

Distributionally Robust Image Classifiers for Stroke Diagnosis in Accelerated MRI*

Boran Hao¹[0000-0001-7922-0513], Guoyao Shen²[0000-0002-5479-8332],
Ruidi Chen^{**1}[0000-0002-1508-1742], Chad W. Farris³[0000-0002-1133-3834],
Stephan W. Anderson³[0000-0002-5367-7459], Xin Zhang²[0000-0002-4413-5084],
and Ioannis Ch. Paschalidis¹[0000-0002-3343-2913]

¹ Department of Electrical and Computer Engineering, Boston University, Boston MA 02215, USA

yannisp@bu.edu

² Department of Mechanical Engineering and the Photonics Center, Boston University, Boston MA 02215, USA

³ Boston Medical Center and Boston University Chobanian & Avedisian School of Medicine, Boston MA 02118, USA

Abstract. *Magnetic Resonance Imaging (MRI)* acceleration techniques using *k-space sub-sampling (KS)* can greatly improve the efficiency of MRI-based stroke diagnosis. Although *Deep Neural Networks (DNN)* have shown great potential on stroke lesion recognition tasks when the MR images are reconstructed from the full k-space, they are vulnerable to the lower quality MR images generated by KS. In this paper, we propose a *Distributionally Robust Learning (DRL)* approach to improve the performance of stroke recognition DNN models when the MR images are reconstructed from the sub-sampled k-space. For *Convolutional Neural Network (CNN)* and *Vision Transformer (ViT)*-based models, our methods improve the stroke classification AUROC and AUPRC by up to 11.91% and 9.32% on the KS-perturbed brain MR images, respectively, compared against *Empirical Risk Minimization (ERM)* and other baseline defensive methods. We further show that DRL models can successfully recognize the stroke cases from highly perturbed MR images where clinicians may fail, which provides a solution for improved diagnosis in an accelerated MRI setting.

Keywords: Stroke diagnosis · MRI acceleration · Distributionally robust optimization · Deep learning.

1 Introduction

Magnetic Resonance Imaging (MRI) has been extensively applied to clinical diagnosis [16]. Compared with *Computed Tomography (CT)*, a brain MRI is more

* This work was supported by the Rajen Kilachand Fund for Integrated Life Science and Engineering and by the NSF under grant CCF-2200052 and IIS-1914792.

** Ruidi Chen contributed to this work while at Boston University, before moving to her current position at Amazon SCOT.

sensitive for multiple stroke types [3], therefore considered as the gold standard for stroke diagnosis. Nevertheless, the long acquisition time for a brain MRI (20 to 30 minutes) imposes challenges, especially in cases of acute stroke where rapid diagnosis is essential and patient movement during this distressing period of time commonly limits evaluation. As a result, MRI acceleration techniques have been developed to achieve more rapid diagnosis, increasing resource availability while reducing costs [14, 18]. A *k-space sub-sampling (KS)* approach serves as a simple MRI acceleration solution [20], compared with other hardware-based acceleration methods. However, the signal loss by KS leads to blurry reconstructed MR images that are less than ideal for a reliable clinical diagnosis.

Artificial Intelligence (AI) plays an increasingly important role in MRI-based diagnosis, for both MR image reconstruction and clinical decision making. *Deep Neural Networks (DNN)* were trained to reconstruct the MR images from the sub-sampled k-space [10, 13], which provides a better reconstruction than the *Inverse Fast Fourier Transform (IFFT)*. Nevertheless, detailed information in the brain may still be lost in the reconstructed MR images due to the signal sparsity in the k-space. On the other hand, traditional *Convolutional Neural Network (CNN)* [15] and the latest *Vision Transformer (ViT)*-based [6] predictive models have shown impressive prediction accuracy on stroke diagnosis tasks, such as slice classification and lesion segmentation [11, 7]. However, these DNNs trained on clean images through *Empirical Risk Minimization (ERM)* are vulnerable to perturbations in the input images [2]. Whatever the reconstruction method used, even the slightest perturbation in accelerated MR images can lead to a wrong stroke prediction from the AI models. Therefore, building robust DNN models to handle the perturbed MR image input is important for MRI acceleration.

