



GAMER MRI: Gated-attention mechanism ranking of multi-contrast MRI in brain pathology

Po-Jui Lu^{a,b,c}, Youngjin Yoo^d, Reza Rahmanzadeh^{a,b,c}, Riccardo Galbusera^{a,b,c}, Matthias Weigel^{a,b,c,e}, Pascal Ceccaldi^d, Thanh D. Nguyen^f, Pascal Spincemaille^f, Yi Wang^f, Alessandro Daducci^g, Francesco La Rosa^{h,i,j}, Meritxell Bach Cuadra^{h,i,j}, Robin Sandkühler^k, Kambiz Nael^{l,m}, Amish Doshi^m, Zahi A. Fayad^{m,n}, Jens Kuhle^{b,c}, Ludwig Kappos^{b,c}, Benjamin Odry^o, Philippe Cattin^k, Eli Gibson^d, Cristina Granziera^{a,b,c}

^a Translational Imaging in Neurology (THINK) Basel, Department of Medicine and Biomedical Engineering, University Hospital Basel and University of Basel, Basel, Switzerland

^b Neurologic Clinic and Policlinic, Departments of Medicine, Clinical Research and Biomedical Engineering, University Hospital Basel and University of Basel, Basel, Switzerland

^c Research Center for Clinical Neuroimmunology and Neuroscience Basel, University Hospital Basel and University of Basel, Basel, Switzerland

^d Digital Technology and Innovation, Siemens Healthineers, Princeton, NJ, USA

^e Division of Radiological Physics, Department of Radiology, University Hospital Basel, Basel, Switzerland

^f Department of Radiology, Weill Cornell Medical College, New York, NY, USA

^g Computer Science Department, University of Verona, Verona, Italy

^h Signal Processing Laboratory (LTS5), Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

ⁱ Medical Image Analysis Laboratory, Center for Biomedical Imaging (CIBM), University of Lausanne, Lausanne, Switzerland

^j Department of Radiology, Lausanne University Hospital and University of Lausanne, Lausanne, Switzerland

^k Center for Medical Image Analysis & Navigation, Department of Biomedical Engineering, University of Basel, Allschwil, Switzerland

^l Department of Radiological Sciences, David Geffen School of Medicine at University of California, Los Angeles, CA, USA

^m Department of Diagnostic, Molecular and Interventional Radiology, Icahn School of Medicine at Mount Sinai, New York, NY, USA

ⁿ BioMedical Engineering and Imaging Institute, Department of Radiology, Icahn School of Medicine at Mount Sinai, New York, NY, USA

^o AI for Clinical Analytics, Covera Health, New York, NY, USA

ARTICLE INFO

Keywords:

Deep learning
Attention mechanism
Relative importance order
Stroke
Multiple sclerosis
Quantitative MRI

ABSTRACT

Introduction: During the last decade, a multitude of novel quantitative and semiquantitative MRI techniques have provided new information about the pathophysiology of neurological diseases. Yet, selection of the most relevant contrasts for a given pathology remains challenging. In this work, we developed and validated a method, Gated-Attention Mechanism Ranking of multi-contrast MRI in brain pathology (GAMER MRI), to rank the relative importance of MR measures in the classification of well understood ischemic stroke lesions. Subsequently, we applied this method to the classification of multiple sclerosis (MS) lesions, where the relative importance of MR measures is less understood.

Methods: GAMER MRI was developed based on the gated attention mechanism, which computes attention weights (AWs) as proxies of importance of hidden features in the classification. In the first two experiments, we used Trace-weighted (Trace), apparent diffusion coefficient (ADC), Fluid-Attenuated Inversion Recovery (FLAIR), and T1-weighted (T1w) images acquired in 904 acute/subacute ischemic stroke patients and in 6,230 healthy controls and patients with other brain pathologies to assess if GAMER MRI could produce clinically meaningful importance orders in two different classification scenarios. In the first experiment, GAMER MRI with a pretrained convolutional neural network (CNN) was used in conjunction with Trace, ADC, and FLAIR to distinguish patients with ischemic stroke from those with other pathologies and healthy controls. In the second experiment, GAMER MRI with a patch-based CNN used Trace, ADC and T1w to differentiate acute ischemic stroke lesions from healthy tissue. The last experiment explored the performance of patch-based CNN with GAMER MRI in ranking the importance of quantitative MRI measures to distinguish two groups of lesions with different pathological characteristics and unknown quantitative MR features. Specifically, GAMER MRI was applied to assess the relative importance of the myelin water fraction (MWF), quantitative susceptibility mapping (QSM), T1 relaxometry map (qT1), and neurite density index (NDI) in distinguishing 750 juxtacortical lesions

<https://doi.org/10.1016/j.nicl.2020.102522>

Received 3 September 2020; Received in revised form 11 November 2020; Accepted 30 November 2020

Available online 3 December 2020

2213-1582/© 2020 The Author(s).

Published by Elsevier Inc.

This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

from 242 periventricular lesions in 47 MS patients. Pair-wise permutation t-tests were used to evaluate the differences between the AWs obtained for each quantitative measure.

Results: In the first experiment, we achieved a mean test AUC of 0.881 and the obtained AWs of FLAIR and the sum of AWs of Trace and ADC were 0.11 and 0.89, respectively, as expected based on previous knowledge. In the second experiment, we achieved a mean test F1 score of 0.895 and a mean AW of Trace = 0.49, of ADC = 0.28, and of T1w = 0.23, thereby confirming the findings of the first experiment. In the third experiment, MS lesion classification achieved test balanced accuracy = 0.777, sensitivity = 0.739, and specificity = 0.814. The mean AWs of T1map, MWF, NDI, and QSM were 0.29, 0.26, 0.24, and 0.22 ($p < 0.001$), respectively.

Conclusions: This work demonstrates that the proposed GAMER MRI might be a useful method to assess the relative importance of MRI measures in neurological diseases with focal pathology. Moreover, the obtained AWs may in fact help to choose the best combination of MR contrasts for a specific classification problem.

1. Introduction

Magnetic resonance imaging (MRI) has proven invaluable for the investigation of the pathophysiology of neurological diseases and guiding neurological diagnoses, prognoses, and evaluation of therapeutics. In fact, during the last decade, numerous fast MRI sequences and quantitative/semiquantitative MRI measures have been developed that provide complementary information to disentangle the pathological mechanisms and characteristics of brain diseases. In addition, specific biomarkers for diagnosis and response to therapy have been identified (Bozzali et al., 2016; González and Schwamm, 2016; Gupta et al., 2017). However, clinical research and practice are still limited by the time required to acquire multiple MR contrasts. It is imperative that these studies be conducted in a time frame compatible with patient tolerance, compliance and in the case of clinical practice, the requirements dictated by the healthcare system. Therefore, the need to address the selection of the most informative MR contrasts is pivotal to avoid uncomfortably lengthy acquisitions, to lower the subsequent possibility of having motion artifacts, and to reduce the related cost.

Deep learning, especially convolutional neural networks (CNN), has proven promising in the segmentation of brain regions or lesions in MR images (Andermatt et al., 2018; Carass et al., 2017; Commowick et al., 2016; La Rosa et al., 2019; Wachinger et al., 2018), classification of brain diseases (Payan and Montana, 2015; Yoo et al., 2018), MR reconstruction (Akçakaya et al., 2019; Schlemper et al., 2018), and prediction of disease prognosis (Saha et al., 2020; Tousignant et al., 2019). The layer-wise neural network (NN) design can identify high-level hidden representations through iterative training, which are pivotal for a given classification task. Some of the deep learning designs specifically enhance the interpretability of the decision made by the NN, such as class activation maps (Selvaraju et al., 2016; Zhou et al., 2016) and Shapley Additive exPlanations (Lundberg and Lee, 2017). Nevertheless, these methods either give importance to the voxels in images or to post-hoc feature importance. On the contrary, the attention mechanism within a NN provides attention weights (AWs) representing the importance of specific features. The concept, which originated in the field of natural language processing, can instruct the NN to attend to useful correlated elements in the text (Bahdanau et al., 2015). One of its variants, the gated attention mechanism, was extended to images and found to successfully assign AWs to non-overlapping patches from histopathological images in the classification of malignant cancer cells (Ilse et al., 2018; Tomczak et al., 2018).

