Learning clinically useful information from images: Past, present and future

Daniel Rueckert, Ben Glocker, Bernhard Kainz

*Biomedical Image Analysis (BioMedIA) Group
Department of Computing, Imperial College London, UK*

Abstract

Over the last decade, research in medical imaging has made significant progress in addressing challenging tasks such as image registration and image segmentation. In particular, the use of model-based approaches has been key in numerous, successful advances in methodology. The advantage of model-based approaches is that they allow the incorporation of prior knowledge acting as a regularisation that favours plausible solutions over implausible ones. More recently, medical imaging has moved away from hand-crafted, and often explicitly designed models towards data-driven, implicit models that are constructed using machine learning techniques. This has led to major improvements in all stages of the medical imaging pipeline, from acquisition and reconstruction to analysis and interpretation. As more and more imaging data is becoming available, *e.g.*, from large population studies, this trend is likely to continue and accelerate. At the same time new developments in machine learning, *e.g.*, deep learning, as well as significant improvements in computing power, *e.g.*, parallelisation on graphics hardware, offer new potential for data-driven, semantic and intelligent medical imaging. This article outlines the work of the BioMedIA group in this area and highlights some of the challenges and opportunities for future work.

Keywords: Semantic imaging, Intelligent imaging, Machine learning

1. Introduction

Over the last decades, the field of medical imaging has made tremendous progress in developing new and improved techniques for comprehensive imaging of anatomy and function with an ever increasing spatio-temporal

resolution acquiring large volumes of imaging data about each patient. At the same time, it has become clear that the way medical images will be used in future to support personalised and precision medicine requires a fundamental shift from qualitative, subjective interpretation by a human expert towards quantitative, objective, and automated analysis of the images within the context of all available other, non-imaging information.

This article aims to summarise some of the past, present and future developments at the Biomedical Image Analysis (BioMedIA) group at Imperial College London that contribute to this new era of medical imaging. We will discuss our previous work in the context of traditional model-based image analysis, before presenting our current efforts of applying machine learning to challenging tasks in the context of semantic and intelligent imaging. We will conclude with a discussion on future trends and opportunities that arise for the upcoming years in the area of biomedical image analysis.

2. The past: From images to models

Much of the progress in medical image analysis over the past decades has relied on model-based approaches where the underlying image analysis task, *e.g.*, segmentation of a particular organ or registration between multi-modal images, has been described in terms of a well defined mathematical model. The main challenge here is the definition of a suitable model that appropriately expresses the prior knowledge about the image analysis problem at hand and constrains the solution space in terms of model parameters. The image analysis task itself is solved via a mathematical optimisation aiming to maximise an objective function (or equivalently minimise an energy function) defined over the employed model. Such energy functions usually include one or more data terms that indicate how well the model fits the data and one or more smoothness terms that act as regularisation. Minimising the energy function for the model parameters is usually referred to as model fitting.

A classical example for a model-based image analysis technique is image registration: The BioMedIA group has made many contributions in the field of non-rigid registration, in particular on approaches using free-form deformations (FFDs) with B-splines [Rueckert et al. (1999)] as the mathematical model for non-rigid transformations. The corresponding energy function consists of an image similarity term (*e.g.*, mutual information or cross correlation) as well as a regularisation term that guarantees additional smoothness for the transformation. A key aspect of the success of such energy-based

models is the fact that modern optimisation techniques, and in particular discrete approaches, provide very efficient tools for inference. For example, discrete optimisation methods that had been successfully applied to image segmentation tasks turned out to be very well suited for FFD-based non-rigid registration [Glocker et al. (2008)] and allow the registration of 3D medical images in time frames that are acceptable in clinical workflows.