In this paper, we introduce a *Distributionally Robust Learning (DRL)*-based approach [4] into the deep MR image classifier training, in order to improve the model robustness to the image perturbation resulting from the signal sparsity in accelerated MRI. Compared with ERM, DRL is an optimization method minimizing the worst-case loss over an ambiguity set, therefore, can tolerate outliers in the data [5]. We implemented DRL to different linear layers in deep CNN/ViT classifiers, and applied a randomized training approach to improve the training efficiency. Our results show that on a real-world dataset, DRL can significantly improve the stroke classification performance of ERM and other baseline defensive training methods, when the signal sparsity and noise in accelerated MRI are generated by the *Cartesian Undersampling (CU)* method [20] and *White Gaussian Noise (WGN)*. We further show that in highly perturbed MR images where the ERM model and even clinicians cannot give a reliable diagnosis, our DRL model can still correctly recognize stroke, which establishes that our method can assist accelerated MRI diagnosis.

2 Methodology

2.1 Distributionally robust learning

We will use the DRL framework under a multi-class classification setting developed by [4] in a DNN-based stroke diagnosis application. We provide a brief overview of the DRL model. Assume that there are K classes, and our goal is to classify an example with an input feature $\mathbf{x} \in \mathbb{R}^d$ to one of the K classes with a one-hot class label $\mathbf{y} \in \{0, 1\}^K$. Logistic regression solves this problem by minimizing the following expected true risk

$$\inf_{\mathbf{B}} \mathbb{E}^{\mathbb{P}^*} [h_{\mathbf{B}}(\mathbf{x}, \mathbf{y})], \quad (1)$$

where $\mathbf{B} \triangleq [\mathbf{w}_1 \cdots \mathbf{w}_K]$ is the coefficient matrix, \mathbb{P}^* is the true distribution of the data (\mathbf{x}, \mathbf{y}) , $h_{\mathbf{B}}(\mathbf{x}, \mathbf{y}) \triangleq \log \mathbf{1}' e^{\mathbf{B}'\mathbf{x}} - \mathbf{y}'\mathbf{B}'\mathbf{x}$ is the loss function to be minimized, and $\mathbb{E}^{\mathbb{P}^*}$ denotes the expectation under the distribution \mathbb{P}^* . Since \mathbb{P}^* is usually unknown, Problem (1) cannot be solved directly. The ERM approach tackles this by replacing the expected true risk by a sample averaged risk. Given N realizations of (\mathbf{x}, \mathbf{y}) , ERM minimizes the following empirical risk

$$\inf_{\mathbf{B}} \frac{1}{N} \sum_{i=1}^N h_{\mathbf{B}}(\mathbf{x}_i, \mathbf{y}_i). \quad (2)$$

ERM can produce unreliable solutions when the samples are contaminated by noise or drawn from an outlying distribution. To obtain robust estimators that can hedge against noise in the training data and generalize well out-of-sample, [4] proposed the DRL framework under the Wasserstein metric. Specifically, it minimizes the worst-case expected loss over a set of probability distributions

$$\inf_{\mathbf{B}} \sup_{\mathbb{Q} \in \Omega} \mathbb{E}^{\mathbb{Q}} [h_{\mathbf{B}}(\mathbf{x}, \mathbf{y})], \quad (3)$$

where Ω contains a set of probability distributions that are close to the empirical distribution $\hat{\mathbb{P}}_N$ measured by the Wasserstein metric, $\Omega \triangleq \{\mathbb{Q} \in \mathcal{P}(\mathcal{Z}) : W_1(\mathbb{Q}, \hat{\mathbb{P}}_N) \leq \epsilon\}$, where \mathcal{Z} is the set of possible values for (\mathbf{x}, \mathbf{y}) , $\mathcal{P}(\mathcal{Z})$ is the space of all probability distributions supported on \mathcal{Z} , ϵ is a pre-specified radius of the ambiguity set Ω , $\hat{\mathbb{P}}_N$ is the empirical distribution that assigns an equal probability $1/N$ to each observed sample $(\mathbf{x}_i, \mathbf{y}_i)$, and $W_1(\mathbb{Q}, \hat{\mathbb{P}}_N)$ is the order-1 Wasserstein distance between \mathbb{Q} and $\hat{\mathbb{P}}_N$ defined as