In this work, we optimized the gated attention mechanism (Ilse et al., 2018) to develop a prototype of a Gated-Attention MEchanism Ranking of multi-contrast MRI in brain pathology (GAMER MRI). GAMER MRI specifically ranks the relative importance of global multi-contrast features, instead of the importance of local single-contrast patches, in the classification of focal lesions. This method was first validated for a clinical application where some MR-measure importance is known (e.g. ischemic stroke) and was then applied to the classification of specific subtypes of MS lesions, which are known to differ for the extent of myelin/axon damage and reparative capacity: this provided new knowledge about which MRI measure – among those sensitive to axon

and myelin integrity – is most suitable to distinguish lesions with different axon/myelin damage and repair in MS.

2. Materials and methods

2.1. MRI data

2.1.1. Stroke data

A total of 7,134 1.5 T and 3 T brain MRI studies obtained from a combination of inpatient and outpatient scanners at the Mount Sinai Hospital, New York, USA were randomly selected as the dataset. These imaging data were accumulated from the Mount Sinai BioMedical Engineering and Imaging Institute's HIPAA compliant Imaging Research Warehouse, including data from 10 scanners produced by two manufacturers (GE and Siemens Healthineers). The dataset consisted of various clinical acquisitions and included healthy controls, patients with subacute and acute infarct stroke, and patients with subacute and acute hemorrhage and mass effect. Among these patients, 904 are subacute and acute infarct stroke patients (defined as group 1) and 6,230 are healthy controls and other patients (defined as group 2). The 2D axial protocol included conventional, isotropically weighted Diffusion Weighted Imaging (DWI), Fluid-Attenuated Inversion Recovery (FLAIR), and T1-weighted images (T1w) from the inversion recovery pulse sequence. The most important mean acquisition parameters are listed in Table 1. Trace-weighted contrast (Trace) and apparent-diffusion coefficient (ADC) were reconstructed on the scanner from DWI.

Acute infarct stroke has distinctive representations on the acquired contrasts (Fig. 1). In the acute phase, hyperintensity is seen on Trace while ADC appears hypointense (Allen et al., 2012). In the subacute infarct stroke phase, both contrasts develop towards pseudo-normality. The segmentation of acute stroke lesions was performed on Trace and ADC by an expert radiologist consulting for Siemens Healthineers.

2.1.2. Multiple sclerosis data

Forty-seven MS patients (33 relapsing-remitting and 14 progressive, 31 females and 16 males, age range = 43.6 ± 14.4 years) were enrolled in the study approved by the local ethics committee of Basel University Hospital. Written consent was obtained prior to the MRI acquisition. Patients underwent a multi-parametric protocol on a 3 T Siemens Healthineers MAGNETOM Prisma MRI system. The 3D protocol included SPACE-based FLAIR, Magnetization-Prepared 2 Rapid Gradient Echoes (MP2RAGE) (Kober et al., 2012; Marques et al., 2010), Fast Acquisition with Spiral Trajectory and T2prep sequence (FAST-T2)

Table 1

Acquisition parameters of each contrast in the stroke dataset. TE: echo time; TR: repetition time; TI: inversion time; FOV: field of view; SR: spatial resolution.

	TE (ms)	TR (ms)	FOV (mm ³)	SR (mm ³)	TI (ms)	b values (s/mm ²)
FLAIR	94	8000	230x230x160	0.72x0.72x5	2460	–
T1w	6.9	2876	179x220x160	0.69x0.69x5	840	–
DWI	113.8	7625	240x240x170	1.02x1.02x5	–	0,1000

(Nguyen et al., 2016), multi-shell Diffusion Weighted Imaging (mDWI), and Multi-Echo Gradient Recalled Echo MRI (ME-GRE) (Wang and Liu, 2015). The most important acquisition parameters are shown in Table 2.

From multi-parametric MRIs, quantitative MR maps (qMRs) were further reconstructed. Quantitative T1 relaxometry map (qT1) was reconstructed from MP2RAGE as in (Kober et al., 2012). Myelin water fraction map (MWF) was reconstructed from FAST-T2 as in (Nguyen et al., 2016). Neurite density index (NDI) from the neurite orientation dispersion and density imaging model (Zhang et al., 2012) was reconstructed from DWI as in (Daducci et al., 2015). Quantitative Susceptibility Map (QSM) was reconstructed from ME-GRE as in (Wang and Liu, 2015). Co-registration between images was performed using FMRIB Software Library (FSL) (Jenkinson et al., 2012) and FreeSurfer (Fischl et al., 2001), and the obtained transformation matrices were later used for finding the correspondence of MS lesions between different qMRs. qMRs were not resampled to the same resolution so that the effect of interpolation in the resampling would not confound the quantitative values reflective of physical characteristics. MS lesions in white matter (WM) show hyperintensities on FLAIR images and in grey matter (GM) are blackholes on MP2RAGE images of the uniform contrast in Fig. 2a. WM lesions were automatically segmented (La Rosa et al., 2020) and manually corrected by two expert raters. Juxtacortical lesions (JCLs) and periventricular lesions (PVLs) were defined as WM lesions located within 3 mm of the boundary between (i) WM and GM and (ii) WM and ventricles, respectively, in Fig. 2b. The aforementioned boundaries were obtained through FreeSurfer processing on MP2RAGE (Fujimoto et al., 2014). In the end, 750 JCLs and 242 PVLs were found with a class-imbalance ratio of 1:3.

2.1.3. Study summary

Fig. 3 summarizes the information about the two datasets and the training, validation and test datasets in the following experiments.

2.2. Gamer MRI

The original gated attention mechanism proposed by Ilse et al. (Ilse et al., 2018) exploits the hidden representations of single-contrast patches to compute the corresponding AWs, which represents the

relative importance among the hidden representations in the classification. The main theorem behind this rationale is the following (Zaheer et al., 2017):

Theorem 1.. A prediction function $f(X)$ for a set of countable elements X is invariant to the permutation of the elements in X , if and only if, for suitable transformations g and h , $f(X)$ can be decomposed as:

$$f(X) = h\left(\sum_{x \in X} g(x)\right) \quad (1)$$

$g()$ and $h()$ were modeled by a NN. Based on (1), the gated attention mechanism is formulated as follows:

$$n = \sum_{l=1}^L a_l m_l = \sum_{x \in X} g(x) \quad (2)$$

$$a_l = \text{softmax}(w^T (\tanh(Um_l) \odot \text{sigm}(Vm_l))) \quad (3)$$

where m_l is the hidden representation of the l^{th} instance, a_l is the AW of the l^{th} instance, U and $V \in R^{K \times M}$ are weights of the fully connected layers (FCs) following the hidden representations, sigm stands for the non-linear sigmoid function, \odot is the element-wise multiplication operator, $w \in R^{1 \times K}$ is the weights of a FC, softmax stands for the softmax function.

Contrary to the original single-contrast approach to model $g()$, GAMER MRI adopted the multi-contrast multi-path approach on different MR contrasts and (2) becomes:

$$\sum_{l=1}^L a_l m_l = \sum_{l=1}^L a_l q_l(x_l) = \sum_{x \in X} g(x) \quad (4)$$

where $q(x)$ is the encoding function of the NN and Equation (3) remains the same. It is a simple variant to extend the meaning of AWs to the assessment of the importance of the MR contrasts in studying diseases and the parallel encoding paths enable the MR contrasts to be ranked by AWs. The core implementation of the gated attention mechanism in the NN was the same as in (Ilse et al., 2018) and formed by a FC followed by the hyperbolic tangent function (the attention layer) and a FC followed by the sigmoid function (the gate layer). The outputs of the attention




Fig. 1. Examples of Trace, ADC, FLAIR and T1w images in the stroke dataset. The lesion is hyperintense on Trace but hypointense on ADC (Allen et al., 2012). On T1w, the lesion is isointense than ADC and is faintly hyperintense on FLAIR.

