

Exploring the similarity of medical imaging classification problems

Veronika Cheplygina^{1,2}, Pim Moeskops¹, Mitko Veta¹, Behdad Dashtbozorg¹,
Josien P.W. Pluim^{1,3}

¹ Medical Image Analysis, Department of Biomedical Engineering, Eindhoven University of Technology, The Netherlands

² Biomedical Imaging Group Rotterdam, Departments of Radiology and Medical Informatics, Erasmus Medical Center, The Netherlands

³ Image Sciences Institute, University Medical Center Utrecht, The Netherlands

Abstract. Supervised learning is ubiquitous in medical image analysis. In this paper we consider the problem of meta-learning – predicting which methods will perform well in an unseen classification problem, given previous experience with other classification problems. We investigate the first step of such an approach: how to quantify the similarity of different classification problems. We characterize datasets sampled from six classification problems by performance ranks of simple classifiers, and define the similarity by the inverse of Euclidean distance in this meta-feature space. We visualize the similarities in a 2D space, where meaningful clusters start to emerge, and show that the proposed representation can be used to classify datasets according to their origin with 89.3% accuracy. These findings, together with the observations of recent trends in machine learning, suggest that meta-learning could be a valuable tool for the medical imaging community.

1 Introduction

Imagine that you are a researcher in medical image analysis, and you have experience with machine learning methods in applications A, B, and C. Now imagine that a colleague, who just started working on application D, asks your advice on what type of methods to use, since trying all of them is too time-consuming. You will probably ask questions like “How much data do you have?” to figure out what advice to give. In other words, your perception of the similarity of problem D with problems A, B and C, is going to influence what kind of “rules of thumb” you will tell your colleague to use.

In machine learning, this type of process is called *meta-learning*, or “learning to learn”. In this meta-learning problem, the samples are the different datasets A, B and C, and the labels are the best-performing methods on each dataset. Given this data, we want to know what the best-performing method for D will be. The first step is to characterize the datasets in a meta-feature space. The meta-features can be defined by properties of the datasets, such as sample size, or by performances of simple classifiers [1–3]. Once the meta-feature space is

defined, dataset similarity can be defined to be inversely proportional to the Euclidean distances within this space, and D can be labeled, for example, by a nearest neighbor classifier.

Despite the potential usefulness of this approach, the popularity of meta-learning has decreased since its peak around 15 years ago. To the best of our knowledge, meta-learning is not widely known in the medical imaging community, although methods for predicting the quality of registration [4] or segmentation [5] can be considered to meta-learn within a single application. In part, meta-learning seems less relevant today because of the superior computational resources. However, with the advent of deep learning and the number of choices to be made in terms of architecture and other parameters, we believe that meta-learning is worth revisiting in the context of its use in applications in medical image analysis.

In this paper we take the first steps towards a meta-learning approach for classification problems in medical image analysis. More specifically, we investigate the construction of a meta-feature space, where datasets known to be similar (i.e. sampled from the same classification problem), form clusters. We represent 120 datasets sampled from six different classification problems by performances of six simple classifiers and propose several methods to embed the datasets into a two-dimensional space. Furthermore, we evaluate whether a classifier is able to predict which classification problem a dataset is sampled from, based on only a few normalized classifier performances. Our results show that even in this simple meta-feature space, clusters are beginning to emerge and 89.3% of the datasets can be classified correctly. We conclude with a discussion of the limitations of our approach, the steps needed for future research, and the potential value for the medical imaging community.

2 Methods

In what follows, we make the distinction between a “classification problem” - a particular database associated with extracted features, and a “dataset” - a subsampled version of this original classification problem.

We assume we are given datasets $\{(D_i, M_i)\}_1^n$, where D_i is a dataset from some supervised classification problem, and M_i is a meta-label that reflects some knowledge about D_i . For example, M_i could be the best-performing (but time-consuming) machine learning method for D_i . For this initial investigation, M_i is defined as the original classification problem D_i is sampled from.