In the context of medical image analysis one can differentiate between mathematical models (such as the FFD model), biomechanically inspired models and data-driven models. Over the last years data-driven models have attracted significant interest, in particular since they offer the advantage that they can model variability of anatomy (both in terms of shape and appearance) and function across subjects in a data-driven fashion. They therefore represent some of the earliest approaches that use learning in the model construction which can be either statistical (*e.g.*, active shape and appearance models) or probabilistic. In this context, we have made a number of contributions that address in particular the challenges of modelling spatio-temporal data in medical images: We have built several spatio-temporal models that characterise brain development in both neonates and fetal subjects across a large population at different ages [Serag et al. (2012)] or for the cardiovascular system [Bai et al. (2015)].

A drawback of statistical and probabilistic models is that they often make simplifying assumptions. For example, statistical models are often based on linear models such as principal component analysis while probabilistic models assume that the distributions that are used to model population variability can be sufficiently accurately modelled using Gaussian distributions. This leads to limitations in the ability to capture pathologies and to model large variability or heterogeneity in shape or appearance. An alternative, powerful approach is to model the shape and appearance of the anatomy via manifold learning. We have shown that such manifold modelling can characterise the typical trajectories of brain development in neonates [Aljabar et al. (2011)] as well as capture useful information for modelling disease progression [Wolz et al. (2012)]. This also provides good examples of the more recent trends that aim for a close integration of machine learning into the design of image analysis algorithms which will be discussed in more detail in the next section.

3. The present:

Convergence of machine learning and image analysis

Until recently, engineering of task-specific models and energy terms for a particular problem and data at hand has been one of the key aspects in developing novel image analysis methodology. Such approaches, however, are highly sensitive to the original input data and do not generalise well to similar tasks or new data. Machine learning aims to alleviate this problem, providing a data-driven approach to analysis tasks such as image segmentation and registration by automatically learning the relevant models, features and data characteristics for a given problem from large image databases. Loosely speaking, instead of trying to specify explicitly a model by hand in order to solve an image analysis problem, the idea is to implicitly learn the model directly from data, *i.e.*, via labelled or unlabelled examples. For example, a general purpose classification method can be utilised to perform a segmentation task by training the classifier on a set of example images with corresponding reference segmentations. The learning process is an optimisation problem in which the right features are selected from the training data that allow the classifier to learn how to make the correct decisions at test time, when applied to new data. In the following, we discuss two particular areas in which machine learning approaches have been successfully applied and which led to significant progress towards automated analysis: *Semantic imaging* and *Intelligent imaging*. While we discuss these areas separately, we observe that both develop in parallel to machine learning at the moment.

3.1. Semantic imaging

Semantic imaging is a term used to categorise image analysis methods that allow to enrich the raw image data with additional information that aid the understanding of the image content. Semantic information can include very detailed, local information such as voxel-wise organ labels, more global information such as as bounding boxes of organs or image-level tags such as information about the patient (*i.e.*, diagnostic labels) or the image itself (*i.e.*, image modality or image acquisition protocol, image quality).

Image segmentation is a particular good example of semantic imaging: Some of our early work has demonstrated that the performance of atlas-based segmentation techniques can be significantly improved by so-called multi-atlas segmentation approaches [Heckemann et al. (2006)]. The key idea of these approaches is to perform segmentation of the target image

using multiple atlases simultaneously. The output of each atlas-based segmentation is then viewed as the output of a weak classifier and the classifiers are fused into a strong classifier using different decision fusion rules [Heckemann et al. (2006)] leading to significantly increased accuracy and robustness compared to single atlas schemes. We have also shown that atlas selection plays an important role in this context: Choosing similar atlases is likely to lead to better segmentation performance [Heckemann et al. (2006)]. However, in practice this leads to problems as the number of atlases that have been annotated by experts is typically limited. One approach to tackle this problem has been proposed by [Wolz et al. (2010)]. Here a manifold representation of the images is first learnt from a large database of unlabelled images. Then the labelled images are mapped into the manifold and multi-atlas segmentation is performed with nearby, unlabelled images. These newly segmented images are then used as new atlases and can be used to segment currently unlabelled images. This process is iterated until all images are labelled. We have shown that such an approach, coupled with some simple steps to prevent the propagation of errors, can be used to efficiently and robustly segment large populations of diverse anatomies.