Table 2

Acquisition parameters of each contrast in the MS dataset. TE: echo time; TR: repetition time; TI: inversion time; FOV: field of view; SR: spatial resolution.

	TE (ms)	TR (ms)	FOV (mm ³)	SR (mm ³)	TI (ms)	Additional Parameters
FLAIR	386	5000	256x256x256	1x1x1	1800	–
MP2RAGE	3	5000	256x256x256	1x1x1	700, 2500	–
ME-GRE	6.7,10.8,14.8,18.9, 22.9,27,31.1, 35.1,39.2,43.2	49	195x240x180	0.75x0.75x3	–	–
FAST-T2	0.5	7.5	240x240x160	1.25x1.25x5	–	T2prep times (ms) 0 (T2prep turned off), 7.5,17.5,67.5,147.5,307.5
mDWI	75	4500	256x256x144	1.8x1.8x1.8	–	b values (s/mm²) 0/12 acquisitions; 700/1000;2000;3000/137 directions in total

layer and the gate layer were element-wise multiplied and connected to a one-neuron FC and the softmax function to generate the normalized AWs. The number of neurons in the attention and gate layers depend on the experiment.

In order to validate our method and rank the importance of MRI features, the following three experiments were conducted: 1. volume-based classification of acute/subacute ischemic stroke vs other strokes and healthy controls; 2. patch-based classification of acute ischemic stroke lesions vs healthy tissue and 3. patch-based classification of JCLs vs PVLs in MS patients.

2.2.1. Pretrained network with GAMER MRI on stroke

To assess the performance of GAMER MRI as a ranking method, we combined GAMER MRI with the feature extracting compartment of an in-house pretrained NN from Siemens Healthineers (Princeton, NJ, USA), for the classification of acute/subacute ischemic stroke vs other patients and healthy controls using volumetric Trace, ADC, and FLAIR. The pretrained NN was trained for the same classification and thus learned how to encode relevant hidden features from Trace, ADC, and FLAIR.

2.2.1.1. Inputs and preprocessing. Trace, ADC, and FLAIR images were considered in this experiment since these contrasts were used for training the pretrained network. Subacute and acute infarct stroke patients were categorized into group 1, while group 2 included other patients and healthy controls. There were 5,002 subjects (group 1: 632 and group 2: 4,370) in the training dataset. The validation dataset had 1,061 subjects (group 1: 141 and group 2: 920) and 1,071 subjects (group 1: 131 and group 2: 940) were in the test dataset.

2.2.1.2. Architecture. The combined NN was built with three main compartments, including the feature extracting compartment of the pretrained NN, GAMER MRI and a classifier, as depicted in Fig. 4. The feature extracting compartment was, for each contrast, composed of two 3D convolutional blocks followed by two dense blocks based on the concept of DenseNet in (Huang et al., 2017). Each convolutional block consisted of a batch normalization layer, leaky ReLU units and a 3D convolutional layer. In each dense block, there were two 3D convolutional blocks with the kernel size of 3x3x3 and 1x1x1. The number of initial features was 16 and the growth rate was 2. The hidden feature vectors from all contrasts were then concatenated as the input to the following GAMER MRI so that the hidden feature vector of each contrast was encoded independently prior to the computation of AWs. In the GAMER MRI, the number of neurons each in the attention layer and in the gate layer was 400. The classifier was one sigmoid neuron receiving the weighted sum of the hidden features and the AWs. The importance of each contrast is represented by the AW.

2.2.1.3. Training Strategy. The combined NN was trained with a cross-entropy loss function and mini-batches. The weighted sampler was used to account for the class imbalance during training. The network

parameters, including the pretrained layers, were updated by the Adam optimizer with decoupled weight decay (AdamW) (Loshchilov and Hutter, 2019). The evaluation metric was the area under the receiver operating characteristic curve (AUC), which was the same metric used in training the pretrained network. To avoid overfitting, data augmentation was independently performed for each contrast on-the-fly. Since there is an inherent randomization in the initialization of network parameters and the split of mini-batches, the assessment of the effect of the random initialization is needed to properly describe the behavior of repeatability. The training, validation, and test datasets were kept the same during training, but the random seed changed in each repetition in the repeatability experiment. The leave-one-out (LOO) experiment on the selection of sequences was also conducted to characterize the method from a different perspective, namely by measuring the drop in the evaluation metrics reflecting the impact of the missing channel.

2.2.2. Patch-based network with GAMER MRI on stroke

The second experiment was performed to assess the ability of the GAMER MRI in a neural network when it was trained from scratch on the stroke dataset. We hypothesize that if GAMER MRI can provide the weights reflective of the current clinical understanding in the classification of acute infarct stroke lesions versus healthy tissues, it can be used in disease studies where the relative importance of MR contrasts is still unknown.

2.2.2.1. Inputs and preprocessing. In consideration of the limited number of existing acute infarct stroke lesion masks and in order to remove the effects of different scanners, Trace, ADC and T1w from 101 acute infarct patients without other pathologies, like hemorrhage, and 237 healthy controls were selected from the stroke dataset for the patch-based experiment. T1w was registered to b0 of DWI because the right correspondence between contrasts is essential to patch sampling. Because acute infarct lesions are of varying sizes, care must be taken when choosing the sampled patch size. Too large of a patch size is detrimental to small lesions. On the other hand, too small of a patch size would under-represent large lesions. Thus, after inspecting a subset of acute infarct stroke images, 24x24 voxels was empirically chosen for 2D patches. For healthy controls, the patches were randomly upsampled three times within the brains so that the healthy brains would not be under-represented by a small number of patches. In the end, 3,355 lesion patches and 9,917 healthy patches were sampled. Patches were divided into training, validation, and test datasets according to the ratios: 0.6, 0.3, and 0.1. As a result, there were 7,234 patches (2,001 lesion patches and 5,233 healthy patches) in the training dataset; 4,531 patches (1,012 lesion patches and 3,519 healthy patches) in the validation dataset; 1,507 patches (342 lesion patches and 1,165 healthy patches) in the test dataset. The patches containing acute infarct stroke lesions were given the label = 1.

2.2.2.2. Architecture. A patch-based multi-contrast CNN with GAMER MRI (NN2) could be decomposed to three compartments as the NN in




Fig. 2. MS lesions and qMRs. In (a), on MP2RAGE, the MS lesion in GM is a black hole (red arrow) and on FLAIR, the MS lesion in WM is hyperintense (green arrow). In (b), qT1, NDI, MWF and QSM reflective of different aspects of the microenvironment illustrate various representations of lesions. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2.2.1.2 (Fig. 5). The feature extracting compartment included three convolutional blocks for each MR contrast. Each convolutional block included a convolutional layer of 128 filters, exponential leaky units, and a batch normalization layer. The number of filters was chosen based on the evaluation metrics without inspecting the AWs prior to the 100-time repetitions. The three connected convolutional blocks were followed by a FC of 128 neurons encoding the hidden feature vector for each contrast. In the GAMER MRI, the number of neurons in the

attention layer and in the gate layer were both 64. The classifier was the same as in 2.2.1.2.

2.2.2.3. *Training strategy.* The NN was trained with a weighted cross-entropy loss function to account for the effect of class imbalance. The mini-batch size was 128 for both training and evaluation. The optimizer was Adam (Kingma and Ba, 2015). The F1 score was chosen as the evaluation metric because the correct identification for positive cases, i.

