



MOTIVATION

- Magnetic Resonance Imaging (MRI) encompasses a set of powerful imaging techniques for understanding brain structure and diagnosing pathology.
- Typical MRI sequence consists of T1-weighted, T2-weighted, and FLAIR sequences.
- Significant equipment costs and acquisition times have inhibited uptake of this critical technology, adversely impacting health equity globally [1].

Is there a way to translate brain T1 scans into T2 scans, thereby obviating T2 acquisitions?


PREVIOUS WORK

- Traditional methods:** Markov random fields [2], regression forests [3], and convolutional neural network-based [4] methods have been proposed. However, these methods do not produce the requisite high-quality translated images.
- GAN-based methods:** Both supervised [5] and unsupervised [6] methods have been proposed for healthy and unhealthy brain MRI. Addition of non-adversarial losses has shown to improve performance [5].

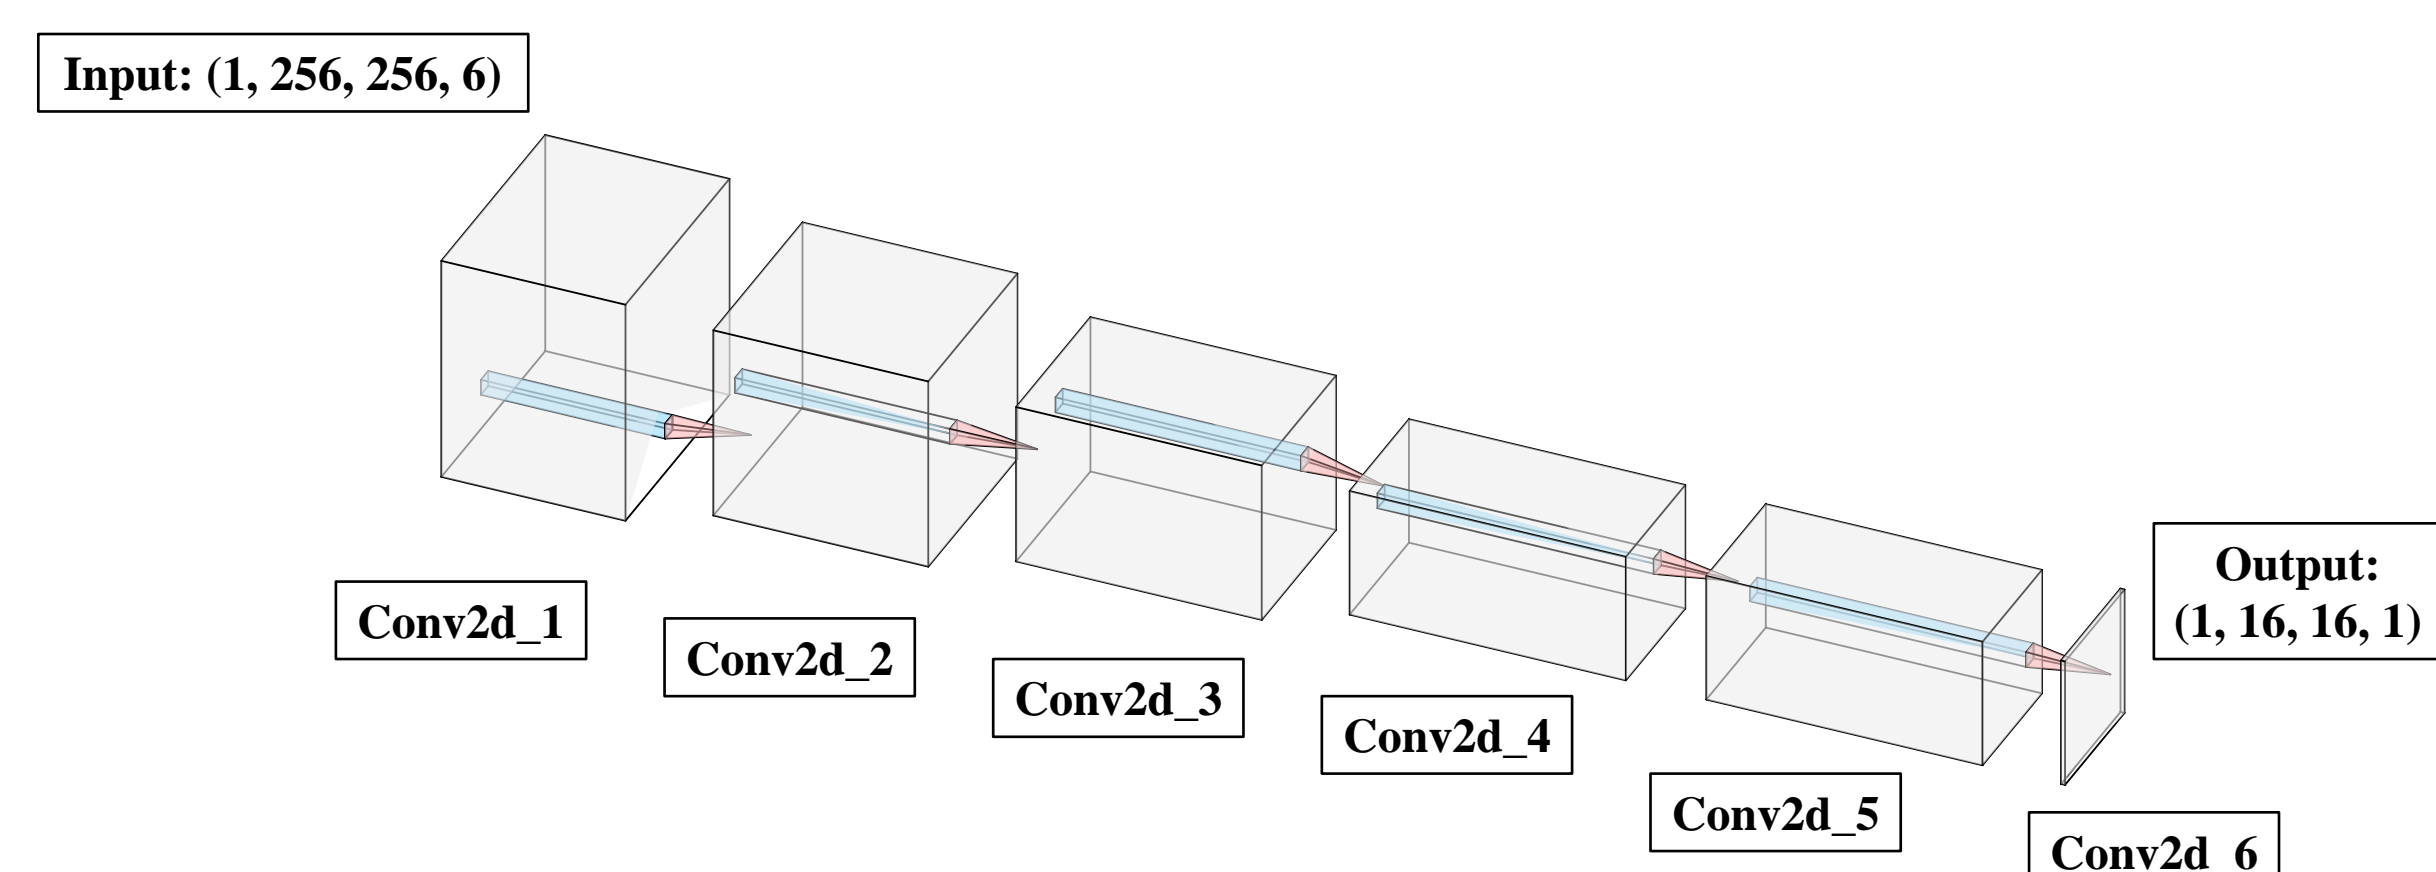
We propose *pTransGAN* which is capable of concurrently translating both healthy and unhealthy brain T1 into T2 scans

