Editorial

Machine learning approaches in medical image analysis: From detection to diagnosis

Marleen de Bruijne

[a] Biomedical Imaging Group Rotterdam, Departments of Medical Informatics, Radiology & Nuclear Medicine, Erasmus MC-University Medical Center Rotterdam, The Netherlands
[b] The Image Section, Department of Computer Science, University of Copenhagen, Denmark

A R T I C L E   I N F O

A B S T R A C T

Machine learning approaches are increasingly successful in image-based diagnosis, disease prognosis, and risk assessment. This paper highlights new research directions and discusses three main challenges related to machine learning in medical imaging: coping with variation in imaging protocols, learning from weak labels, and interpretation and evaluation of results.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Supervised learning techniques, which learn a mapping from input data to output (labels) from a set of training examples, have shown great promise in medical image analysis. Pattern classification has already been used for decades to detect, and later characterize, abnormalities such as masses in mammograms and nodules in chest radiographs based on features describing local image appearance (Giger et al., 2008). With improvements in computer hardware it has become feasible to train more and more complex models on more data, and in the last few years, the use of supervised learning in image segmentation, recognition, and registration has accelerated. Trained appearance models are replacing simple intensity and gradient models as a component in segmentation systems, and statistical shape models that describe the typical shape and shape variations in a set of training shapes have replaced free form deformable models in many cases. Several new methods learn to diagnose disease in a fully data driven manner, using multivariate classification or regression to directly map from imaging data to diagnosis. These techniques are not restricted by current knowledge on disease-related radiological patterns and often have higher diagnostic accuracy than more traditional quantitative analysis based on simple volume or density measures.

Supervised quantification approaches can not only assist in diagnosis, but are also increasingly used to predict future disease onset or progression. Models are then trained on data from longitudinal studies in which the disease status years after the acquisition of the baseline image is known. For example at Erasmus MC, Achterberg et al. (2014) showed that hippocampal shape classification in a healthy elderly population is predictive of onset of dementia symptoms up to ten years later. van Engelen et al. (2014) used multivariate sparse Cox regression to take time to event into account in the model and found that changes in plaque texture and volume in ultrasound images of the carotid artery could predict future vascular events better than traditional risk factors could.

Possibly the most widespread application of machine learning based diagnosis appearing in publications is in neurodegenerative diseases, where researchers aim to diagnose Alzheimer's disease or other forms of dementia, or predict conversion from mild cognitive impairment (MCI) to dementia, based on brain MR images. This is likely driven, at least in part, by the availability of large datasets with diagnostic labels, such as the Alzheimer's Disease Neuroimaging Initiative (ADNI) and Open Access Series of Imaging Studies (OASIS).

Another example where availability of data has altered the course of research is the detection of diabetic retinopathy in retinal fundus photographs. Many early papers focused on optimizing detection and segmentation of retinal vessels, for which several smaller public databases with ground truth were available. A recent Kaggle competition on diabetic retinopathy detection[1] changed the field by providing 35,000 images with expert visual

---

E-mail address: marleen@di.ku.dk, marleen.de.bruijne@gmail.com

[1] https://www.kaggle.com/c/diabetic-retinopathy-detection.

scores for training. This has drawn attention from data scientists around the world with no or little prior experience in medical image analysis. Many of the 661 participating teams used no specific pre-processing or segmentation but still obtained very good results. The top performing contributions all used different layouts of convolutional networks, with extensive data augmentation to increase the amount of training data even further, and achieved performance scores surpassing those previously reported for human experts.

We need to keep in mind that this example is a specific task, performed on 2D images. Differential diagnosis or quantification based on full 3D or 4D, possibly multi-modal, imaging data would require even larger training sets to describe all biological variation adequately. Additional domain specific knowledge will therefore still be needed in many cases. Results of another big data challenge which have just become available at the time of writing this paper, the "2015 Data Science Bowl"[2], seem to point in that direction. The challenge was to automatically measure end-systolic and end-diastolic volumes from dynamic cardiac MRIs. While the list of best performing algorithms is again dominated by convolutional neural networks, all top teams also performed specific preprocessing steps to detect the relevant regions and align image sequences. Nonetheless, these two examples suggest that supplying general purpose machine learning algorithms with a large amount of training data can lead to large improvements over current state-of-the-art performance in medical image analysis and computer aided diagnosis.

