

---

# Unsupervised deep clustering for predictive texture pattern discovery in medical images

---

Matthias Perkonigg

Daniel Sobotka

Ahmed Ba-Ssalamah

Georg Langs

Computational Imaging Research Lab (CIR)  
Department of Biomedical Imaging and Image-guided Therapy  
Medical University of Vienna, Austria

## Abstract

Predictive marker patterns in imaging data are a means to quantify disease and progression, but their identification is challenging, if the underlying biology is poorly understood. Here, we present a method to identify predictive texture patterns in medical images in an unsupervised way. Based on deep clustering networks, we simultaneously encode and cluster medical image patches in a low-dimensional latent space. The resulting clusters serve as features for disease staging, linking them to the underlying disease. We evaluate the method on 70 T1-weighted magnetic resonance images of patients with different stages of liver steatosis. The deep clustering approach is able to find predictive clusters with a stable ranking, differentiating between low and high steatosis with an F1-Score of 0.78.

## 1 Introduction

Marker patterns serve prediction and quantification of disease status, and their machine learning based discovery is relevant if the underlying biological processes are poorly understood, or no a priori defined markers are known. Here, we pose the identification of marker pattern candidates in imaging data of patients with a particular disease as a simultaneous dimensionality reduction and clustering problem. We perform unsupervised deep clustering of texture patterns in medical images and identify those patterns linked to the stages of an underlying disease present in the population. The discovery of frequent image patterns contributing to classification enables the analysis of regions where those patterns occur, and is more interpretable than global image features. Several relevant approaches exist [2, 12, 13, 3], and unsupervised clustering of deep features in the medical domain has shown promise for unsupervised segmentation [7] or image categorization [11]. We build on Deep Clustering Networks (DCN) jointly learning and clustering latent space features instead of performing these steps sequentially [12].

We perform deep convolutional clustering on medical image patches (image segments) in a real-world problem setting to find distinct, predictive texture patterns that can be used for the staging of liver steatosis in MRI data. Steatosis is an accumulation of fat within the liver. It is one of the characteristics of Nonalcoholic Fatty Liver Disease (NAFLD) [5]. Steatosis is graded into four stages (Grade 0: <5%, Grade 1: 5%-33%, Grade 2: 33%-66%, Grade 3: >66%) and can be assessed by biopsy, ultrasound or MRI-Proton Density Fat Fraction (PDFF) and MR elastography (MRE) [4]. In this work we link texture patterns in T1 weighted magnetic resonance imaging (MRI) data discovered by our unsupervised approach to the steatosis grade.

## 2 Methods

The proposed method can be divided into four major steps: First, we randomly extract 2D axial patches within an anatomical structure - in this case the liver - from image volumes. Then, we simultaneously encode and cluster appearance variability in these patches in a latent space with DCN, and apply the trained DCN to calculate a signature of relative cluster proportions for each case. Finally, we link signatures to disease staging by training a random forest classifier [1], and scoring the contribution of each cluster to classification accuracy.

**Encoding and clustering of image patches with convolutional DCNs** Building upon the DCN [12] approach we use a deep convolutional autoencoder to encode a set of image patches in a latent space. Three convolutional layers (50, 20 and 10 feature maps respectively) with receptive field of 3x3 and ReLU activation followed by a 2x2 MaxPooling are used as an encoder to a 20 dimensional latent space. Deconvolution in the decoder is performed by upsampling followed by a convolutional layer in the reversed order. In addition to the reconstruction loss of the conventional autoencoder, DCN uses a second loss term measures the 'cluster-friendliness' derived from a k-means algorithm [6]. The full loss function is

$$\min_{W, Z, M, \{s_i\}} \sum_{n=1}^N (l(\mathbf{g}(\mathbf{f}(x_i), x_i)) + \lambda \|\mathbf{f}(x_i) - M s_i\|_2^2), \quad (1)$$

where  $W$  and  $Z$  are the parameters of the encoder ( $\mathbf{f}$ ) and decoder ( $\mathbf{g}$ ) respectively.  $M$  are the centroids of the clusters and  $\{s_i\}$  is the cluster membership of patch  $i$ .  $l$  is a reconstruction loss function, in our approach we use the mean-squared error.  $\lambda$  is a hyperparameter that regulates the trade-off between reconstruction quality and 'cluster-friendliness' [12].