ARCHITECTURES

Generator: 6 U-blocks with skip connections



Discriminator: 70x70 PatchGAN



LOSS FUNCTION AND DATA

$$\mathcal{L}_{pTransGAN} = \lambda_{cGAN} \mathcal{L}_{cGAN} + \lambda_{L1} \mathcal{L}_{L1} + \lambda_{style} \mathcal{L}_{style} + \lambda_{content} \mathcal{L}_{content}$$

$$\mathcal{L}_{style} = \sum_{i=1}^{Total\ Blocks} \lambda_{style, i} * \frac{1}{4d_i^2} * ||Gram_i(y) - Gram_i(\hat{y})||_F^2$$

$$\mathcal{L}_{content} = \sum_{i=1}^{Total\ Blocks} \lambda_{content, i} * ||F_i(y) - F_i(\hat{y})||_F^2$$

- Discriminator trains on adversarial and reconstruction errors. Generator trains on style and content losses.
- Internal feature activations of a pre-trained VGG-19 after inferring generated and real images are used to calculate the non-adversarial style and content losses.
- Batch size of 1, epochs = 100.
- $\eta_{discriminator} = 0.008$, $\eta_{generator} = 0.002$.
- Paired T1 and T2 healthy (IXI dataset [7]) and unhealthy (BRaTS2020 [8]) brain MRI in axial direction were used.

Dataset	Number of images (Number of patients)		
	Training	Hyperparameter optimization	Testing
IXI (Healthy)	5071 (n=461)	638 (n=58)	638 (n=58)
BRaTS2020 (Unhealthy)	6482 (n=463)	434 (n=31)	434 (n=31)

EXPERIMENTS

- Experiment 1:** We performed an ablation study to investigate the effect of adding non-adversarial loss functions.
- Experiment 2:** We trained *pTransGAN* on healthy data and tested it on unhealthy data (and vice-versa) to see the effect of domain shift on *pTransGAN*.
- Experiment 3:** We introduce and test a simultaneous training protocol (a single *pTransGAN* model is trained three times on a healthy data and three times on an unhealthy data for each iteration).
- PSNR, SSIM, MSE, LPIPS, UQI, VIF metrics were used to judge the quality of translated images.

RESULTS

- High level brain features like the cerebellum are sharper after adding non-adversarial losses (Figure 1). Only small improvements are seen in metrics though (Table 1).
- pTransGAN* trained on unhealthy data creates accurate translations of glioma. Global shape, tumor boundary, and contrast are reproduced (Figure 1)
- pTransGAN* trained on healthy data cannot produce good translations of unhealthy T1 MRI (Table 1). Addition of non-adversarial losses improves performance.

Table 1: Average metrics on test set for experiments. Adv: adversarial loss, L1: reconstruction loss, s: style loss, c: content loss

Train set	Healthy	Healthy	Healthy	Healthy	Unhealthy	Unhealthy	Sim.	Sim.
Test set	Healthy	Healthy	Unhealthy	Unhealthy	Unhealthy	Healthy	Healthy	Unhealthy
Loss	adv, L1	+ s, c	adv, L1	+ s, c	+ s, c	+ s, c	+ s, c	+ s, c
PSNR (dB)	24.7	24.9	17.9	18.2	19.5	17.9	24.1	19.9
SSIM	0.875	0.883	0.572	0.685	0.831	0.63	0.851	0.825
MSE	0.003	0.004	0.02	0.017	0.012	0.016	0.004	0.011
LPIPS	0.061	0.06	0.256	0.197	0.113	0.265	0.065	0.132
UQI	0.941	0.952	0.501	0.551	0.835	0.749	0.885	0.862
VIF	0.882	0.875	0.778	0.601	0.752	0.559	0.821	0.762