An enormous amount of data that could potentially be used for training exists: clinical experts assess many thousands of MRI and CT scans every day. In OECD countries alone, over 200 million CT and MRI scans are acquired per year[3] and radiographs and ultrasound images are acquired even much more frequently. Making part of this data available to train computer aided diagnosis algorithms could have tremendous impact.

In this paper I discuss three of the main challenges in approaching diagnosis with machine learning techniques and highlight several interesting research directions.

2. Varying imaging protocols

The main obstacle currently preventing wider use of machine learning in medical imaging is a lack of representative training data. While supervised learning techniques have shown much promise in relatively constrained experiments with standardized imaging protocols, their performance may quickly deteriorate on new images that are acquired under slightly different conditions. These techniques operate under the assumption that both train and test datasets are random samples drawn from the same distribution. In practice however, the available training data is often acquired earlier with a different imaging protocol, different scanner model, or from a different patient population, which would violate this assumption. An example of typical differences that can be found in multi-center MRI studies is given in Fig. 1.

One approach to cope with these issues, which is gaining increasing interest, is to apply transfer learning or domain adaptation techniques. We discern two classes of approaches that both aim to make train and test distributions more similar: weighting and feature space transformation techniques.

In weighting based transfer learning, training data with slightly different properties from the target data to analyze is used next to some labeled target data. A transfer classifier or regressor is then trained on all samples, but the additional, different-distribution samples receive a lower weight than the labeled target data. These different-distribution samples can help to regularize a classifier in a data driven manner — better than an uninformed regularizer — which makes it possible to train a reliable model with fewer labeled target samples. A similar effect can be achieved using the parameters of a classifier trained on different data to regularize a classifier on the target samples, as is done for instance in adaptive SVM. Such approaches may be easier to share between institutes as they do not require access to the original data samples that produced the classifier. Alternatively, samples, images, or image sets can be weighted in a fully unsupervised manner e.g. based on feature distribution similarity (van Opbroek et al., 2015b) or sample similarity (Heimann et al., 2014) with the target data.

In our research, we found that weighting based transfer learning approaches can significantly improve classification accuracy in MRI segmentation problems when few labeled target samples are available (van Opbroek et al., 2015a; van Engelen et al., 2015). However, the number of labeled target samples at which a classifier trained on only those samples performs as good as the transfer learning approaches was in these experiments quite low — several hundred labeled voxels distributed over all classes, up to a few well chosen, fully annotated images (Fig. 2). This depends of course on the data distribution and the model complexity. We would expect that with more complex representations, such as an increased number of image features or the representations obtained using 3D deep neural networks, the benefit of transfer learning becomes more clear. For example, in a different application using marginal space learning to localize ultrasound transducers in fluoroscopy sequences, Heimann et al. (2014) could completely eliminate localization errors by augmenting training sequences with synthetic data and subsequently downweighting less realistic synthetic images using a domain adaptation technique. Moreover, there is clearly still room for improvement in current methods; many general purpose transfer learning techniques are available but few explicitly take (medical) image properties into account.

While approaches based on sample or image weighting can compensate for some changes in distribution, they assume that the conditional distribution of the labels given the feature vectors is similar between the target data and (at least part of) the training data. This will often not be the case, for instance if intensity scale or contrast varies between images and the derived image features are not invariant to such transformations. A first step to address this will typically be image contrast normalization or standardization of image features to zero mean and unit variance, if necessary followed by a correction for intensity inhomogeneities. To compensate for further differences in distributions, a range of supervised and unsupervised techniques have been proposed in the machine learning and computer vision literature to project data into a latent space where distributions are more similar, for instance by minimizing the so-called Maximum Mean Discrepancy between distributions in a kernel space. An important remaining issue is that although transfer learning often improves results on similar tasks, without sufficient labeled target data it is not possible to detect negative transfer which undermines performance. Transfer learning techniques that could guarantee that the result of the transfer technique is never worse than the supervised solution, such as recently proposed for semi-supervised learning (Loog, 2016), are therefore of great interest.

The approaches discussed so far use training data from different sources more wisely and can compensate for possible differences between distributions. An alternative strategy would be to collect a very large and heterogeneous database for each task that contains all possible variations in imaging protocols, similar to the approach taken in the diabetic retinopathy competition described earlier. Combined with a sufficiently rich feature representation

---