$$W_1(\mathbb{Q}, \hat{\mathbb{P}}_N) \triangleq \min_{\Pi \in \mathcal{P}(\mathcal{Z} \times \mathcal{Z})} \left\{ \int_{\mathcal{Z} \times \mathcal{Z}} l(\mathbf{z}_1, \mathbf{z}_2) \Pi(d\mathbf{z}_1, d\mathbf{z}_2) \right\},$$

where Π is the joint distribution of $\mathbf{z}_1 \triangleq (\mathbf{x}_1, \mathbf{y}_1)$ and $\mathbf{z}_2 \triangleq (\mathbf{x}_2, \mathbf{y}_2)$ with marginals \mathbb{Q} and $\hat{\mathbb{P}}_N$, respectively, and l is a distance metric on the data space.

An equivalent reformulation (4) of (3) was developed by [4] when $l(\mathbf{z}_1, \mathbf{z}_2) \triangleq \|\mathbf{W}^{1/2}(\mathbf{x}_1 - \mathbf{x}_2)\|_2 + M\|\mathbf{y}_1 - \mathbf{y}_2\|_2$, where \mathbf{W} is a positive semidefinite weight matrix to account for any transformation on the input feature \mathbf{x} and can be estimated from data using metric learning (see Sec. 2.2) and with $M \rightarrow \infty$:

$$\inf_{\mathbf{B}} \frac{1}{N} \sum_{i=1}^N h_{\mathbf{B}}(\mathbf{x}_i, \mathbf{y}_i) + \epsilon 2^{1/2} \|\mathbf{W}^{-1/2} \mathbf{B}\|_2, \quad (4)$$

where $\|\mathbf{W}^{-1/2} \mathbf{B}\|_2$ is the induced ℓ_2 -norm of $\mathbf{W}^{-1/2} \mathbf{B}$.

2.2 DRL for deep stroke diagnosis networks

We apply DRL to ViT and CNN-based MR image stroke classification models, in order to enhance their robustness against image perturbations in accelerated MRI. We apply the DRL reformulation (4) to the last layer B , and certain intermediate linear layers in a deep MR image classifier. For a ViT model (cf. Fig. 1), we apply DRL to the patch projection layer P and the final linear classification layer B , in order to predict if an MR image slice is normal (label 0) or depicts a stroke lesion (label 1). To speed up the training process, during each epoch, we randomly pick one layer L to train while keeping all other layers frozen. A validation set \mathcal{V} is used to tune the hyperparameters (e.g., regularization coefficients). To account for the non-linear transformation of the raw image resulted from all layers before L , we solve the following *Linear Matrix Inequality (LMI)* problem [4] to estimate a weight matrix \mathbf{W} :

$$\begin{aligned} \min_{\mathbf{W} \succ 0} \quad & \sum_{\mathbf{x}_i \in \mathcal{D}} \sum_{t=1}^T \|\mathbf{W}^{1/2}(\phi_L^{(t)}(\tilde{\mathbf{x}}_i) - \phi_L^{(t)}(\mathbf{x}_i))\|_2^2 \\ \text{s.t.} \quad & \frac{1}{|\mathcal{S}|} \sum_{(i,j) \in \mathcal{S}} \sum_{t=1}^T \|\mathbf{W}^{1/2}(\phi_L^{(t)}(\tilde{\mathbf{x}}_i) - \phi_L^{(t)}(\tilde{\mathbf{x}}_j))\|_2^2 \geq c, \\ & \frac{1}{|\mathcal{S}|} \sum_{(i,j) \in \mathcal{S}} \sum_{t=1}^T \|\mathbf{W}^{1/2}(\phi_L^{(t)}(\mathbf{x}_i) - \phi_L^{(t)}(\mathbf{x}_j))\|_2^2 \geq c, \end{aligned} \quad (5)$$

where $\tilde{\mathbf{x}}_i$ is the perturbed version of an MR image slice \mathbf{x}_i , \mathcal{D} the training set, $\mathcal{S} \triangleq \{(i, j) | \mathbf{x}_i, \mathbf{x}_j \in \mathcal{D}, \mathbf{y}_i \neq \mathbf{y}_j\}$, $|\mathcal{S}|$ denotes the cardinality of the set \mathcal{S} , ϕ_L is the input to L and $\phi_L^{(t)}$ is the t -th hidden state (i.e., the vector representation for each instance in the sequence, output by and fed into different layers in ViT) in the sequence ϕ_L of length T , and c is a fixed parameter. $T = 1$ if L refers to the B layer. The intuition of (5) is that in the transformed feature space, distance between the clean and perturbed version of a slice will be minimized, while slices from different classes (normal and stroke) are sufficiently far away. For a CNN model, we only applied DRL to the final linear layer B .