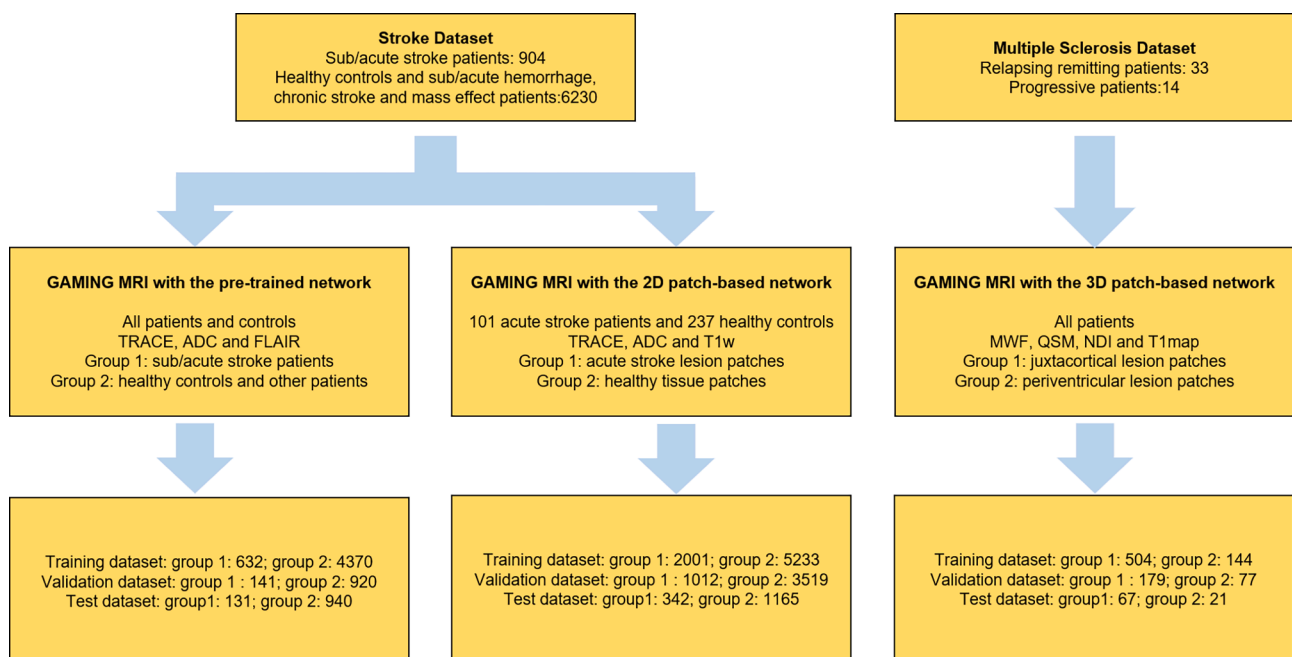


Fig. 3. Study summary.

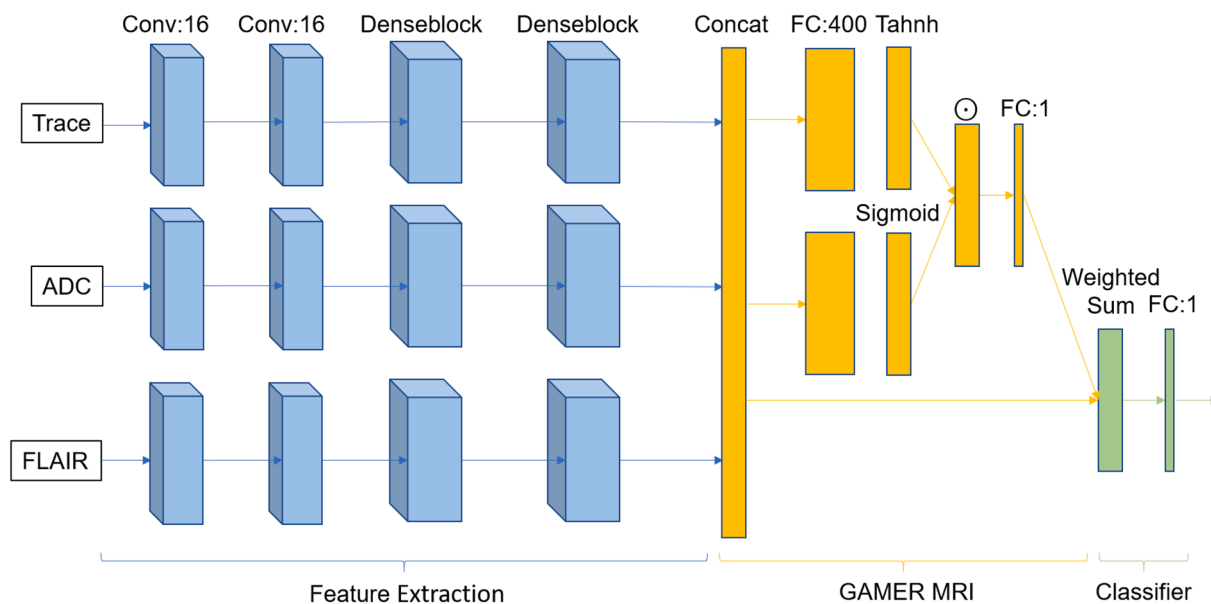


Fig. 4. The network structure consists of the pretrained feature extraction, the GAMER MRI and the classifier. Conv stands for the 3D convolutional block. FC is the fully connected layer. Concat is the concatenating layer. \odot represents an element-wise multiplication.

e. acute infarct stroke, was more important than healthy tissue. To avoid overfitting, data augmentation and early stopping was performed.

To appropriately characterize the performance, in addition to the aforementioned training strategy on the different random initializations, the next level of assessment was to split training, validation, and test datasets differently in different repetitions to make sure the power of the method does not come from the split.

2.2.3. Patch-based network with GAMER MRI on MS lesions

The third experiment was the evaluation of the mechanism on the classification of JCLs and PVLs from the MS dataset using qMRs, where the relative importance is unknown in the clinic.

2.2.3.1. *Inputs and preprocessing.* 3D patches close to $5 \times 5 \times 5 \text{ mm}^3$ were chosen as samples for training the neural network for the following three reasons: JCLs and PVLs are defined within 3 mm regions, the minimal slice thickness of qMRs is 5 mm, and various resolutions. This led to different patch sizes for each qMR to avoid confounding of the quantitative values by the interpolation in the registration. Considering the defined JCLs and PVLs being in the WM, each qMR was masked by the WM mask. Lesion patches were divided into training, validation, and test datasets following the ratios: 0.6, 0.3 and 0.1. Therefore, there were 648 lesion patches (504 JCLs and 144 PVLs) in the training dataset, 256 lesion patches (179 JCLs and 77 PVLs) in the validation dataset, and 88 lesion patches (67 JCLs and 21 PVLs) in the test dataset.

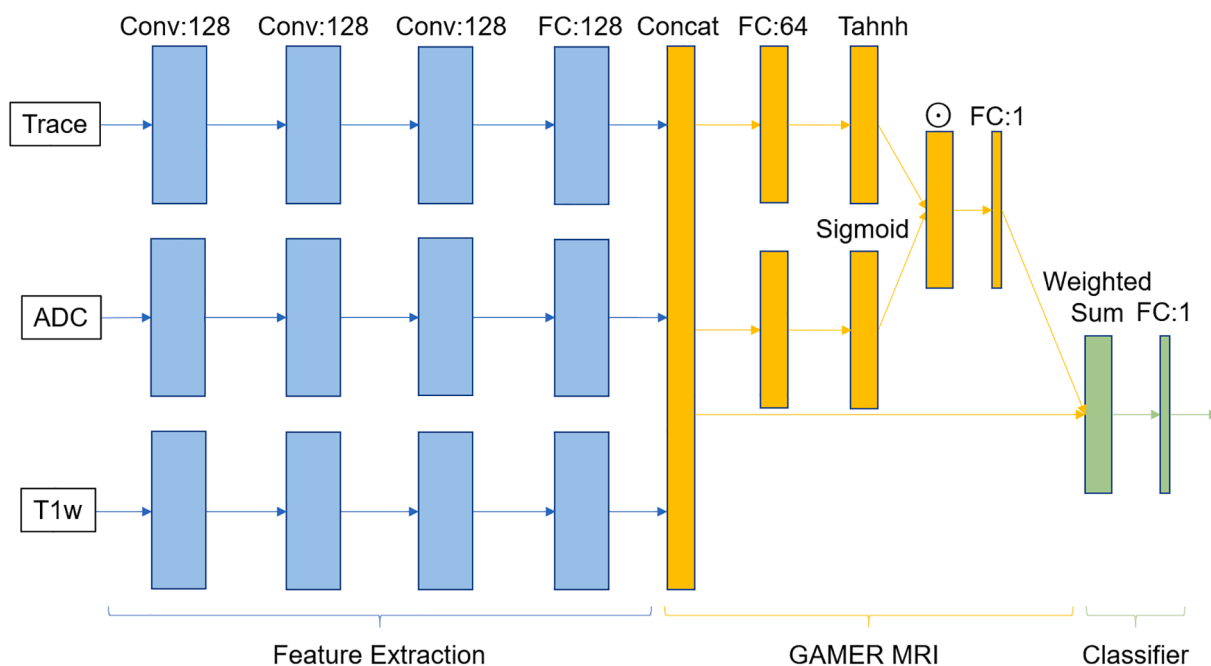


Fig. 5. The network structure is composed of the feature extraction, the GAMER MRI and the classifier. Conv stands for a convolutional block of 2D convolutional filters. FC is the fully connected layer. Concat is the concatenating layer. \odot represents the element-wise multiplication.