We represent each dataset D_i by performances of k simple classifiers. Each of the n D_i s is therefore represented by an k -dimensional vector \mathbf{x}_i which together form a $n \times k$ meta-dataset A_k . Due to different inherent class overlap, the values in A_k might not be meaningful to compare to each other. To address this problem, we propose to transform the values of each \mathbf{x}_i by:

- Normalizing the values to zero mean and unit variance, creating a meta-dataset N_k

- Ranking the values between 1 and k , creating a meta-dataset R_k . In cases of ties, we use average ranks.

The final step is to embed the meta-datasets A_k , N_k and R_k in a 2D space for visualization, to obtain A_2 , N_2 and R_2 . We use two types of embedding: multi-dimensional scaling (MDS) [6] and t-stochastic nearest neighbor embedding (t-SNE) [7]. These embeddings can help to understand complementary properties of the data. MDS emphasizes large distances, and is therefore good at pointing out outliers, whereas t-SNE emphasizes small distances, potentially creating more meaningful visualizations. An overview of the approach is shown in Fig. 1.

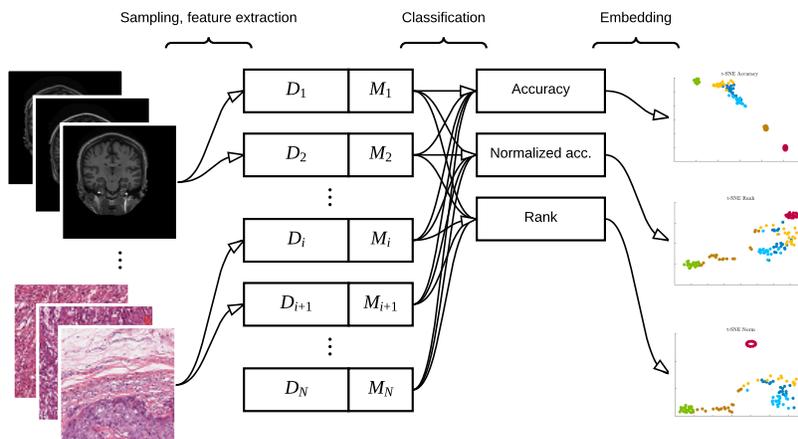


Fig. 1. Overview of the method.

3 Experiments

Table 1. Classification problems described by type of image, number of images (subjects), number of classes and number and type of features: “Classical” = intensity and texture-based, “CNN” = defined by output of the fully-connected layer of a convolutional neural network.

Dataset	Type	Images	Classes	Features
Tissue [8]	Brain MR	20	7	768, CNN [9]
Mitosis, MitosisNorm [10]	Histopathology	12	2	200, CNN [11]
Vessel [12]	Retinal	20	2	29, Classical [13]
ArteryVein [14]	Retinal	20	2	30, Classical [14]
Microaneurysm [15]	Retinal	381	2	30, Classical

Data and Setup. We sample datasets from six classification problems, described in Table 1. The problems are segmentation or detection problems, the number of samples (pixels or voxels) is therefore much higher than the number of images. We sample each classification problem 20 times by selecting 70% of the subjects for training, and 30% of the subjects for testing, to generate $n = 120$ datasets for the embedding. For each of dataset, we do the following:

- Subsample $\{100, 300, 1K, 3K, 10K\}$ pixels/voxels from the training subjects
- Train $k = 6$ classifiers on each training set: nearest mean, linear discriminant, quadratic discriminant, logistic regression, 1-nearest neighbor and decision tree
- Subsample 10K pixels/voxels from the test subjects
- Evaluate accuracy on the test set, to obtain 5×6 accuracies
- Transform the accuracies by ranking or normalizing
- Average the ranks/normalized accuracies over the 5 training sizes

We use balanced sampling for both training and test sets, in order to keep performances across datasets comparable (and to remove the “easy” differences between the datasets). The classifiers are chosen mainly due to their speed and diversity (3 linear and 3 non-linear classifiers). The training sizes are varied to get a better estimation of each classifier’s performance.

Embedding. We use the MDS algorithm with default parameters⁴, and t-SNE⁵ with perplexity = 5. Because of the stochastic nature of t-SNE, we run the algorithm 10 times, and select the embedding that returns the lowest error. We apply each embedding method to A_k , N_k and R_k , creating embeddings A_2^{tsne} , A_2^{mds} and so forth.

