

1580

q-Space Deep Learning for Alzheimer's Disease Diagnosis: Global Prediction and Weakly-Supervised Localization

Vladimir Golkov¹, Phillip Swazinna¹, Marcel M. Schmitt¹, Qadeer A. Khan¹, Chantal M.W. Tax², Marat Serahlazau¹, Francesco Pasa^{1,3}, Franz Pfeiffer³, Geert Jan Biessels⁴, Alexander Leemans⁵, and Daniel Cremers¹

¹Department of Informatics, Technical University of Munich, Munich, Germany, ²CUBRIC, Cardiff University, Cardiff, United Kingdom, ³Physics Department, Technical University of Munich, Munich, Germany, ⁴Department of Neurology, University Medical Center Utrecht, Utrecht, Netherlands, ⁵Image Sciences Institute, University Medical Center Utrecht, Utrecht, Netherlands

Synopsis

Most diffusion MRI approaches rely on comparably long scan time and a suboptimal processing pipeline with handcrafted physical/mathematical representations. They can be outperformed by recent handcrafted-representation-free methods. For instance, q-space deep learning (q-DL) allows unprecedentedly short scan times and optimized voxel-wise tissue characterization. We reformulate q-DL such that it estimates global (i.e. scan-wise rather than voxel-wise) information. We use this formulation to distinguish Alzheimer's disease (AD) patients from healthy controls based solely on raw q-space data without handcrafted representations such as DTI. Classification quality is very promising. Weakly-supervised localization techniques indicate that the neural network attends to AD-relevant brain areas.

Introduction

Most diffusion MRI approaches rely on relatively long scan time and a suboptimal processing pipeline with handcrafted physical/mathematical representations. They can be outperformed by recent handcrafted-representation-free methods.^{1,2} For instance, q-space deep learning¹ (q-DL) allows unprecedentedly short scan times and optimized voxel-wise tissue characterization.^{1,2} Here we reformulate q-DL such that it estimates global (i.e. scan-wise rather than voxel-wise) information. We use this formulation of q-DL to distinguish patients with Alzheimer's disease (AD) from healthy controls based solely on raw q-space data without any handcrafted representations such as diffusion tensor imaging.

Methods

Data: The data were as follows:^{3,4} 47 AD patients, 58 healthy controls, one $b=0$ image (averaged over 3 repetitions), 45 diffusion directions ($b=1200\text{s/mm}^2$, single-shot SE-EPI, $TR=6638\text{ms}$, $TE=73\text{ms}$, voxels $1.72\text{mm}\times 1.72\text{mm}\times 2.5\text{mm}$, matrix 128×128 , 48 axial slices), motion/distortion-corrected using ExploreDTI.⁵ Following deep learning terminology, we refer to each of these 45+1 contrasts (diffusion-weighted images) as "channels". To study the effect of scan time reduction, separate experiments were performed using sets of 46/30/23/15/7/4 randomly selected channels. For convenient neural network training, so-called feature scaling was performed by dividing each channel by the corresponding channel mean taken over all scans. To prevent overfitting on intensity values, each scan was additionally divided by its mean intensity, and during each training iteration multiplied by a random value between 0.5 and 1.5.

Deformations: For scan-wise AD prediction, we consider each scan a sample (in the machine-learning sense). With a relatively small number of samples (i.e. scans) and large number of features (i.e. voxel values), the neural network would overfit unique but disease-unrelated local image patterns. To circumvent this problem, we generated "additional" images (sized $128\times 128\times 48\times 46$) by augmenting the data with random elastic spatial deformations of the original images (sized $128\times 128\times 48\times 46$). Thus, the local information varies, while the global context and q-space (important for diagnosis) remain similar. Deformation is done by generating a random coarse $8\times 8\times 4$ vector grid, upsampling it to $128\times 128\times 48$ and applying the resulting deformation field to the image. All channels of an image are transformed jointly. Fig. 1 illustrates the influence of the coarse grid parameters. Fig. 2 shows an example on real data.

