is refined in a successive manner. Segmentation reached best results in several of the organs. Kechichian et al. [4] employed multiple graph cut optimization for multi–organ segmentation. Spatial relationships of organs are modelled and registration was done using SURF key points to reach good segmentation results.

5 Discussions at the Workshop

The large number of over 80 participants at the workshop also led to a large number of very interesting questions during the discussions after the talks and also the lunch and coffee breaks. Many comments after the two invited talks highlighted both the importance of data availability and systematic testing. For commercial applications it is clear that robustness is much more important than only pure performance on very specific data sets as pointed out by Sean Zhou. The case of face recognitions also highlights the importance that standardised and publicly available data sets have as well as standardized performance comparisons on the development of algorithms. This can really show advances over the years and it was highlighted that popular believe on best techniques often does not correspond to the reality of systematic evaluations of it.

The discussions also made clear that theoretical novelty is not necessarily the main point when building real applications in medical imaging as stability is important or *redundancy* as Erik Learned–Miller emphasised. Several of the approaches show that clinical impact and importance gain in importance in the field and that computer vision and machine learning approaches can now well be applied to large and heterogeneous data sets in medical imaging. Registration and segmentation remain very important underlying techniques that can help clinical applications. The session on translational medical imaging also highlights that there are many potential application areas with a potential real impact. At this point we would also like to thank the speakers and the workshop participants for the many discussions and exchange of ideas.

6 Conclusions

The fourth edition of the workshop on medical computer vision at MICCAI was a clear success. High quality papers and posters were presented and many discussions on challenges and techniques in medical imaging emerged at the workshop. The workshop gives a forum for exchange at the crossing of medical imaging, computer vision, machine learning and techniques to manage large data sets of heterogeneous nature. Based on the positive experience we foresee to again hold similar workshops at MICCAI in the coming years to follow up on developments in this quickly changing research area.

7 Acknowledgments

This work was supported by the EU in the FP7 through the VISCERAL (318068) and Khresmoi (257528) projects.

References

1. Cabrera Lozoya, R., Komatsu, Y., Berte, K.B., Relan, J., Cochet, H., Haïssaguerre, M., Jaïs, P., Ayache, N., Sermesant, M.: Confidence–based training for clinical data uncertainty in image–based prediction of cardiac ablation targets. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
2. Gass, T., Goksul, O., Szekely, G.: Multi–atlas segmentation and landmark localization in images with large fields of view. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
3. Harmouche, R., Ross, J., Washko, G., San Jose Estepar, R.: Pectoralis muscle segmentation on CT images based on a bayesian graph cuts with a subject–tailored atlas. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
4. Kechichian, R., Valette, S., Sdika, M., Desvignes, M.: Automatic 3D multiorgan segmentation via clustering and graph cut using spatial relations and hierarchically–registered atlases. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
5. Langs, G., Müller, H., Menze, B.H., Hanbury, A.: Visceral: Towards large data in medical imaging – challenges and directions. In: Medical Content–based Retrieval for Clinical Decision Support. MCBR–CDS 2012 (October 2012)
6. Li, X., Huang, C., Jia, F., Li, Z., Fang, C., Fan, Y., Hu, Q.: Automatic liver segmentation using statistical prior models and free–form deformation. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
7. Lugauer, F., Zheng, Y., Hornegger, J., Kelm, M.: Precise lumen segmentation in coronary computed tomography angiography. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
8. Ma, G., Gao, Y., Wu, G., Wu, L., Shen, D.: Atlas–guided multi–channel forest learning for human brain labeling. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
9. Maraci, M.A., Noble, A., Napolitano, R., Papageorghiou, A.: Object classification in an ultrasound video using LP–SIFT features. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
10. Menze, H., B., Langs, G., Lu, L., Montillo, A., Tu, Z., Criminisi, A.: Medical Computer Vision: Recognition techniques and applications in medical imaging. Proceedings of the MICCAI 2012 Workshop on Medical Computer Vision (MCV 2012), Lecture Notes in Computer Science, vol. 7766. Springer (February 2013)

11. Menze, H., B., Langs, G., Montillo, A., Kelm, M., Ḿuller, H., Tu, Z.: Medical Computer Vision: Large DAta in Medical Imaging. Proceedings of the MICCAI 2013 Workshop on Medical Computer Vision (MCV 2013), Lecture Notes in Computer Science, vol. 8331. Springer (February 2014)
12. Menze, B.H., Langs, G., Tu, Z., Criminisi, A. (eds.): Medical Computer Vision: Recognition techniques and applications in medical imaging — MICCAI–MCV 2010, LNCS, vol. 6533. Springer, Beijing, China (December 2010)
13. Rivest-Henault, D., Ghose, S., Pluim, J., Fripp, J., Dowling, J.: Fast multiatlas selection using composition of transformations for radiation therapy planning. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
14. Schlegl, T., Langs, G.: Unsupervised pre–training across image domains improves lung tissue classification. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
15. Shao, Y., Gao, Y., Yang, X., Shen, D.: CT prostate segmentation by boundary regression. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
16. Spanier, A., Joskowicz, L.: Rule–based ventral cavity multi–organ automatic segmentation in CT scans. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
17. Stühler, E.: Dementia–related features in longitudinal MRI: Tracking keypoints over time. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
18. Jimenez del Toro, O., Ḿuller, H.: Evaluation of a hierarchic anatomical segmentation approach as part of the VISCERAL anatomy benchmarks. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
19. van Tulder, G., de Bruijne, M.: Learning features for tissue classification with the classification restricted boltzmann machine. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
20. Wang, L., Gao, Y., Shi, F., Li, G., Gilmore, J., Lin, W., Shen, D.: LINKS: Learning–based multi–source integration framework for segmentation of infant brain images. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
21. Wu, G., Wang, L., Gilmore, J., Lin, Weili an Shen, D.: Joint segmentation and registration for infant brain images. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)
22. Zikic, D., Glocker, B., Criminisi, A.: Classifier–based multi–atlas label propagation with test–specific atlas weighting for correspondence–free scenarios. In: MICCAI workshop on Medical Computer Vision. Lecture Notes in Computer Science, Springer (2014)