Classification. To quantify the utility of each embedding, we also perform a classification experiment. We train a 1-nearest neighbor classifier to distinguish between the different classification problems, based on the meta-representation. The classifiers are trained with a random subset of $\{5, 10, 20, 40, 60\}$ meta-samples 5 times, and accuracy is evaluated on the remaining meta-samples.

4 Results

Embedding. The embeddings are shown in Fig. 2. The embeddings based on accuracy are the best at discovering the true structure. Although we sampled the datasets in a balanced way in an attempt to remove some differences in accuracy, it is clear that some problems have larger or smaller class overlap, and therefore consistently lower or higher performances. Both t-SNE and MDS are able to recover this structure.

Looking at N_2 and R_2 , t-SNE appears to be slightly better at separating clusters of datasets from the same classification problem. This is particularly clear when looking at the ArteryVein datasets. Visually it is difficult to judge

⁴ <http://prtools.org/>

⁵ <https://lvdmaaten.github.io/tsne/>

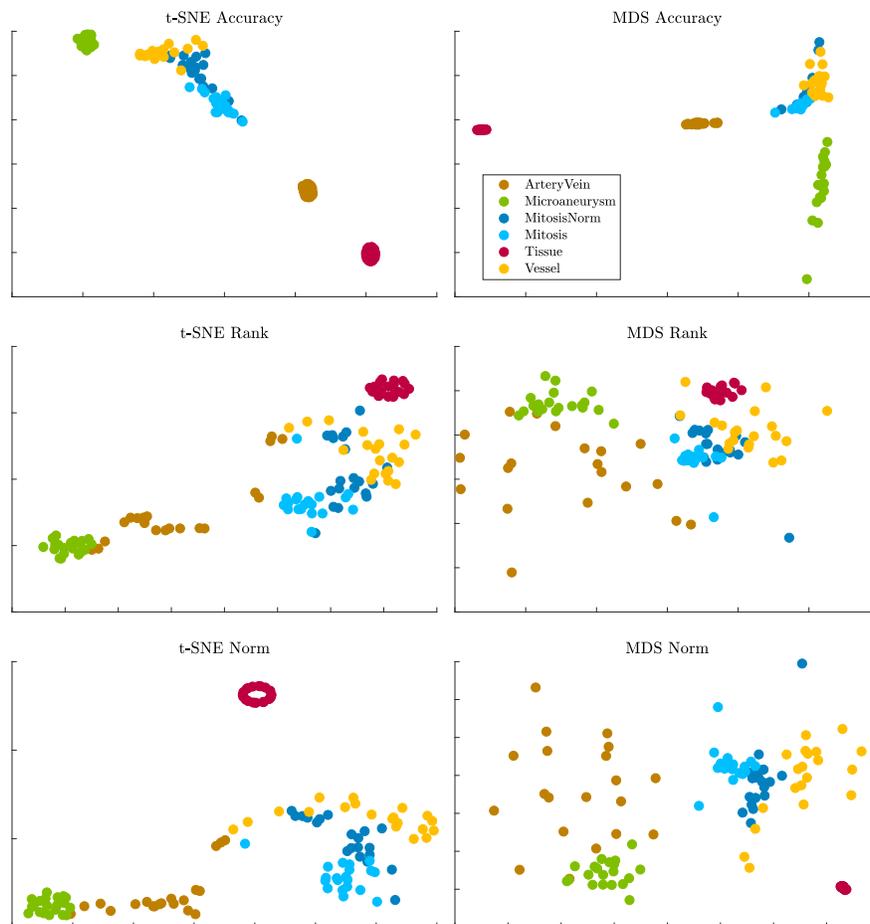


Fig. 2. 2D embeddings of the datasets with t-SNE (left) and MDS (right), based on accuracy (top), ranks (middle) and scaling (bottom).

whether N_2 or R_2 provides a more meaningful embedding, which is why it is instructive to look at how a classifier would perform with each of the embeddings.

Looking at the different clusters, several patterns start to emerge. ArteryVein and Microaneurysm are quite similar to each other, likely to the similarity of the images and the features used. Furthermore, Mitosis and MitosisNorm datasets are quite similar to each other, which is expected, because the same images but with different normalization are used. The Tissue dataset is often the most isolated from the others, being the only dataset based on 3D MR images.