We chose two approaches to generate the perturbation in accelerated MRI. *Cartesian Undersampling (CU)* perturbation [20] keeps only the central and a few randomly-sampled parts of the k-space; the corresponding reconstructed MR image only keeps the main structural information in the brain and introduces misalignments. A smaller central fraction f used in k-space indicates a larger perturbation. Noise might be introduced during the signal transmission, so we add *White Gaussian Noise (WGN)* as another type of perturbation, where the standard deviation σ is regarded as the perturbation intensity. To show the strength of DRL, in addition to ERM, we also apply *Brute-force Adversarial Training (BAT)* [2] and *Projected Gradient Descent (PGD)* [12] as baseline methods. Among all of the current defensive training methods which improve the model robustness against perturbations, BAT and PGD are representative. BAT adds noisy samples into the training set, therefore is simple and widely used. PGD is known to be robust to a wide range of image perturbations, and is considered a state-of-the-art method.

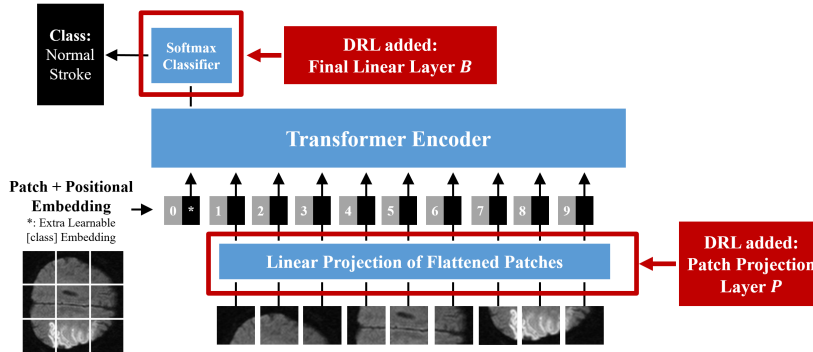


Fig. 1: Adding DRL into different layers in ViT.

3 Experiments

3.1 Experimental materials and settings

Our dataset included MRI brain scans from 226 patients performed at an urban tertiary referral academic medical center that is a comprehensive stroke center. Clinical scans of adult patients aged 18-89 years with recent (acute or subacute) strokes were identified between 1/1/2013 and 1/1/2021 for inclusion in this study via a search of the Philips Performance Bridge. Scans meeting this criteria were downloaded and simultaneously anonymized to preserve patient anonymity and prevent disclosure of protected health information as part of this IRB exempt study. No patient demographic information was retained for the scans, as it was considered to represent an unnecessary risk for accidental release of protected health information. The diffusion weighted images with a gradient of $B=1000$ were utilized for the analysis (see the Supplement⁴ for information about the MRI scanner and parameter settings). Each MR image contains multiple slices, and every slice was annotated as normal or stroke by a board-certified neuro-radiologist with a subspecialty certification. Annotation of the strokes was performed on the diffusion weighted images using ITK-SNAP (ver. 3.80) [19], and all included MRI examinations were reviewed by the neuroradiologist during the annotation process to ensure that the images were of diagnostic quality without significant motion degradation or other artifacts. To avoid the dependency among the slices from the same subject, we applied a 2-d acquisition during the MR imaging, and implemented a slice-level MR image preprocessing. While the whole dataset includes 4,883 (74.7%) normal slices and 1,650 (25.3%) stroke slices, we further randomly split them into training/validation/test sets using the ratio 80%/10%/10%. For the training set, we implemented data augmentation strategies by rotating or flipping each slice. Finally, the training/validation/test set contains 31,356/653/654 slices, correspondingly.

⁴ Supplement and source code are available at https://github.com/noc-lab/drl_mri.