Network: Five-fold cross-validation was performed: The dataset was split into five equally-sized subsets, and in each fold three of the subsets were used for training, one was used for validation (early stopping) and one for testing (results are reported on the test set of each of the five folds). The neural network architecture was C128-P-C256-P-C512-GP-FC2000-FC1, where Cn is a convolutional layer with n filters sized $3\times 3\times 3$, P is a $2\times 2\times 2$ max-pooling layer, GP is a global-pooling layer, and FC n is a fully-connected layer with n units. Hidden-layer nonlinearities: $\text{ReLU}(z)=\max\{z,0\}$, output nonlinearity: sigmoid, trained with binary cross-entropy loss, Adam algorithm,⁶ learning rate $2\cdot 10^{-5}$, to distinguish AD patients from healthy controls.

Results

Figure 3 shows the receiver operating characteristic (ROC) for the test sets from the five cross-validation folds. Area under the curve (AUC) in the five cross-validation folds ranges between 0.85 and 0.96. Dependency of AUC on channel number is illustrated in Fig. 4.

Gradient-based class activation mapping⁷ (Grad-CAM) is a "weakly-supervised localization" technique to visualize on what spatial image area the network bases its decision. Its combination with saliency maps (e.g. guided backpropagation⁸) that model which inputs strongly influence the prediction, is called Guided Grad-CAM.⁷ We use these techniques to examine which brain areas drive the network's decision. Visualization of Guided Backpropagation⁸ and (Guided) Grad-CAM⁷ for an AD patient that the network correctly classified are shown in Fig. 5.

Discussion

We can thus report that one can train a network directly on q-space data, without any handcrafted representations needed when classifying AD. The network can compensate for missing channels, i.e. providing good predictions for relatively short scan times. AUC drops substantially only when extremely few channels are exposed to the network.

Since the network is able to make good predictions with less channels, it is likely that the essential information is included in few channels. Potentially the network is also able to reconstruct missing information from the exposed subset of channels. Even though the diagnosis was the only output target information available to the network, it seems to focus on brain regions that are prone to AD as can be seen in Fig. 5. This behavior is highly desired since it shows the network's ability to generalize to unseen data in a meaningful way.

Conclusions

We conclude, (1) that it is possible to directly use raw q-space data as inputs for a convolutional neural network and obtain good classification results when detecting AD, and (2) that the convolutional neural networks seem to learn with weak supervision which areas of the brain are affected by AD.


Acknowledgements

This study was partly funded by the European Research Council (ERC Consolidator Grant "3DReloaded"), Deutsche Telekom Foundation, German Research Council (DFG, LA 2804/1-1) and the contribution of Chantal Tax is supported by an FC-EW grant (No. 612.001.104) from the Dutch Scientific Foundation (NWO). The authors would like to thank the members of the Utrecht Vascular Cognitive Impairment Study Group for providing the diffusion MRI data (in alphabetical order by department): University Medical Center Utrecht, the Netherlands, Department of Neurology: E. van den Berg, G.J. Biessels, M. Brundel, W.H. Bouvy, S.M. Heringa, L.J. Kappelle, Y.D. Reijmer; Department of Radiology/Image Sciences Institute: J. de Bresser, H.J. Kuijff, A. Leemans, P.R. Luijten, W.P.Th.M. Mali, M.A. Viergever, K.L. Vincken, J.J.M. Zwanenburg; Department of Geriatrics: H.L. Koek, J.E. de Wit; Hospital Diaconessenhuis Zeist, the Netherlands: M. Hamaker, R. Faaij, M. Pleizier, E. Vriens.