Not all similarities can be explained by prior knowledge about the datasets. For example, we would expect the Vessel to be similar to the ArteryVein and Microaneurysm datasets, but in most embeddings it is actually more similar to Mitosis and MitosisNorm. This suggests that larger sample sizes and/or more meta-features are needed to better characterize these datasets.

Classification. The results of the 1-NN classifier, trained to distinguish between different classification problems, are shown in Fig. 3 (left). The performances confirm that the t-SNE embeddings are better than MDS. As expected, A_2 is the best embedding. Between N_2 and R_2 , which were difficult to assess visually, N_2 performs better. When trained on half (=60) of the meta-samples, 1-NN misclassifies 10.7% of the remaining meta-samples, which is low for a 6-class classification problem.

To assess which samples are misclassified, we examine the confusion matrix of this classifier in Fig. 3 (right). Most confusion can be found between Mitosis and MitosisNorm, ArteryVein and Microaneurysm and between Vessel and the Mitosis datasets, as would be expected from the embeddings.

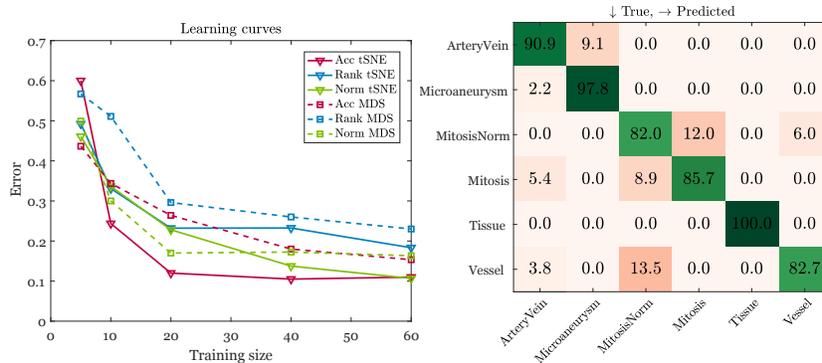


Fig. 3. Left: Learning curves of 1-NN classifiers trained on different-size samples from six meta-datasets. Right: confusion matrix of classifier trained on 60 samples from $N_2^{t_{sne}}$. Each cell shows what % of the true class (row) is classified as (column).

5 Discussion and Conclusions

We presented an approach to quantify the similarity of medical imaging datasets, by representing each dataset by performances of six linear and non-linear classifiers. Even though we used small samples from each dataset and only six classifiers, this representation produced meaningful clusters and was reasonably successful (89.3% accuracy) in predicting the origin of each dataset. This demonstrates the potential of using this representation in a meta-learning approach, with the goal of predicting which machine learning method will be effective for a previously unseen dataset.

The main limitation is the use of artificial meta-labels, based on each dataset’s original classification problem. Ideally the labels should reflect the best-performing method on the dataset. However, we would not expect the best-performing method to change for different samplings of the same dataset. Therefore, since we observe clusters with these artificial meta-labels, we also expect to observe clusters if more realistic meta-labels are used. Validating the approach with classifier-based meta-labels is the next step for a more practical application of this method.

Furthermore, we considered features as immutable properties of the classification problem. By considering the features as fixed, our approach would only be able to predict which classifier to apply to these already extracted features. However, due to recent advances in CNNs, where no explicit distinction is made between feature extraction and classification, we would want to start with the raw images. A challenge that needs to be addressed is how to represent the datasets at this point: for example, performances of classifiers on features extracted by CNNs pretrained on external data, or by some intrinsic, non-classifier-based characteristics. In a practical application, perceived similarity of the images (for example, obtained via crowdsourcing) could be used in addition to these features.

Despite these limitations, we believe this study reaches a more important goal: that of increasing awareness about meta-learning, which is largely overlooked by the medical imaging community. One opportunity is to use meta-learning jointly with transfer learning or domain adaptation, which have similar goals of transferring knowledge from a source dataset to a target dataset. For example, pretraining a CNN on the source, and extracting features on the target, is a form of transfer. In this context, meta-learning could be used to study which source datasets should be used for the transfer: for example, a single most similar source, or a selection of several, diverse sources.