The network is trained in an alternating gradient descent scheme of (1) updating the network parameters, (2) updating the cluster assignments of the training patches and (3) updating the cluster centroids [12].

**Calculation of the signature of an image region** After training the convolutional DCN (cDCN), we apply it to a whole region of interest (ROI, e.g., anatomical structure) by parsing it with a sliding window, and assigning each position to one of the learned clusters. For each ROI the vector of the relative proportion of each cluster forms the *signature* of the ROI. This approach is similar to Bag-of-Visual-Words [9].

**Linking signatures to clinical parameters** If signature components are linked to clinical parameters, they hold meaningful pattern variations of the ROI. We calculate the signatures for all ROIs in the training data, and train a classifier to predict steatosis grade from the signature. We used Random Forests [1] and LASSO regression [10] to quantify the link of individual signature components to disease stage, by quantifying its contribution to classifier accuracy, yielding an interpretable set of predictive components.

## 3 Results

The method is evaluated on multi-center data of T1-weighted MRI volumes of 70 NASH patients for whom diagnosis and true steatosis grades are confirmed by biopsy (grade 0/1/2/3; n=1/20/32/17). We evaluate classification of individual grades and of low/high steatosis (grades 0&1/2&3). It is worth noting that in T1-weighted MRIs, other than in e.g. MRI-PDFF, the steatosis grade is not directly quantifiable [4]. By linking signatures of T1-weighted MRIs to steatosis grade, we demonstrate that our approach can identify patterns that capture disease progression, even if no a priori known marker exists in the data. Prior to analysis, livers are segmented with a U-Net [8] and masks are manually corrected. For training the cDCN 50.000 image patches with a size of 14x14mm (resampled to 32x32 pixels) are extracted at random positions within the livers of the dataset. We train the DCN to cluster the patches into 10 clusters. The number of clusters is a hyperparameter of DCN and is fixed after preliminary experiments. Signatures for all MRI volumes are calculated within the liver. We applied classification and predictive signature component identification using LASSO regression and a random forest classifier using a Leave-One-Out cross validation scheme.

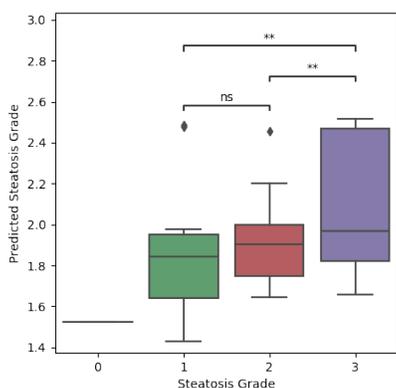


Figure 1: Regression values with corresponding ground truth steatosis grades.

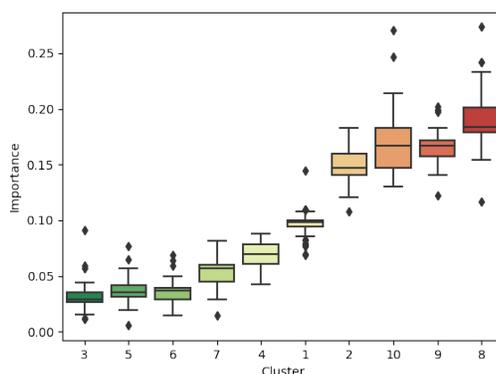


Figure 2: Feature Importance of the random forest classifier for low/high steatosis classification.

The results of predicting grade from signature are shown in Figure 1. Steatosis grade 3 can be reliably separated from all other grades and predicted values are correlated with true steatosis grade estimated from biopsy. Low/high steatosis classification reaches an accuracy of 0.70, Sensitivity of 0.78, Specificity of 0.52 and an F-Score of 0.78. Figure 2 shows the feature importance and its variability across leave-one-out samples in the population for a random forest classifier between low/high steatosis. Clusters 2, 10, 9 and 8 are consistently the most important features in the Leave-One-Out cross validation. To illustrate the presence and distribution of signature components, in Figure 3 those four clusters are shown on MR images from patients with different grades of steatosis. Two of the main predictive components (8 and 9) focus on regions around vessels.

## 4 Discussion

We applied unsupervised clustering to a real-world problem of finding predictive features in medical imaging. We show how clusters can be linked to clinical parameters and demonstrate that the identified predictive signature components are stable and interpretable by humans.