References

1. V. Golkov et al. q-Space deep learning: twelve-fold shorter and model-free diffusion MRI scans. *IEEE Trans Med Imag.* 2016;35(5):1344–1351.
2. V. Golkov et al. Model-free novelty-based diffusion MRI. *Proc IEEE Int Symp Biomed Imag (ISBI)*. 2016:1233–1236.
3. Y.D. Reijmer, A. Leemans, K. Caeyenberghs, et al. Disruption of cerebral networks and cognitive impairment in Alzheimer disease. *Neurology* 2013;80(15):1370–1377.
4. M. Bach, F.B. Laun, A. Leemans, et al. Methodological considerations on tract-based spatial statistics (TBSS). *NeuroImage* 2014;100:358–369.
5. A. Leemans, B. Jeurissen, J. Sijbers, D.K. Jones. ExploreDTI: a graphical toolbox for processing, analyzing, and visualizing diffusion MR data. *ISMRM* 2009:3537.
6. D. Kingma, J. Ba. Adam: a method for stochastic optimization. *Proc Int C Learning Representations (ICLR)* 2015:arXiv:1412.6980.
7. R.R. Selvaraju et al. Grad-CAM: visual explanations from deep networks via gradient-based localization. *Proc Int C Comp Vision (ICCV)* 2017:618–626.
8. J.T. Springenberg, A. Dosovitskiy, T. Brox, M. Riedmiller. Striving for simplicity: the all convolutional net. *ICLR Workshops* 2015.


Figures




2D slice of random elastic deformations of a 128x128x48 checkerboard image using deformation vectors with different coarse-grid densities and standard deviations. Grid density 8x8x4 and standard deviation of 1 voxel were used for MRI data augmentation.




Original 128x128 2D MRI slice (left) and randomly elastically deformed version for training with data augmentation (right).



ROC for Alzheimer's disease classification using global q-DL for the five folds of cross-validation.



Dependency of AUC on selected number of channels (diffusion-weighted images) that are exposed to the network during training and testing. Error bars show standard deviation across five folds of cross-validation. AUC remains high even with fewer channels, but drops when the network is trained on very few channels.



Visualization of Grad-CAM, Guided Grad-CAM and Guided Backpropagation overlaid with the b=0 channel of an AD patient. Highlighted regions mark brain areas that were important for the network's classification decision.

q-Space Deep Learning for Alzheimer's Disease Diagnosis: Global Prediction and Weakly-Supervised Localization

Vladimir Golkov¹ Phillip Swazinna¹ Marcel M. Schmitt¹ Qadeer A. Khan¹ Chantal M.W. Tax² Marat Serahlazau¹ Francesco Pasa^{1,3} Franz Pfeiffer³ Geert Jan Biessels⁴ Alexander Leemans⁵ Daniel Cremers¹
¹ Computer Vision Group, Technical University of Munich, Germany ² CUBRIC, Cardiff University, UK ³ Physics Department, Technical University of Munich, Germany
⁴ Department of Neurology, University Medical Center Utrecht, Netherlands ⁵ Image Sciences Institute, University Medical Center Utrecht, Netherlands

Goals

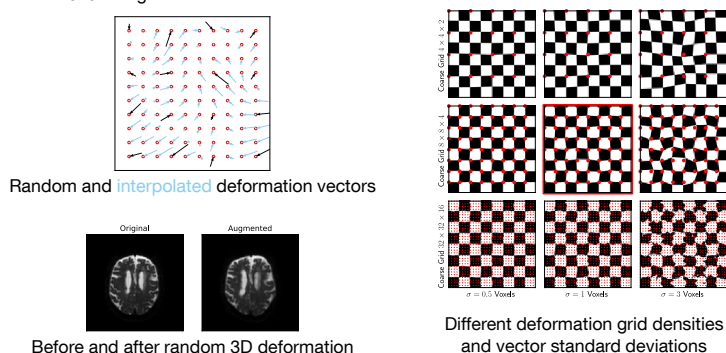
- q-Space deep learning [1,2]:
 - Prediction of tissue properties directly from q-space data
 - Usually voxel-wise
- Here:
 - AD/healthy label for entire scan (not voxel-wise)
 - ConvNet with global prediction
 - Weakly-supervised learning: voxel-wise reasons for global decision