Another opportunity is to learn more general rules of thumb for “what works when” by running the same feature extraction and classification pipelines on different medical imaging datasets, such as challenge datasets. In machine learning, this type of comparison is already being facilitated by OpenML [16], a large experiment database allows running different classification pipelines on already extracted features. We believe that a similar concept for medical imaging, that would also include preprocessing and feature extraction steps, would be a valuable resource for the community.

References

1. Vilalta, R., Drissi, Y.: A Perspective View and Survey of Meta-Learning. *Artificial Intelligence Review* **18** (2002) 77–95
2. Duin, R.P.W., Pekalska, E., Tax, D.M.J.: The characterization of classification problems by classifier disagreements. In: *International Conference on Pattern Recognition*. Volume 1. (2004) 141–143
3. Cheplygina, V., Tax, D.M.J.: Characterizing multiple instance datasets. In Fergan, A., Pelillo, M., , Loog, M., eds.: *Similarity-Based Pattern Recognition: Third International Workshop, SIMBAD 2015, Copenhagen, Denmark, October 12-14, 2015*. Proceedings, Springer International Publishing (2015) 15–27
4. Muenzing, S.E.A., van Ginneken, B., Viergever, M.A., Pluim, J.P.W.: DIRBoost—an algorithm for boosting deformable image registration: Application to lung CT intra-subject registration. *Medical Image Analysis* **18**(3) (2014) 449–459
5. Gurari, D., Jain, S.D., Betke, M., Grauman, K.: Pull the Plug? Predicting If Computers or Humans Should Segment Images. In: *Computer Vision and Pattern Recognition*. (2016) 382–391
6. Cox, T.F., Cox, M.A.: *Multidimensional scaling*. CRC Press (2000)
7. van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *Journal of Machine Learning Research* **9** (2008) 2579–2605
8. Landman, B.A., et al.: *MICCAI 2012 Workshop on Multi-Atlas Labeling*. CreateSpace Independent Publishing Platform (2012)
9. Moeskops, P., Viergever, M.A., Mendrik, A.M., de Vries, L.S., Benders, M.J., Išgum, I.: Automatic segmentation of MR brain images with a convolutional neural network. *IEEE Transactions on Medical Imaging* **35**(5) (2016) 1252–1261
10. Veta, M., Van Diest, P.J., Willems, S.M., Wang, H., Madabhushi, A., Cruz-Roa, A., Gonzalez, F., Larsen, A.B., Vestergaard, J.S., Dahl, A.B., et al.: Assessment of algorithms for mitosis detection in breast cancer histopathology images. *Medical Image Analysis* **20**(1) (2015) 237–248
11. Veta, M., van Diest, P.J., Jiwa, M., Al-Janabi, S., Pluim, J.P.W.: Mitosis counting in breast cancer: Object-level interobserver agreement and comparison to an automatic method. *PLoS ONE* **11**(8) (2016) e0161286
12. Staal, J., Abràmoff, M.D., Niemeijer, M., Viergever, M.A., van Ginneken, B.: Ridge-based vessel segmentation in color images of the retina. *IEEE Transactions on Medical Imaging* **23**(4) (2004) 501–509
13. Zhang, J., Dashtbozorg, B., Bekkers, E., Pluim, J.P.W., Duits, R., ter Haar Romeny, B.M.: Robust retinal vessel segmentation via locally adaptive derivative frames in orientation scores. *IEEE Transactions on Medical Imaging* **35**(12) (2016) 2631–2644
14. Dashtbozorg, B., Mendonça, A.M., Campilho, A.: An automatic graph-based approach for artery/vein classification in retinal images. *IEEE Transactions on Image Processing* **23**(3) (2014) 1073–1083
15. Decencièrè, E., et al.: TeleOphta: Machine learning and image processing methods for teleophthalmology. *IRBM* **34**(2) (2013) 196–203
16. Vanschoren, J., Van Rijn, J.N., Bischl, B., Torgo, L.: OpenML: networked science in machine learning. *ACM SIGKDD Explorations Newsletter* **15**(2) (2014) 49–60