An alternative approach is to relax the requirements for similar atlases by measuring similarity in a more local fashion and by also accounting for possible misregistration between the atlas and target image. This has led to the development of numerous patch-based segmentation approaches, including several from our own group. These methods currently provide state-of-the-art performance in many applications including brain [Tong et al. (2013)], heart [Bai et al. (2013)] and abdominal organ segmentation [Wolz et al. (2013); Tong et al. (2015)].

So called correspondence-free approaches remove the need for image registration by learning local classifiers with contextual information, for example, as done in our work based on the popular machine learning method of Random Forests. This general approach can be widely applied to many analysis tasks including efficient multi-atlas label propagation [Zikic et al. (2014)], robust anatomy detection and localisation in arbitrary field-of-view images

Figure 1: Multi-organ, semantic segmentation of whole-body MRI using a Random Forest classification.

[Criminisi et al. (2013)], and effective supervised nearest neighbour retrieval for patch-based segmentation and global image label prediction [Konukoglu et al. (2013)]. A visual example from our work on multi-organ segmentation in whole-body MRI is illustrated in Figure 1.

A particular area where correspondence-free approaches play an important role is fetal image analysis. The arbitrary orientation of the fetus inside the womb and the large possible range of developmental age makes establishing a semantic understanding of the scene very challenging. [Keraudren et al. (2015)] show for example how steerable image features and a Random Forest classification can be used to establish a mapping to a standardised fetal body coordinate system, which defines a semantic neighbourhood relationship for the localisation and segmentation of randomly oriented fetal organs. Figure 2 illustrates this process.

Figure 2: Starting from a segmentation of the fetal brain (purple) we can localize other fetal organs by transforming them into a standard body coordinate system, thus establishing a spatial relationship [Keraudren et al. (2015)], which can be learned together with image features. The panel on the left shows the initial predictions of organ class likelihoods and the panel on right shows the final segmentation. This approach inherently provides semantic understanding of a very complex scene.

The other area where correspondences between training and test images cannot be assumed is in the segmentation of pathologies such as lesions or tumours. Brain lesions, for example, are characterised by very heterogeneous appearance with largely varying locations, frequency, shapes and extents. In this context, machine learning techniques need to be especially effective in generalising from training data. We have recently employed deep learning techniques and in particular fully 3D convolutional neural networks for brain lesion segmentation which have been shown to yield excellent results on very challenging lesion segmentation tasks including multi-modal data from patients with traumatic brain injuries, brain tumours, and stroke [Kamnitsas et al. (2016)]. Deep learning is an emerging technology and considered one of the most promising directions for semantic imaging. Figure 3 sketches our network architecture where intermediate layers correspond to interesting intermediate representations of the brain image data. It is worth highlighting that those intermediate representations are learned directly from the data and the only information provided at training time are multi-modal MR images with corresponding lesion maps. Such intermediate representations could reveal novel imaging biomarkers and provide clinically useful information for better understanding of complex diseases.

Figure 3: Lesion segmentation with deep convolutional neural networks [Kamnitsas et al. (2016)]. Intermediate layers of the network learn automatically high-level features such as spatial location within the brain and separation of graymatter, whitematter, and ventricles. Exploring and visualising those intermediate representations provides valuable insights about black-box machine learning. Additionally, it might reveal novel imaging biomarkers for complex, multi-modal image data.

3.2. Intelligent imaging

Machine learning also plays an increasingly important role in other parts of the medical imaging pipeline, for example in the acquisition and reconstruction of medical images. Many of these techniques aim to make the image acquisition and reconstruction process more intelligent. A good example of this are compressed sensing (CS) approaches that allow the reconstruction of images from undersampled data. These approaches couple clever sampling strategies with reconstruction approaches that exploit sparsity to increase the efficiency of the acquisition and reconstruction stage. While conventional CS approaches typically rely on fixed basis transforms for sparse modelling, which are only able to guarantee suboptimal sparse modelling, one can exploit ideas from dictionary learning to learn dictionaries that are able to optimally adapt sparse modelling to the data being acquired. We demonstrated this idea in the context of cardiac MR images where the use of spatio-temporal dictionaries that are directly learned from the acquired data significantly outperforms fixed basis transforms, thus leading to improved image reconstruction of dynamic MR images of the heart [Caballero et al. (2014b)]. This framework offers significant flexibility and can be extended to enable so-called application driven MR imaging [Caballero et al. (2014a)]. In this approach we demonstrate that jointly reconstructing and segmenting cardiac cine MR images from undersampled data is more efficient compared to the traditional approach in which the images are first reconstructed and subsequently segmented.