We implemented DRL to both CNN and ViT models. For the CNN model, we used a ResNet-18 [9] architecture, while for the ViT model, we first pre-trained a 4-layer ViT using a self-supervised pre-training method called Masked Autoencoder (MAE) [8], using the T1/T2-weighted brain MR images in the IXI

dataset [1]. MAE pre-training first randomly masks 75% of the image patches in an MRI slice input, and then uses a ViT encoder-decoder architecture to reconstruct the masked MRI patches, in order to learn the dependency among different locations in the brain. After 400 pre-training epochs, an overall satisfying reconstruction result can be observed in Fig. 2.

To evaluate the binary classification performance of different models, we use the *Area Under the Receiver Operating Characteristic (AUROC)* curve as our main metric. As our dataset is unbalanced, we also considered the *Area Under Precision-Recall Curve (AUPRC)*. We ran the experiments 3 times using different random seeds. The training of our DNNs were implemented on 3 NVIDIA RTX A6000 (48GB VRAM) GPUs, and each DRL training epoch can be completed within 0.03 GPU hours. We used a learning rate of 1×10^{-5} and batchsize of 128 for DRL training, while no weight decay was applied. To solve the LMI problem in (5), we used SDPT3 v4.0 [17] as the solver. We set the CU perturbation with the acceleration factor of 4, 6, 8, 12 with the central fraction of 8%, 6%, 4% and 2% in k-space respectively, and the remaining parts were chosen randomly in the peripheral region accordingly.

3.2 Results

We show the stroke classification AUROC in Fig. 3. When the k-space subsampling fraction decreased and the signal became sparser, the performance of both ViT and CNN models trained under ERM dropped significantly, from around 95% to below 80%. DRL significantly improved the AUROC of the ERM-based ViT model from 74.5% to 83.1% when the MR images were under extreme CU perturbation, while only slightly influenced model performance on the clean test set. For WGN, the largest improvement brought to ERM-based ViT model was 16.9%. Although we only applied DRL to the last layer of the CNN model, the improvement against ERM was still remarkable, up to 11.9%/4.9% for CU/WGN. With BAT and PGD adversarial training, the corresponding ViT or CNN models were also improved, though when DRL was combined with BAT and PGD, the model robustness can be further enhanced. Table 1 shows the maximum AUROC and AUPRC improvement that DRL can bring to different baseline methods. For ViT and CNN models, the AUROC improvement w.r.t BAT/PGD defensive methods is up to 23.9%/12.2%, respectively. Note that the




Fig. 2: MAE pre-trained models reconstruct the masked MR images. This example uses a T2-weighted image.

perturbed MRI samples used to implement BAT were the same as those used by DRL, which shows that DRL is a more effective way to exploit the information in adversarial samples, compared to simply adding the blurry images into the training set. For CU perturbation, our best combined model using DRL improved the AUROC/AUPRC of the ERM model by up to 15%/12.5%, while this improvement under WGN perturbation was up to 18.8%/36.2%.