We link the signature components identified in T1w MRI data to steatosis grade. This grade can be non-invasively assessed [4], but results demonstrate that cDCN can identify predictive features in data where no a priori markers are known. Ground-truth steatosis grade based on biopsy is limited to a single position in the liver, while we assess the steatosis grade over the whole liver, possibly leading to discrepancy due to heterogeneity in the liver. Future work will extend our approach to

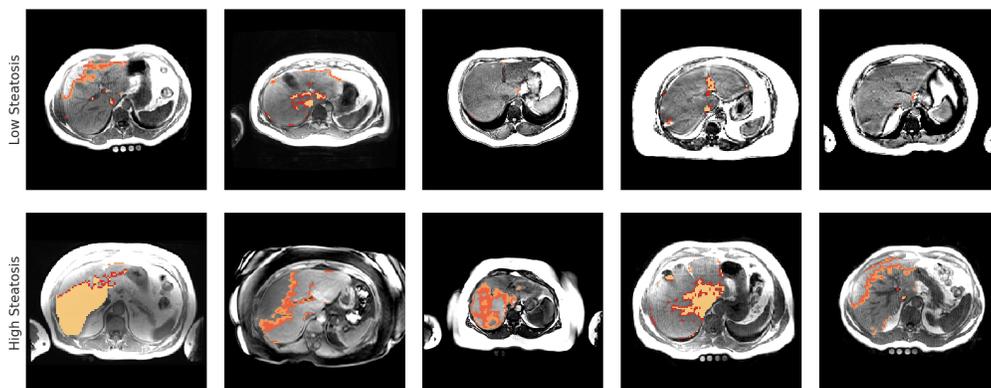


Figure 3: Examples of a low and a high steatosis patients with overlays showing the distribution of the four most informative signature components in these cases.

multi-parametric MRI and link clusters to other clinical parameters such as inflammation. Aside from linking signature components to known changes associated with a disease, they offer the possibility to identify new hypotheses regarding imaging biomarker candidates.

## Acknowledgments

This work was supported by the Austrian Science Fund (FWF) I 2714B31, OENB 18207 and Novartis Pharmaceuticals Corporation.

## References

- [1] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [2] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 132–149, 2018.
- [3] Bert De Brabandere and Davy Neven. Semantic Instance Segmentation with a Discriminative Loss Function. *arXiv 1708.02551v1*.
- [4] Parambir S. Dulai, Claude B. Sirlin, and Rohit Loomba. MRI and MRE for non-invasive quantitative assessment of hepatic steatosis and fibrosis in NAFLD and NASH: Clinical trials to clinical practice. *Journal of Hepatology*, 65(5):1006–1016, nov 2016.
- [5] David E Kleiner, Elizabeth M Brunt, Mark Van Natta, Cynthia Behling, Melissa J Contos, Oscar W Cummings, Linda D Ferrell, Yao-Chang Liu, Michael S Torbenson, Aynur Unalp-Arida, et al. Design and validation of a histological scoring system for nonalcoholic fatty liver disease. *Hepatology*, 41(6):1313–1321, 2005.
- [6] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- [7] Takayasu Moriya, Holger R Roth, Shota Nakamura, Hirohisa Oda, Kai Nagara, Masahiro Oda, and Kensaku Mori. Unsupervised segmentation of 3d medical images based on clustering and deep representation learning. In *Medical Imaging 2018: Biomedical Applications in Molecular, Structural, and Functional Imaging*, volume 10578, page 1057820. International Society for Optics and Photonics, 2018.
- [8] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [9] Josef Sivic and Andrew Zisserman. Video google: A text retrieval approach to object matching in videos. In *Proceedings Ninth IEEE International Conference on Computer Vision*, page 1470. IEEE, 2003.
- [10] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [11] Xiaosong Wang, Le Lu, Hoo-Chang Shin, Lauren Kim, Mohammadhadi Bagheri, Isabella Noguees, Jianhua Yao, and Ronald M Summers. Unsupervised joint mining of deep features and image labels for large-scale radiology image categorization and scene recognition. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 998–1007. IEEE, 2017.
- [12] Bo Yang, Xiao Fu, Nicholas D Sidiropoulos, and Mingyi Hong. Towards k-means-friendly spaces: Simultaneous deep learning and clustering. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3861–3870. JMLR. org, 2017.
- [13] Jianwei Yang, Devi Parikh, and Dhruv Batra. Joint unsupervised learning of deep representations and image clusters. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5147–5156, 2016.