Another example of the intelligent imaging paradigm is fetal MRI. Traditionally, fetal MRI is often limited to the acquisition of 2D images of the 3D fetal anatomy to avoid motion artefacts. We have pioneered a number of approaches that use image registration to compensate for fetal motion during the acquisition [Kainz et al. (2015b)]. These approaches not only compensate for motion but exploit multiple motion-corrupted acquisitions in order to reconstruct high-resolution, super-resolved 3D volumes of fetal organs such as the brain, lungs and the cardiac vasculature. Until recently, these methods could not be applied outside the fetal brain because of the assumption of rigid motion in the 2D to 3D registration step of slice-to-volume registration (SVR) methods. To solve this problem we have developed a patch-to-volume registration (PVR) approach, which employs a flexible subdivision of the input space into overlapping and partly rigid image patches [Kainz et al. (2015a)]. This way the motion compensation problem is solved for each patch independently without requiring an a-priori defined region of inter-

Figure 4: Multi-planar reconstruction of a multislice single-shot fast spin echo (ssFSE) MRI stack in a normal fetus. The reconstructed planes orthogonal to the acquisition plane (axial view) show the extent of fetal movement between slices. Following the application of PVR to multiple overlapping stacks of 2D images, a fully navigable 3D volume of the fetal thorax shows clinically relevant information, providing detailed views of the extracardiac vasculature.

est. Acceleration using graphics processing units (GPUs) is used to enable reconstructions within clinically acceptable time frames. An example of a reconstruction of the fetal thorax and extracardiac vasculature is shown in Figure 4.

4. The Future: Challenges and opportunities

At the moment we can observe three major developments in medical image analysis: The first one is the advent of large-scale, publicly accessible databases including both population studies and clinical cases. A good example for large amounts of potential training data is the UK Biobank¹ initiative which is collecting comprehensive imaging information at population level about the multiple organ systems with MRI, together with non-imaging information including lifestyle, demographic and genomic data. The second one is the increasing availability of high-performance computing including cloud and GPU computing that enable the use of much more complex image analysis and machine learning approaches. Finally, machine learning techniques, in particular deep learning, have made significant progress over the last years leading with increased ability to learning of powerful feature representations that are useful in medical imaging. This dramatically improved

¹<http://www.ukbiobank.ac.uk/>, last accessed: 03 June 2016

the state-of-the-art as evidenced by winning several of the computational challenges organised by the medical image analysis community.

We believe that these three developments are likely to accelerate the trend towards integrated, data-driven and learning-based medical image analysis. However, this trend poses also several important research challenges for the scientific community. One of these challenges is the lack of annotated or labelled data that are required to train supervised machine learning techniques. To avoid this bottleneck *weakly* annotated data, *i.e.*, data beyond the traditional pixel-wise manual expert labelling, is likely to be increasingly exploited. A potential approach to achieve this may lie in the use of crowd-sourcing for annotation tasks that do not require extensive expert knowledge. Learning based approaches may also make increasing use of abstract data linked to images such as *omics* data, text and electronic health records. Another challenge is the interpretability of machine learning systems, especially those that are used for computer-aided diagnosis. Future approaches will have to move towards solutions that are interpretable by clinical experts in order to be accepted in general clinical practice. Visualisation and interpretation of the learned intermediate representations as shown in our work on lesion segmentation [Kamnitsas et al. (2016)] is one direction worth further exploration.