2.2.3.2. Architecture. A patch-based multi-contrast CNN with a GAMER MRI similar to the NN in 2.2.2.2, was built (Fig. 6). The feature extraction compartment included two convolutional blocks followed by a FC, as in 2.2.2.2, for each qMR. The convolutional layer in the convolutional block had 32 filters and the FC has 16 neurons encoding the hidden feature vector for each qMR. The criterion to choose the number of filters was the same as in 2.2.2.2. The hidden feature vectors from all qMRs were then concatenated as the input to the following GAMER MRI. In GAMER MRI, the number of neurons in the attention layer and gate layer were both 32. The classifier was the same as in 2.2.2.2.

2.2.3.3. Training strategy. The loss function and the mini-batch size, data augmentation, early stopping and the learning-rate-reduce-plateau scheduler were the same as in 2.2.2.3. The optimizer and the evaluation metric were the same as in 2.2.1.3.

In addition to the characteristics evaluated in the previous two experiments, resampling patches prior to the split of datasets was performed. To avoid sampling bias in the patch-based classification, randomly resampling patches is pivotal for reproducibility. The pairwise one-sided 10,000 permutation t-tests were performed on the obtained orders of AW of all repetitions and the multiple comparison problem was tackled by Bonferroni correction. See the [supplementary data](#) for further details.

2.3. Data and code availability statement

The datasets, provided by the Mount Sinai Hospital and Basel University Hospital, used in this study are not publicly available because the IRB of the study limits access to the data. The code used for training the models has dependencies on Siemens' internal tooling, infrastructure and hardware, and its release is therefore not feasible. However, the architecture, layer details and hyperparameters are described in

sufficient details in the manuscript to support replication with non-proprietary libraries.

3. Results

3.1. Pretrained network with GAMER MRI on the stroke dataset

Validation and test results of the NN in 2.2.1 using three different random seeds for the random sampler, which led to different initializations and split of mini-batches, are given in Table 3. In each repetition, the mean AW (mAW) were averaged over the AWs of the corrected predicted samples. The reported mean AWs (rmAWs) were the average of all mAW across repetitions.

The LOO experiment was conducted twice for each pair combination of Trace, ADC, and

FLAIR. The drops in validation AUC were averaged across the repetitions and compared between combinations in Table 4.

Table 3

Pretrained network with GAMER MRI on the stroke dataset: Mean validation and test results over three repetitions. The mean area under the curve (AUC) is reported to show the classification performance and the sum of reported mean attention weights (rmAWs) of Trace and ADC and the rmAW of FLAIR are shown to provide the importance ranking of the MRI metrics.

Pretrained-network with GAMER MRI in stroke		
Dataset	Validation	Test
AUC	0.919	0.881
AWs	Trace + ADC, FLAIR (0.890, 0.110)	Trace + ADC, FLAIR (0.886, 0.114)

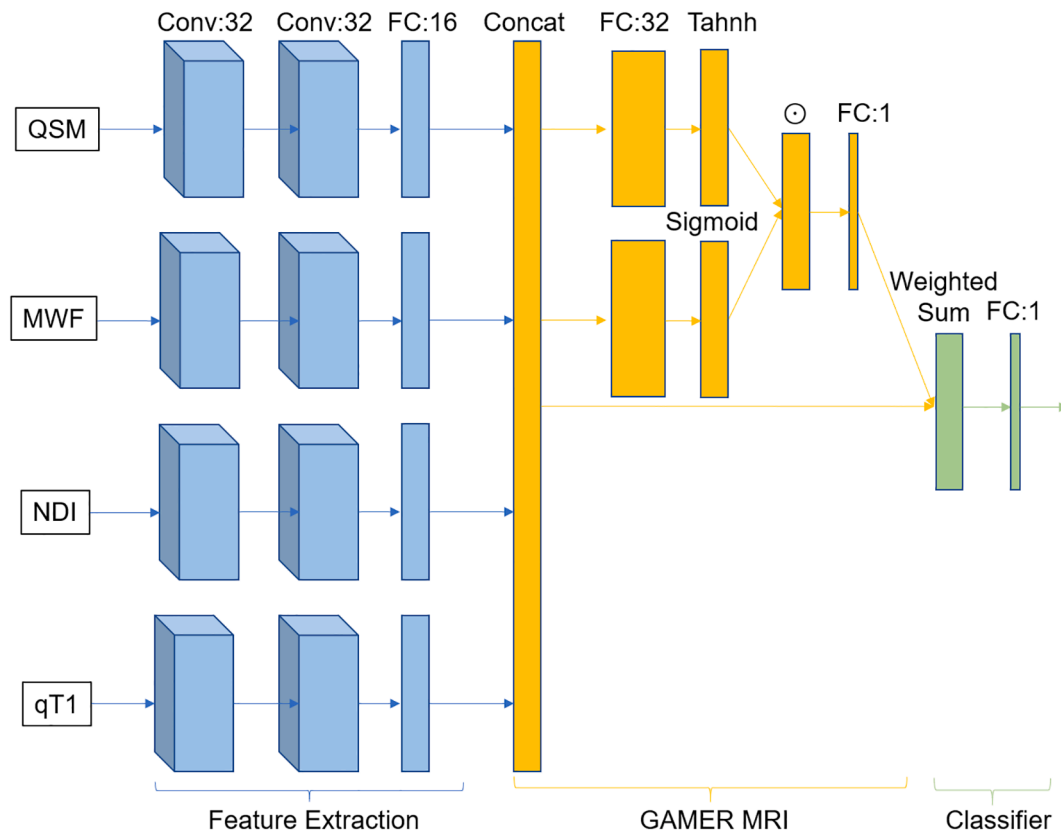


Fig. 6. The network structure includes feature extraction, GAMER MRI and classifier. Conv stands for a convolutional block of 3D convolutional filters. FC is the fully connected layer. Concat is the concatenating layer. \odot represents the element-wise multiplication.

Table 4

The validation result of the leave-one-out experiment. The averaged AUC, the drop in performance, and the rmAWs are reported. For the experiment using only Trace and ADC, the rmAWs varied greatly in the repetitions. The rmAWs in the first repetition were (Trace, ADC)=(0.871, 0.129) and in the second repetition were (0.138, 0.861).

Leave-one-out Experiment			
Input Contrasts	ADC, FLAIR	Trace, FLAIR	Trace, ADC
AUC	0.866	0.914	0.91
Drop in AUC	0.052	0.004	0.008
AWs	ADC, FLAIR (0.843, 0.157)	Trace, FLAIR (0.866, 0.134)	Trace, ADC (0.505*, 0.495*)

3.2. Patch-based network with GAMER MRI on the stroke dataset

The NN in 2.2.2 was evaluated 100 times using the same training, validation and test datasets with different random initialization of the network and mini-batches. Furthermore, the NN was evaluated 100 times using different splits of training, validation, and test datasets; respective validation results are reported in Table 5.

3.3. Patch-based network with GAMER MRI on MS lesions

The NN in 2.2.3 was trained 100 times on the resampled datasets to ensure reproducibility of the method for classification of MS lesions where the important order of sequences is unknown. The performance of the repetition experiment on the validation and test datasets is reported in Table 6. In Table 7, we report the mean and standard deviation of the mean AWs, which is defined as in 3.1. Furthermore, the results of permutation *t*-test on the obtained order of AWs are reported.