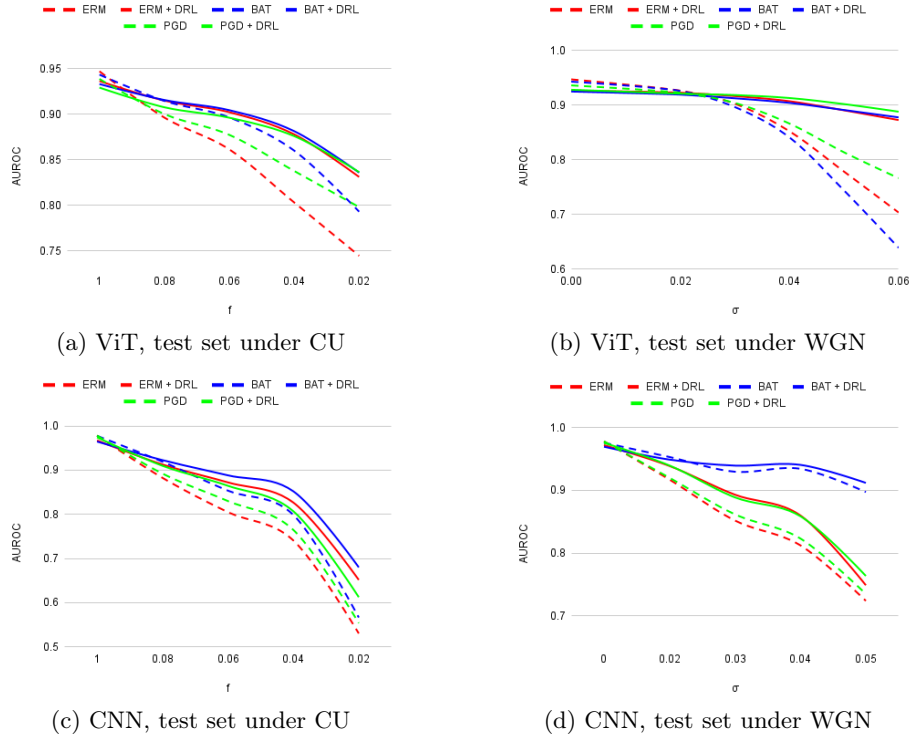


Fig. 3: AUROC of different methods using ViT and CNN models.

Under CU perturbation, we further show that our DRL model can recognize stroke while clinicians may fail to. In Fig. 4, the stroke MRI slices from the test sets are under different levels of CU perturbations. For both ERM- and DRL-based ViT models, we maximized the F1 score of the stroke class on the training set to calculate the optimal decision threshold for stroke prediction, in order to balance the precision and recall. When the k-space signal becomes more sparse, the reconstructed MRI slices get more blurry and the lesion areas become less recognizable, even for human eyes. As a result, the ERM model fails to detect stroke under high perturbation levels. Nevertheless, the DRL model can tolerate more intense CU perturbation and recognize stroke slices that may even be misclassified by clinicians, which reveals its value in improving the diagnosis in

Table 1: AUROC and AUPRC improvement of DRL over three baseline methods using ViT/CNN models under CU and WGN perturbations on the test set.

	ViT				CNN				
	CU		WGN		CU		WGN		
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	
ERM+DRL	AUROC	8.65%	0.00%	16.87%	0.04%	11.91%	0.18%	4.92%	0.24%
	AUPRC	8.31%	0.02%	29.77%	0.10%	8.65%	0.15%	6.81%	0.79%
BAT+DRL	AUROC	4.29%	0.00%	23.90%	0.05%	11.26%	0.07%	1.42%	0.04%
	AUPRC	4.37%	0.01%	39.75%	0.12%	9.32%	0.11%	3.59%	0.22%
PGD+DRL	AUROC	3.83%	0.01%	12.19%	0.05%	6.00%	0.14%	3.45%	0.28%
	AUPRC	2.88%	0.02%	18.68%	0.06%	3.85%	0.21%	4.13%	0.77%

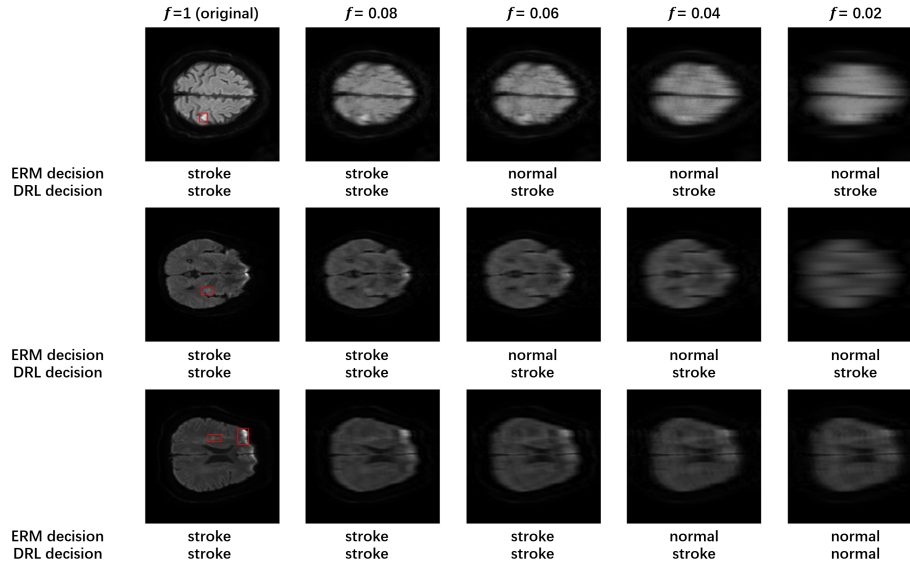


Fig. 4: Stroke slices where the ERM model and clinicians may fail when the CU perturbation is large, while the DRL model succeeds. The red boxes indicate the stroke lesion areas in the clean images.

an accelerated MRI mode. We verified the effectiveness of our approach on the actual clinical scans acquired for clinical care and not just for research purposes, suggesting that the methods and findings in the current study should be generalizable to routine clinical practice conditions and potentially other types of clinical image-based diagnosis (e.g., brain tumor) as well. In addition, our DRL

framework does not necessarily need to be used in isolation, rather it can also be combined with other performance boosting methods in accelerated MRI to further improve them, just like for BAT and PGD.