Overfitting of Intensities → Data Augmentation 1

- Prevent overfitting of image intensity values:
 - Divide each scan by its mean intensity
 - Then multiply each scan by a random number between 0.5 and 1.5 (different in each training iteration)


Overfitting of Patches → Data Augmentation 2

- Few scans, many voxels → danger of overfitting irrelevant image patterns
- Random elastic spatial deformations → more variability of irrelevant patterns, less overfitting



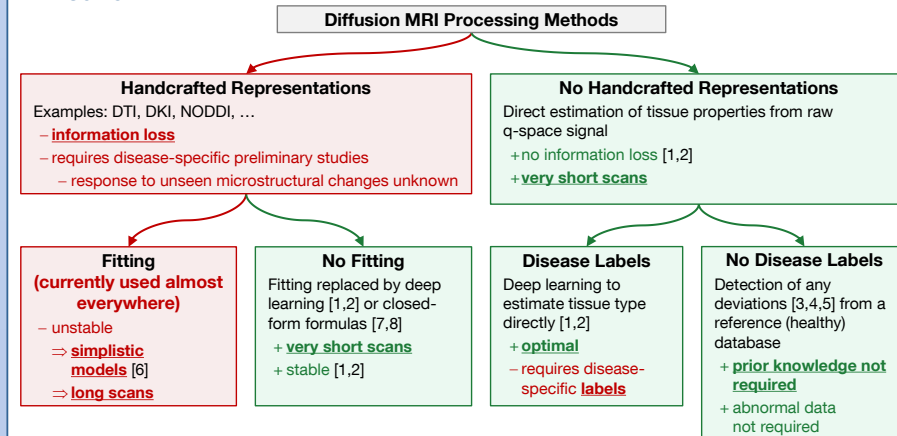
Overfitting of Distortion Artifacts → Cropping

- Reliable classification, but network attends to distorted regions (probably each scan has unique "overfittable" features there)



- Solution: Crop training images
- This reveals the next "overfitting problem"

Context




Weakly-Supervised Learning Methods


- One label for every **scan**, one prediction for every **voxel**
- Global supervised learning:
 - One label for every **scan**, one prediction for every **scan**
- Local supervised learning:
 - One label for every **voxel**, one prediction for every **voxel**
- Global label: Alzheimer's disease
- Local prediction: which areas of given image influence classifier the most

Class Activation Mapping (CAM) [12,13]

- ConvNets (illustration: 2D) often have a global pooling layer
- Train with global pooling (→ global prediction)



- Test without global pooling (→ local prediction)



- Fully-connected layers (if any) can be considered as 1x1 convolutional layers
- Result: local prediction of a network that was trained for global prediction
- Fewer fully-connected layers → more well-founded

Guided Backpropagation [14]

- Which voxel intensities should be changed how to strongly influence the prediction


Data [9,10]

- 47 AD patients, 58 healthy controls, 5-fold cross-validation
- One $b=0$ image (averaged over 3 repetitions), 45 diffusion directions ($b=1200\text{s/mm}^2$)
- Single-shot SE-EPI, TR=6638ms, TE=73ms, voxels $1.72\text{mm} \times 1.72\text{mm} \times 2.5\text{mm}$, matrix 128×128 , 48 axial slices, motion/distortion-corrected using ExploreDTI [11]

Network and Training


- 3D ConvNet: C128-P-C256-P-C512-GP-FC2000-FC1
 - C_n – 3D convolutional layer with n $3 \times 3 \times 3$ filters
 - P – 3D $2 \times 2 \times 2$ max-pooling
 - GP – 3D global pooling
 - FC_n – fully-connected layer with n neurons
- Hidden layers: ReLU nonlinearity, output: sigmoid
- Binary cross-entropy loss (target: AD/healthy label), Adam optimizer [15], learning rate $2 \cdot 10^{-5}$