Another important trend that is likely to continue and even accelerate is the closer integration of joint acquisition, reconstruction and analysis. The current sequential process of acquisition, reconstruction, processing and interpretation has several disadvantages: First, it is prone to error propagation which is unavoidable in serial data manipulation. A further drawback is the limited efficiency of such a sequential framework. Finally, the sequential paradigm offers no possibility that can steer the image acquisition process based on already acquired information. The integration of image acquisition, reconstruction and analysis with close coupling and feedback between the different stages of the entire imaging pipeline enables what we call intelligent imaging. As mentioned above, we have already shown that different stages of the imaging pipeline can be extended to enable application driven MR imaging [Caballero et al. (2014a)]. Extending such approaches may offer significantly more efficient ways of extracting clinically useful information: For example, one may directly attempt to recover clinically useful information, *e.g.*, the ejection fraction of the left ventricle of the heart from under-sampled MR acquisitions bypassing reconstruction and segmentation all together. It is also likely that such integrated approaches are more robust towards errors

and allow the quantification of uncertainties in the analysis. Finally, the combination of such approaches with developments in parallel computing on GPUs offers the potential for steerable intelligent imaging where real-time feedback from the imaging pipeline steers the image acquisition to optimise the extraction of clinically useful information.

One of the biggest key challenges, however, for our field remains the slow uptake of cutting edge methodology by both clinicians and industry. There are several reasons for this: A major problem is the lack of robustness of many algorithms to cope with pathologies or large and unexpected anatomical variability. Similarly, algorithms often struggle when images have been acquired with different acquisition parameters or different protocols. Furthermore, computational demands of algorithms make it hard to integrate compute intensive methods into standard clinical workflows where high-performance infrastructure might not be available. However, there are also emerging developments that are likely to help the field to overcome these challenges. For example, cloud computing and developments in specialised processing hardware can help address the problem of bringing compute intensive algorithms into clinical environments.

Another reason for slow uptake is that the field might need to focus more on providing *clinically useful solutions*. Much of our research is still focused on the development of components for image analysis (*e.g.*, registration or segmentation). However, clinicians are primarily focused on the clinical utility of the solutions we produce. Thus, we need to move towards the development of integrated solutions that solve clinical problems. In addition to technological challenges, there are a number of real-world problems we are facing as a research community: Our healthcare systems face enormous challenges in terms of resource and cost efficiency; low-cost and affordable medical imaging solutions are required for developing countries; objective and quantifiable information is required to support personalised and precision medicine. These challenges provide an enormous wealth of opportunities for researchers in the field of medical image analysis over the next decades.

Acknowledgements

The authors would like to thank all current and previous members of the BioMedIA group who are the main contributors to the works discussed in this article. We would also like to thank our many clinical and non-clinical collaborators and funders for an exceptional strong support over many years.

References

- Aljabar, P., Wolz, R., Srinivasan, L., Counsell, S., Rutherford, M., Edwards, A., Hajnal, J., Rueckert, D., 2011. A combined manifold learning analysis of shape and appearance to characterize neonatal brain development. *IEEE Trans. Med. Imaging* 30, 2072–2086.
- Bai, W., Shi, W., Ledig, C., Rueckert, D., 2015. Multi-atlas segmentation with augmented features for cardiac MR images. *Medical Image Analysis* 19, 98–109.
- Bai, W., Shi, W., O’Regan, D., Tong, T., Wang, H., Jamil-Copley, S., Peters, N., Rueckert, D., 2013. A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: Application to cardiac MR images. *IEEE Trans Med Imaging* 32, 1302–1315.
- Caballero, J., Bai, W., Price, A., Rueckert, D., Hajnal, J., 2014a. Application-driven MRI: joint reconstruction and segmentation from undersampled MRI data. *Med Image Comput Comput Assist Interv* 17, 106–113.
- Caballero, J., Price, A., Rueckert, D., Hajnal, J., 2014b. Dictionary learning and time sparsity for dynamic MR data reconstruction. *IEEE Trans Med Imaging* 33, 979–994.
- Criminisi, A., Robertson, D., Pauly, O., Glocker, B., Konukoglu, E., Shotton, J., Mateus, D., Möller, A. M., Nekolla, S., Navab, N., 2013. Anatomy detection and localization in 3D medical images. In: *Decision Forests for Computer Vision and Medical Image Analysis*. pp. 193–209.
- Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N., 2008. Dense image registration through MRFs and efficient linear programming. *Medical Image Analysis* 12 (6), 731–741.
- Heckemann, R., Hajnal, J., Aljabar, P., Rueckert, D., Hammers, A., 2006. Automatic anatomical brain MRI segmentation combining label propagation and decision fusion. *NeuroImage* 33, 115–126.
- Kainz, B., Alansary, A., Malamateniou, C., Keraudren, K., Rutherford, M., Hajnal, J., Rueckert, D., 2015a. Flexible reconstruction and correction of