4. Discussion

We developed a Gated-Attention MEchanism Ranking of multi-contrast MRI in brain pathology (GAMER MRI) and demonstrated its ability to rank the relative importance of MRI contrasts / qMRs in the three different classification scenarios including the differentiation of well-studied infarct strokes and that of less understood MS lesions.

4.1. Pretrained network with GAMER MRI on stroke

To accomplish the classification task, the NN should be able to extract unique and common information from the input contrasts. We demonstrated in 3.1 that GAMER MRI could utilize the unique and common information from each contrast to provide the AW as a proxy of the importance of each contrast. The mean AUC in this experiment was comparable to the performance of the original pre-trained network in a similar classification task. In addition, the mean AUC of validation and test datasets (Table 3) indicated that the combination of a pretrained encoder and GAMER MRI well performed. Because the AWs of the correctly classified samples formed the correct pattern with the hidden features for the classifier to make the right decision, we then proceeded to average those AWs to obtain the mAW for each repetition: in fact, considering the AWs of the incorrectly classified samples would not have reflected their real importance in the identification of stroke lesions. The consistent ratio between the sum of rmAWs of Trace and ADC and the rmAW of FLAIR showed that FLAIR was less important compared to the other two contrasts in the given classification task. This is in line with the relative clinical importance of these contrasts for the diagnosis of acute and subacute infarct stroke (González and Schwamm, 2016).

We observed an inconsistent ratio between the rmAWs of Trace and that of ADC, which is probably due to the strong correlation between the contrasts. Because of the known evolvement of the infarct stroke representation from the acute to the subacute stage on Trace and ADC, the representations become pseudonormal and similar. This leads to

Table 5 The validation and test result of the patch-based acute stroke classification. The averaged F1 score and the rmAWs and the standard deviation of mAWs across 100 repetitions are reported. The order in repetitions obtained by 100 repetitions shows the number of the corresponding order of attention weights from the largest to the smallest.

Patch-based network with GAMER MRI in stroke			
Dataset	Validation		
Training Strategy	Different Random Initialization		
F1 score	0.917	0.895	
AWs	Trace 0.5 ± 0.043	ADC 0.293 ± 0.036	TIw 0.207 ± 0.028
Order in Repetitions	Trace, ADC and TIw 98	Trace, TIw and ADC 2	Other combinations —
AWs	Trace 0.281 ± 0.039	ADC 0.278 ± 0.040	TIw 0.229 ± 0.030
Order in Repetitions	Trace, TIw and ADC 82	Trace, TIw and ADC 18	Other combinations —
Dataset	Test		
Training Strategy	Different Random Initialization		
F1 score	0.885	0.895	
AWs	Trace 0.507 ± 0.043	ADC 0.284 ± 0.035	TIw 0.208 ± 0.028
Order in Repetitions	Trace, ADC and TIw 94	Trace, TIw and ADC 6	Other combinations —
AWs	Trace 0.278 ± 0.040	ADC 0.278 ± 0.040	TIw 0.228 ± 0.029
Order in Repetitions	Trace, ADC and TIw 80	Trace, TIw and ADC 19	Other combinations 1

Table 6

Performance of the patch-based network on MS lesions. The average mean and standard deviation of the metrics as a percentage of 100 repetitions are reported.

Mean metrics (%)	Balanced Accuracy	Sensitivity	Specificity	F1 score
Validation dataset	78.34 ± 3.09	74.21 ± 7.86	82.49 ± 5.12	64.66 ± 4.00
Test dataset	77.65 ± 5.49	73.90 ± 10.10	81.41 ± 8.21	65.09 ± 8.93

Table 7

The rmAWs, the standard deviation of mAWs and the statistical test on the pairwise comparison. The upper section shows the results of the validation dataset and in the lower section are the results of the test dataset. ***: corrected $p < 0.001$.

Validation				
Contrast	AW	Statistical test	P value	Significance
qT1	0.285 ± 0.027	qT1 > MWF	0.0001	***
MWF	0.256 ± 0.015	MWF > NDI	0.0001	***
NDI	0.241 ± 0.014	NDI > QSM	0.0001	***
QSM	0.218 ± 0.022	—	—	—
Test				
Contrast	AW	Statistical test	P value	Significance
qT1	0.284 ± 0.030	qT1 > MWF	0.0001	***
MWF	0.256 ± 0.016	MWF > NDI	0.0001	***
NDI	0.241 ± 0.021	NDI > QSM	0.0001	***
QSM	0.218 ± 0.023	—	—	—

stronger correlation between the information brought by Trace and ADC in addition to the intrinsic physical correlation between these two contrasts.

In Table 4, the obvious drop in performance of the LOO experiment when Trace was excluded indicates that Trace provides more unique information than ADC and FLAIR. Indeed, in the two LOO experiments including FLAIR, the rmAWs of FLAIR were smaller suggesting its relative lower importance in this classification, which echoes the result in Table 3. In the experiment performed without using FLAIR images, the inconsistent ratio of AWs between Trace and ADC bolsters the implication of confounds caused by the strong correlation between these two contrasts. One main assumption behind the previous interpretation is that the amount of extractable information is the same across the LOO experiments. The comparison of the obtained results with the current clinical understanding of stroke lesions indicates this assumption is valid.

Considering the results in 3.1, we obtained empirical evidence that the proposed variant of the attention mechanism can provide AWs representative of the importance of non-correlated contrasts.

4.2. Patch-based model with attention mechanism

We then aimed to assess whether GAMER MRI combined with a NN, which is trained from scratch for the classification of acute infarct stroke, is able to provide AWs representative of the known relative importance of MR contrasts in clinical practice.

The obtained F1 scores on both validation and test datasets (Table 5) indicated that the NN combined with GAMER MRI learned information necessary to classify the patch containing acute infarct stroke lesions regardless of being trained with different random initialization or different split of training, validation, and test datasets.

In both repetition experiments, the rmAWs of Trace, ADC, and T1w showed that Trace carried more characteristic information than ADC followed by T1w images, which is in accord with the clinical understanding that on Trace, acute infarct stroke is hyperintense and more obvious than on ADC, and that the light hypointense appearance of an acute stroke in T1w images is less evident than the drop in signal often

observed in ADC maps (González and Schwamm, 2016). In addition, while the number of repetitions slightly varied across the different experimental conditions, the importance order of Trace, ADC and T1w remained consistent (Table 5). This demonstrates that the NN was able to extract relevant information in the given samples. The order of importance of Trace, T1w, and ADC might be a result of the patches having similar information on Trace and ADC. As a result, the NN would learn less unique information from ADC leading to its lower mAW than the one of Trace.

The consistent results in 3.2 validated the assumption that the AWs obtained with GAMER MRI can be used to assess the relative importance of MRI contrasts without the restriction on an informative pre-trained NN. Also considering the results in 3.1, which demonstrated that GAMER MRI could obtain a clinically meaningful ranking of MRI contrasts, the method may be well applicable to neurological diseases that are less understood.

4.3. MS patch-based model with attention mechanism

In this last scenario, we aimed at assessing if the GAMER MRI could be applied to other MRI measures and diseases, where the relative importance of measures is less understood. Therefore, we studied whether the GAMER MRI could rank myelin/axonal sensitive measures such as qT1, MWF, NDI, and QSM to classify lesions that are known to have different myelin and axonal content, such as lesions located near to the ventricles (PVL: lower myelin and axonal content) and next to the cortex (JCL: relatively higher myelin and axonal content) (Goldschmidt et al., 2009; Tonietto, 2018).