Good Classification, Relevant Brain Regions




"Over"fitting of Macrostructure & Open Questions

- Network attends to ventricles (enlarged in AD, easy macrostructural features) instead of q-space (microstructural features)



- This is further supported by the fact that just the $b=0$ image yields good results



- Open question: smaller receptive fields to use only q-space information, but for global prediction

Discussion & Conclusions

- q-Space deep learning for global (scan-wise) prediction yields good results
- No handcrafted representations (such as DTI or NODDI)
- Diagnosis directly from raw q-space measurements, data-driven
- ConvNet with large receptive field uses macrostructural features (image space, ventricles) more than microstructural features (q-space in brain voxels) for Alzheimer's disease classification
- Open question: study q-space information – small (e.g. minimal [1]) receptive field combined with scan-wise labels

References

- V. Golkov et al. *q-Space Deep Learning for Twelve-Fold Shorter and Model-Free Diffusion MRI Scans*. MICCAI 2015
 - V. Golkov et al. *q-Space Deep Learning: Twelve-Fold Shorter and Model-Free Diffusion MRI Scans*. IEEE TMI 2016
 - V. Golkov et al. *Model-Free Novelty-Based Diffusion MRI*. ISBI 2016
 - V. Golkov et al. *q-Space Novelty Detection in Short Diffusion MRI Scans of Multiple Sclerosis*. ISMRM 2018 (#5378)
 - A. Vasiliev et al. *q-Space Novelty Detection with Variational Autoencoders*. arXiv:1806.02997 2018
 - U. Ferizi et al. *White matter compartment models for in vivo diffusion MRI at 300 mT/m*. NeuroImage 2015
 - B. Hansen et al. *Experimentally and computationally fast method for estimation of a mean kurtosis*. MRM 2013
 - B. Hansen et al. *Fast imaging of mean, axial and radial diffusion kurtosis*. NeuroImage 2016
 - Y.D. Reijmer, A. Leemans, K. Caeyenberghs, et al. *Disruption of cerebral networks and cognitive impairment in Alzheimer disease*. Neurology 2013
 - M. Bach, F.B. Laun, A. Leemans, et al. *Methodological considerations on tract-based spatial statistics (TBSS)*. NeuroImage 2014
 - A. Leemans, B. Jeurissen, J. Sijbers, D.K. Jones. *ExploreDTI: a graphical toolbox for processing, analyzing, and visualizing diffusion MR data*. ISMRM 2009
 - B. Zhou et al. *Learning Deep Features for Discriminative Localization*. CVPR 2016
 - R.R. Selvaraju et al. *Grad-CAM: visual explanations from deep networks via gradient-based localization*. ICCV 2017
 - J.T. Springenberg, A. Dosovitskiy, T. Brox, M. Riedmiller. *Striving for simplicity: the all convolutional net*. ICLR Workshops 2015
 - D. Kingma, J. Ba. *Adam: a method for stochastic optimization*. ICLR 2015
- Acknowledgments: This study was partly funded by the European Research Council, Deutsche Telekom Foundation, German Research Council (DFG, LA 2804/1-1), and the contribution of Chantal Tax is supported by an FC-EW grant (No. 612.001.104) from the Dutch Scientific Foundation (NWO). The authors would like to thank the members of the Utrecht Vascular Cognitive Impairment Study Group for providing the diffusion MRI data (in alphabetical order by department): University Medical Center Utrecht, the Netherlands, Department of Neurology: E. van den Berg, G.J. Biessels, M. Brundel, W.H. Bouvy, S.M. Heringa, L.J. Kappelle, Y.D. Reijmer; Department of Radiology/Image Sciences Institute: J. de Bresser, H.J. Kuijff, A. Leemans, P.R. Luijten, W.P.Th.M. Mali, M.A. Viergever, K.L. Vincken, J.J.M. Zwambag; Department of Geriatrics: H.L. Koek, J.E. de Wit; Hospital Diaconessenhuis Zeist, the Netherlands: M. Hamaker, R. Faaij, M. Pleizier, E. Vriens.