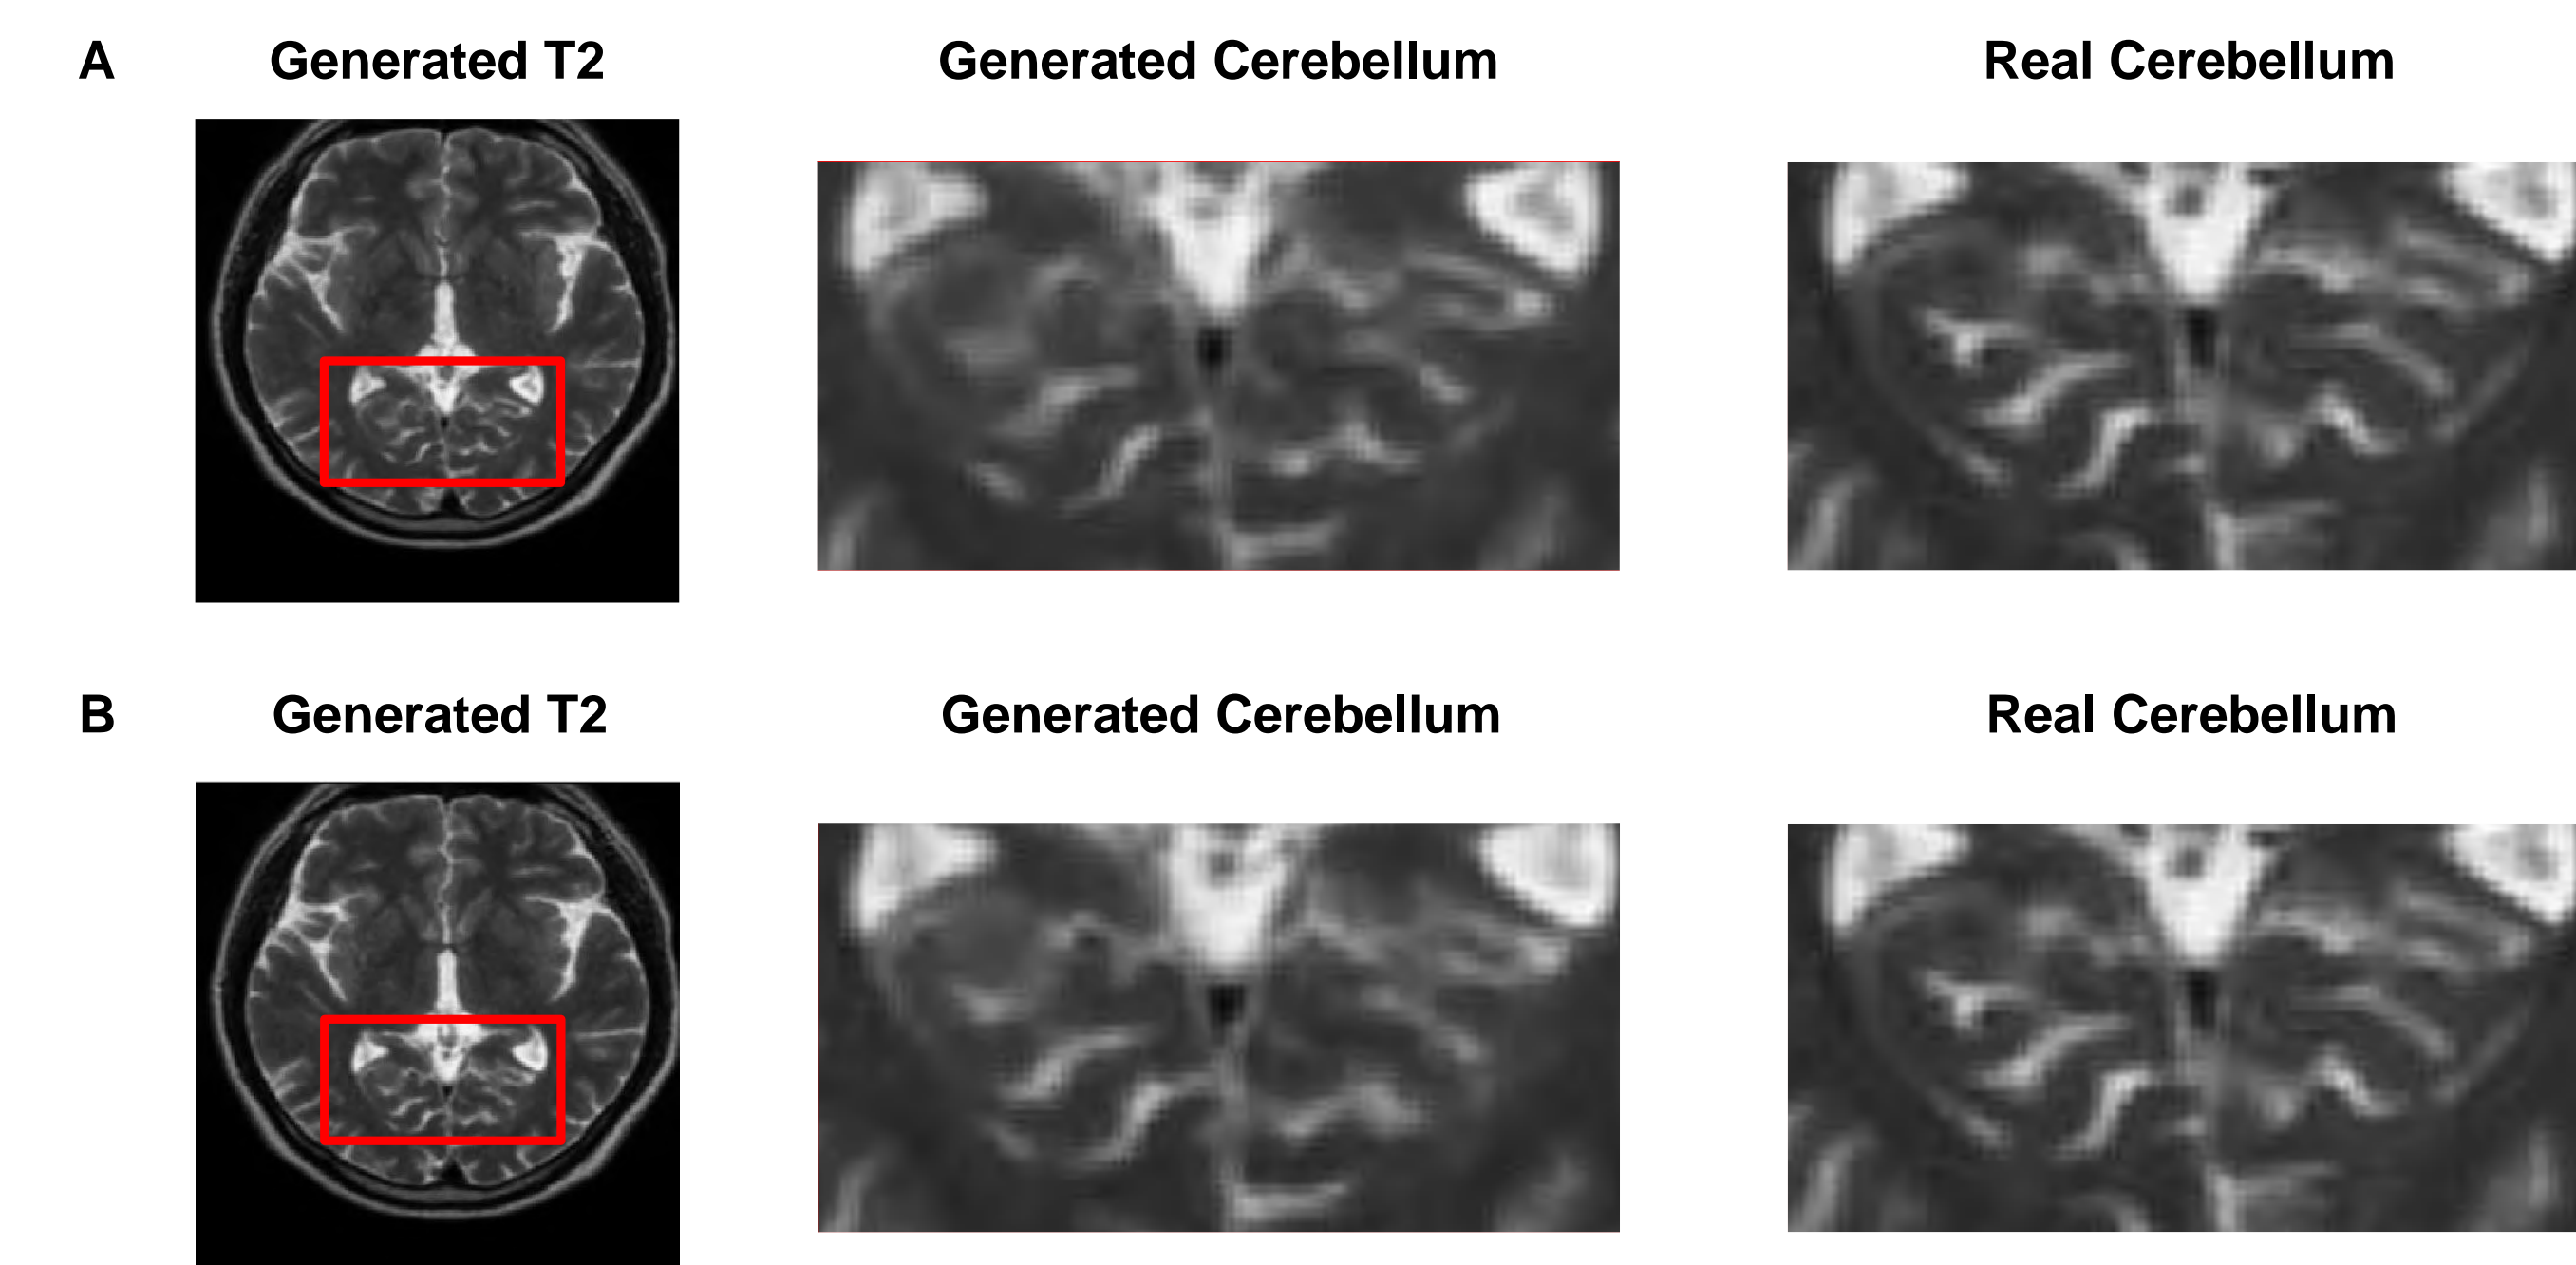


Figure 1: (A) T2 generated without non-adversarial losses (B) T2 generated with non-adversarial losses.




Figure 2: *pTransGAN* can accurately translate unhealthy T1 scans

- A single *pTransGAN* model trained using the simultaneous training protocol can produce translations of both healthy and unhealthy scans.

CONCLUSIONS

- Addition of non-adversarial losses produces higher quality translations and improves model robustness
- With simultaneous training protocol, a single model can translate healthy and unhealthy T1 scans

REFERENCES

- Stewart, *Statistica.com*, 2019
- Rousseau, *ECCV*, 2008
- Jog, *IEEE*, 2013
- Nyugen, *MICCAI*, 2015
- Armanious, *Comp. Med. Img. & Graphics*, 2020.
- Armanious, *IEEE*, 2019
- brain-development.org/ixi-dataset/
- med.upenn.edu/cbica/brats2020/