For both the validation and the test datasets, the network exhibited a moderate performance (Table 6): balanced accuracy was ca 78% - with a specificity that was slightly higher than the sensitivity (74% vs 82%), and the F1 score was ca 65%. In this experiment, different than in the previous one, we have assessed the network performance by using other summary measures than the F1 score: this is essentially because the F1 score does not consider true negative results, hence it may not equally consider lesions, whose characteristics are not completely understood (i. e. JCLs and PVLs). The multiple statistical tests on pairwise rmAWs showed that the metric best discriminating PVL vs JCL microstructure is qT1 followed by MWF, NDI, and QSM. qT1 quantifies the overall microstructural tissue damage within MS lesions (Bonnier et al., 2014), whereas MWF and NDI provide specific information about myelin and axonal content (Nguyen et al., 2016; Zhang et al., 2012). The order of importance reflects the overall difference in myelin/axonal content revealed in pathological studies (Goldschmidt et al., 2009), which qT1 depicts with the highest sensitivity. Hence, through this experiment, we could establish the reliability of GAMER MRI in a context where the relative contribution of MR measures to the discrimination of focal pathology is not clear.

Compared to the results obtained on the stroke dataset, the smaller differences between rmAWs of different qMRs might be caused by the smaller size of MS lesion datasets and/or higher similarities between lesion groups. A much larger effect is expected if an increased number of samples in datasets is included. Another potential underlying cause of this difference is the fact that the applied qMRs have in part redundant information. Indeed, the microstructural environment measured by qT1 encompasses the myelin content and neuro-axonal integrity measured by MWF and NDI. On the other hand, QSM measures both iron deposition and myelin properties since it is sensitive to susceptibility effect due to paramagnetic substances and to the orientation of myelin sheaths. Besides, it has to be considered that – different than the contrasts applied in stroke (e.g. Trace) – qMRs in the MS experiment could not sharply delineate the boundary of MS lesions, hereby reflecting the local variations surrounding the focal damage. Despite all this, however, GAMER-MRI still demonstrated a statistically significant difference between rmAWs of the qMRs.

4.4. Guideline on GAMER MRI

In consideration of the obtained results, we propose to use GAMER MRI as follows:

1. Train and evaluate the method multiple times to see if there is strong or mild correlation between the resultant AWs of input measures. If there is strong correlation, an ablation study should be performed to remove the correlated modality showing a smaller drop in performance. Train and evaluate the method on the remaining measures to obtain AWs.
2. If there is no strong correlation, the importance order based on the mean AWs across the repetitions is recommended.

4.5. Conclusion

Our work shows that GAMER MRI provides a clinically meaningful order of importance for MR-based features in the classification of infarct strokes. In addition, even though qMRs in the classification of JCLs and PVLs in MS had redundant information, GAMER MRI still managed to reveal a close but significant order of importance. Considering this importance order, it may be possible to reduce the number of input MRI measures while retaining most of the useful information.

The main disadvantage of this method is the need for multiple evaluations since the criteria on the so-called strong correlation is based on the AWs, not just on the input contrasts, as shown in the experiments of NN2. Future work will be required to remove this constraint. Furthermore, future work should center on combining the proxy quantification of the importance of qMRs with the values of qMRs to form meaningfully combined patterns for further studies since qMRs characterize different physical processes and physiological environments.

Declaration of competing interest

Part of the work was performed while Po-Jui Lu was doing internship in Siemens Healthineers, Princeton, USA. Youngjin Yoo, Pascal Ceccaldi and Eli Gibson are employed by Siemens Healthineers, Princeton, USA. Benjamin Odry is employed by Covera Health, New York, USA. Matthias Weigel has received research funding by Biogen for developing spinal cord MRI. Kambiz Nael has consulted for Olea Medical outside the scope of this work. Zahi Fayad has a research grant from Siemens Healthineers and is founder and board member of Trained Therapeutix Discovery.

CRediT authorship contribution statement

Po-Jui Lu: Conceptualization, Methodology, Investigation, Formal analysis, Writing - original draft. **Youngjin Yoo:** Methodology, Resources, Writing - review & editing. **Reza Rahmzadeh:** Data curation, Writing - review & editing. **Riccardo Galbusera:** Data curation, Writing - review & editing. **Matthias Weigel:** Resources, Data curation, Writing - review & editing. **Pascal Ceccaldi:** Resources, Writing - review & editing. **Thanh D. Nguyen:** Resources, Writing - review & editing. **Pascal Spincemaille:** Resources, Writing - review & editing. **Yi Wang:** Resources, Writing - review & editing. **Alessandro Daducci:** Resources, Writing - review & editing. **Francesco La Rosa:** Data curation, Writing - review & editing. **Meritxell Bach Cuadra:** Resources, Writing - review & editing. **Robin Sandkühler:** Conceptualization, Writing - review & editing. **Kambiz Nael:** Resources, Writing - review & editing. **Amish Doshi:** Resources, Writing - review & editing. **Zahi A. Fayad:** Resources, Writing - review & editing. **Jens Kuhle:** Writing - review & editing. **Ludwig Kappos:** Writing - review & editing. **Benjamin Odry:** Methodology, Writing - review & editing. **Philippe Cattin:** Supervision, Writing - review & editing. **Eli Gibson:** Supervision, Writing - review & editing, Methodology. **Cristina Granziera:** Supervision, Writing - review & editing, Conceptualization, Funding acquisition, Resources.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement and Disclaimer

We would like to acknowledge all the patients and healthy controls in this project. This project is supported by Swiss National Funds PZ00P3_154508, PZ00P3_131914 and PP00P3_176984. Francesco La Rosa and Meritxell Bach Cuadra are supported by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie project TRABIT (agreement No 765148) and by the Centre d'Imagerie BioMedicale (CIBM). We would also thank the Mount Sinai Hospital for providing the stroke dataset and Basel University Hospital for acquiring the MS dataset. The concepts and information presented in this paper are based on research results and GAMER MRI is not commercially available.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.nicl.2020.102522>.

References

- M. Akçakaya S. Moeller S. Weingärtner K. Uğurbil Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging 2019 Reson. Med Magn 10.1002/mrm.27420.
- Allen, L.M., Hasso, A.N., Handwerker, J., Farid, H., 2012. Sequence-specific MR imaging findings that are useful in dating ischemic stroke. Radiographics. <https://doi.org/10.1148/rg.325115760>.
- Andermatt, S., Pezold, S., Cattin, P.C., 2018. Automated Segmentation of Multiple Sclerosis Lesions Using Multi-dimensional Gated Recurrent Units. pp. 31–42. https://doi.org/10.1007/978-3-319-75238-9_3.
- Bahdanau, D., Cho, K.H., Bengio, Y., 2015. Neural machine translation by jointly learning to align and translate. in: 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings.
- Bonnier, G., Roche, A., Romascano, D., Simioni, S., Meskaldji, D., Rotzinger, D., Lin, Y.C., Menegaz, G., Schlupe, M., Du Pasquier, R., Sumpf, T.J., Frahm, J., Thiran, J.P., Krueger, G., Granziera, C., 2014. Advanced MRI unravels the nature of tissue alterations in early multiple sclerosis. Ann. Clin. Transl. Neurol. 1, 423–432. <https://doi.org/10.1002/acn3.68>.
- M. Bozzali L. Serra M. Cercignani Quantitative MRI to understand Alzheimer's disease pathophysiology 2016 Opin. Neurol Curr 10.1097/WCO.0000000000000345.
- Carass, A., Roy, S., Jog, A., Cuzzocreo, J.L., Magrath, E., Gherman, A., Button, J., Nguyen, J., Prados, F., Sudre, C.H., Jorge Cardoso, M., Cawley, N., Ciccarelli, O., Wheeler-Kingshott, C.A.M., Ourselin, S., Catanese, L., Deshpande, H., Maurel, P., Commowick, O., Barillot, C., Tomas-Fernandez, X., Warfield, S.K., Vaidya, S., Chunduru, A., Muthuganapathy, R., Krishnamurthi, G., Jesson, A., Arbel, T., Maier, O., Handels, H., Ithme, L.O., Unay, D., Jain, S., Sima, D.M., Smeets, D., Ghafoorian, M., Platel, B., Birenbaum, A., Greenspan, H., Bazin, P.L., Calabresi, P.A., Crainiceanu, C.M., Ellingsen, L.M., Reich, D.S., Prince, J.L., Pham, D.L., 2017. Longitudinal multiple sclerosis lesion segmentation: Resource and challenge. Neuroimage 148, 77–102. <https://doi.org/10.1016/j.neuroimage.2016.12.064>.
- Commowick, O., Cervenansky, F., Ameli, R., 2016. MSSEG Challenge Proceedings: Multiple Sclerosis Lesions Segmentation Challenge Using a Data Management and Processing Infrastructure. Miccai.
- Daducci, A., Canales-Rodríguez, E.J., Zhang, H., Dyrby, T.B., Alexander, D.C., Thiran, J.P., 2015. Accelerated Microstructure Imaging via Convex Optimization (AMICO) from diffusion MRI data. Neuroimage. <https://doi.org/10.1016/j.neuroimage.2014.10.026>.
- Fischl, B., Liu, A., Dale, A.M., 2001. Automated manifold surgery: Constructing geometrically accurate and topologically correct models of the human cerebral cortex. IEEE Trans. Med. Imaging. doi 10 (1109/42), 906426.
- Fujimoto, K., Polimeni, J.R., van der Kouwe, A.J.W., Reuter, M., Kober, T., Benner, T., Fischl, B., Wald, L.L., 2014. Quantitative comparison of cortical surface reconstructions from MP2RAGE and multi-echo MPRAGE data at 3 and 7T. Neuroimage 90, 60–73. <https://doi.org/10.1016/j.neuroimage.2013.12.012>.
- Goldschmidt, T., Antel, J., König, F.B., Brück, W., Kuhlmann, T., 2009. Remyelination capacity of the MS brain decreases with disease chronicity. Neurology 72, 1914–1921. <https://doi.org/10.1212/WNL.0b013e3181a8260a>.
- González, R.G., Schwamm, L.H., 2016. Imaging acute ischemic stroke. pp. 293–315. <https://doi.org/10.1016/B978-0-444-53485-9.00016-7>.
- Gupta, A., Al-Dasuqi, K., Xia, F., Askin, G., Zhao, Y., Delgado, D., Wang, Y., 2017. The use of noncontrast quantitative MRI to detect gadolinium-enhancing multiple sclerosis