4 Conclusions

In this study, we implemented a DRL-based robust learning approach to improve the robustness of deep image classifiers, in order to achieve more accurate stroke classification from brain MR images reconstructed from a sub-sampled k-space. Our work can potentially be applied to accelerate and improve time-critical stroke diagnosis. Future work can apply DRL to more MRI diagnosis tasks (e.g., lesion area segmentation), justifying its effectiveness on more types of sub-sampling methods in MRI acceleration.

References

1. IXI dataset. <https://brain-development.org/ixi-dataset/>
2. Akhtar, N., Mian, A.: Threat of adversarial attacks on deep learning in computer vision: A survey. *Ieee Access* **6**, 14410–14430 (2018)
3. Alberts, M.J., Horner, J., Gray, L., Brazer, S.R.: Aspiration after stroke: lesion analysis by brain mri. *Dysphagia* **7**, 170–173 (1992)
4. Chen, R., Hao, B., Paschalidis, I.C.: Distributionally robust multiclass classification and applications in deep image classifiers. In: *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1–2. IEEE (2023)
5. Chen, R., Paschalidis, I.C.: A robust learning approach for regression models based on distributionally robust optimization. *The Journal of Machine Learning Research* **19**(1), 517–564 (2018)
6. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
7. Gu, Y., Piao, Z., Yoo, S.J.: Sthardnet: Swin transformer with hardnet for mri segmentation. *Applied Sciences* **12**(1), 468 (2022)
8. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16000–16009 (2022)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
10. Hyun, C.M., Kim, H.P., Lee, S.M., Lee, S., Seo, J.K.: Deep learning for undersampled mri reconstruction. *Physics in Medicine & Biology* **63**(13), 135007 (2018)
11. Liu, L., Chen, S., Zhang, F., Wu, F.X., Pan, Y., Wang, J.: Deep convolutional neural network for automatically segmenting acute ischemic stroke lesion in multi-modality mri. *Neural Computing and Applications* **32**, 6545–6558 (2020)
12. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083* (2017)

13. Montalt-Tordera, J., Muthurangu, V., Hauptmann, A., Steeden, J.A.: Machine learning in magnetic resonance imaging: image reconstruction. *Physica Medica* **83**, 79–87 (2021)
14. Plein, S., Ryf, S., Schwitter, J., Radjenovic, A., Boesiger, P., Kozerke, S.: Dynamic contrast-enhanced myocardial perfusion mri accelerated with k-t sense. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **58**(4), 777–785 (2007)
15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. pp. 234–241. Springer (2015)
16. Stahl, R., Krug, R., Kelley, D.A., Zuo, J., Ma, C.B., Majumdar, S., Link, T.M.: Assessment of cartilage-dedicated sequences at ultra-high-field mri: comparison of imaging performance and diagnostic confidence between 3.0 and 7.0 t with respect to osteoarthritis-induced changes at the knee joint. *Skeletal radiology* **38**, 771–783 (2009)
17. Toh, K.C., Todd, M.J., Tütüncü, R.H.: SDPT3 — a MATLAB software package for semidefinite programming, version 1.3. *Optimization methods and software* **11**(1-4), 545–581 (1999)
18. Wang, H., Peng, H., Chang, Y., Liang, D.: A survey of gpu-based acceleration techniques in mri reconstructions. *Quantitative imaging in medicine and surgery* **8**(2), 196 (2018)
19. Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., Gerig, G.: User-guided 3d active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *Neuroimage* **31**(3), 1116–1128 (2006)
20. Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M.J., De-fazio, A., Stern, R., Johnson, P., Bruno, M., et al.: fastmri: An open dataset and benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839* (2018)