- unpredictable motion from stacks of 2D images. In: *Med Image Comput Comput Assist Interv.* pp. 555–562.
- Kainz, B., Steinberger, M., Wein, W., Kuklisova-Murgasova, M., Malamateniou, C., Keraudren, K., Torsney-Weir, T., Rutherford, M. A., Aljabar, P., Hajnal, J. V., Rueckert, D., 2015b. Fast volume reconstruction from motion corrupted stacks of 2D slices. *IEEE Trans. Med. Imaging* 34 (9), 1901–1913.
- Kamnitsas, K., Ledig, C., Newcombe, V. F., Simpson, J. P., Kane, A. D., Menon, D. K., Rueckert, D., Glocker, B., 2016. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *arXiv preprint arXiv:1603.05959*.
- Keraudren, K., Kainz, B., Oktay, O., Kyriakopoulou, V., Rutherford, M., Hajnal, J., Rueckert, D., 2015. Automated localization of fetal organs in MRI using random forests with steerable features. In: *Med Image Comput Comput Assist Interv.* pp. 620–627.
- Konukoglu, E., Glocker, B., Zikic, D., Criminisi, A., 2013. Neighbourhood approximation using randomized forests. *Medical Image Analysis* 17 (7), 790–804.
- Rueckert, D., Sonoda, L., Hayes, I., Hill, D., Leach, M., Hawkes, D., 08 1999. Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Trans Med Imaging* 18 (8), 712–21.
- Serag, A., Aljabar, P., Ball, G., Counsell, S., Boardman, J., Rutherford, M., Edwards, A., Hajnal, J., Rueckert, D., Feb 2012. Construction of a consistent high-definition spatio-temporal atlas of the developing brain using adaptive kernel regression. *NeuroImage* 59 (3), 2255?–2265.
- Tong, T., Wolz, R., Coupe, P., Hajnal, J., Rueckert, D., 2013. Segmentation of MR images via discriminative dictionary learning and sparse coding: Application to hippocampus labeling. *NeuroImage* 76, 11–23.
- Tong, T., Wolz, R., Wang, Z., Gao, Q., Misawa, K., Fujiwara, M., Mori, K., Hajnal, J. V., Rueckert, D., 2015. Discriminative dictionary learning for abdominal multi-organ segmentation. *Medical Image Analysis* 23 (1), 92–104.

- Wolz, R., Aljabar, P., Hajnal, J., Hammers, A., Rueckert, D., 2010. LEAP: Learning embeddings for atlas propagation. *NeuroImage* 49, 1316–1325.
- Wolz, R., Aljabar, P., Hajnal, J., Lotjonen, J., Rueckert, D., 2012. Non-linear dimensionality reduction combining MR imaging with non-imaging information. *Medical Image Analysis* 16 (4), 819–839.
- Wolz, R., Chu, C., Misawa, K., Fujiwara, M., Mori, K., Rueckert, D., 2013. Automated abdominal multi-organ segmentation with subject-specific atlas generation. *IEEE Trans Med Imaging* 32, 1723–1730.
- Zikic, D., Glocker, B., Criminisi, A., 2014. Encoding atlases by randomized classification forests for efficient multi-atlas label propagation. *Medical Image Analysis* 18 (8), 1262–1273.