- brain lesions: A systematic review and meta-analysis. *Am. J. Neuroradiol.* <https://doi.org/10.3174/ajnr.A5209>.
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks, in: Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017. <https://doi.org/10.1109/CVPR.2017.243>.
- Ilse, M., Tomczak, J.M., Welling, M., 2018. Attention-based Deep Multiple Instance Learning.
- Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. Review FSL. *Neuroimage.* <https://doi.org/10.1016/j.neuroimage.2011.09.015>.
- D.P. Kingma J.L. Ba Adam: A method for stochastic optimization in: 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings 2015.
- T. Kober C. Granziera D. Ribes P. Browaeys M. Schluep R. Meuli R. Frackowiak R. Gruetter G. Krueger MP2RAGE multiple sclerosis magnetic resonance imaging at 3 T 2012 *Radiol Invest* 10.1097/RLI.0b013e31824600e9.
- La Rosa, F., Fartaria, M.J., Kober, T., Richiardi, J., Granziera, C., Thiran, J.-P., Cuadra, M. B., 2019. Shallow vs Deep Learning Architectures for White Matter Lesion Segmentation in the Early Stages of Multiple Sclerosis. pp. 142–151. https://doi.org/10.1007/978-3-030-11723-8_14.
- La Rosa, Francesco, Abdulkadir, Ahmed, Fartaria, Mário João, Rahmzadeh, Reza, Lu, Po-Jui, Galbusera, Riccardo, Barakovic, Muhamed, Thiran, Jean-Philippe, Granziera, Cristina, Cuadra, Merixtell Bach, 2020. Multiple sclerosis cortical and WM lesion segmentation at 3T MRI: a deep learning method based on FLAIR and MP2RAGE. *NeuroImage* 27, 102335.
- Loshchilov, I., Hutter, F., 2019. Decoupled weight decay regularization, in: 7th International Conference on Learning Representations, ICLR 2019.
- S.M. Lundberg S.I. Lee A unified approach to interpreting model predictions in 2017 Advances in Neural Information Processing Systems.
- Marques, J.P., Kober, T., Krueger, G., van der Zwaag, W., Van de Moortele, P.F., Gruetter, R., 2010. MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *Neuroimage.* <https://doi.org/10.1016/j.neuroimage.2009.10.002>.
- Nguyen, T.D., Deh, K., Monohan, E., Pandya, S., Spincemaille, P., Raj, A., Wang, Y., Gauthier, S.A., 2016. Feasibility and reproducibility of whole brain myelin water mapping in 4 minutes using fast acquisition with spiral trajectory and adiabatic T2prep (FAST-T2) at 3T. *Magn. Reson. Med.* 76, 456–465. <https://doi.org/10.1002/mrm.25877>.
- Payan, A., Montana, G., 2015. Predicting Alzheimer's disease a neuroimaging study with 3D convolutional neural networks, in: ICPRAM 2015–4th International Conference on Pattern Recognition Applications and Methods, Proceedings.
- Saha, S., Pagnozzi, A., Bourgeat, P., George, J.M., Bradford, D.K., Colditz, P.B., Boyd, R. N., Rose, S.E., Frupp, J., Pannek, K., 2020. Predicting motor outcome in preterm infants from very early brain diffusion MRI using a deep learning convolutional neural network (CNN) model. *Neuroimage.* <https://doi.org/10.1016/j.neuroimage.2020.116807>.
- J. Schlemper J. Caballero J.V. Hajnal A.N. Price D. Rueckert A Deep Cascade of Convolutional Neural Networks for Dynamic MR Image Reconstruction 2018 Med. Imaging IEEE Trans 10.1109/TMI.2017.2760978.
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2016. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. *Int. J. Comput. Vis.* 128, 336–359. <https://doi.org/10.1007/s11263-019-01228-7>.
- Tomczak, J.M., Ilse, M., Welling, M., Jansen, M., Coleman, H.G., Lucas, M., de Laat, K., de Bruin, M., Marquering, H., van der Wel, M.J., de Boer, O.J., Heijink, C.D.S., Meijer, S.L., 2018. Histopathological classification of precursor lesions of esophageal adenocarcinoma: A Deep Multiple Instance Learning Approach. *Med. Imaging with Deep Learn.* 3–5.
- Toniello, M., 2018. Periventricular remyelination is associated with grey matter atrophy in MS. *ECTRIMS*.
- Tousignant, A., Paul LemaitreLemaitre, M., Doina Precup, C., Arnold, D.L., 2019. Prediction of Disease Progression in Multiple Sclerosis Patients using Deep Learning Analysis of MRI Data Tal Arbel 3, Proceedings of Machine Learning Research.
- Wachinger, C., Reuter, M., Klein, T., 2018. DeepNAT: Deep convolutional neural network for segmenting neuroanatomy. *Neuroimage* 170, 434–445. <https://doi.org/10.1016/j.neuroimage.2017.02.035>.
- Wang, Y., Liu, T., 2015. Quantitative susceptibility mapping (QSM): Decoding MRI data for a tissue magnetic biomarker. *Magn. Reson. Med.* 73, 82–101. <https://doi.org/10.1002/mrm.25358>.
- Yoo, Y., Tang, L.Y.W., Brosch, T., Li, D.K.B., Kolind, S., Vavasour, I., Rauscher, A., MacKay, A.L., Traboulsee, A., Tam, R.C., 2018. Deep learning of joint myelin and T1w MRI features in normal-appearing brain tissue to distinguish between multiple sclerosis patients and healthy controls. *NeuroImage Clin.* 17, 169–178. <https://doi.org/10.1016/j.nicl.2017.10.015>.
- M. Zaheer S. Kottur S. Ravanbakhsh B. Póczos R. Salakhutdinov A.J. Smola Deep sets in 2017 Advances in Neural Information Processing Systems.
- Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C., 2012. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. *Neuroimage* 61, 1000–1016. <https://doi.org/10.1016/j.neuroimage.2012.03.072>.
- B. Zhou A. Khosla A. Lapedriza A. Oliva A. Torralba Learning Deep Features for Discriminative Localization in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE 2016 2921 2929 10.1109/CVPR.2016.319.