Self-Supervised CSF Inpainting with Synthetic Atrophy for Improved Accuracy Validation of Cortical Surface Analyses

Author name(s) withheld EMAIL(S) WITHHELD
Address withheld

Editors: Under Review for MIDL 2023

Abstract

Accuracy validation of cortical thickness measurement is a difficult problem due to the lack of ground truth data. To address this need, many methods have been developed to synthetically induce gray matter (GM) atrophy in an MRI via deformable registration, creating a set of images with known changes in cortical thickness. However, these methods often cause blurring in atrophied regions, and cannot simulate realistic atrophy within deep sulci where cerebrospinal fluid (CSF) is obscured or absent. In this paper, we present a solution using a self-supervised inpainting model to generate CSF in these regions and create images with sharper GM/CSF boundaries. Specifically, we introduce a novel, 3D GAN model that incorporates patch-based dropout training, edge map priors, and sinusoidal positional encoding, all of which are established methods previously limited to 2D domains. We show that our framework significantly improves the quality of the resulting synthetic images and is adaptable to unseen data with fine-tuning. We also demonstrate that our resulting dataset can be employed for accuracy validation of cortical segmentation and thickness measurement.

Keywords: Inpainting, Self-supervised, Sinusoidal positional encoding, Edge detection, Synthetic atrophy, Multi-modal MRI, Cortical thickness measurement

1. Introduction

Cortical segmentation and thickness measurement are powerful tools for measuring changes in the gray matter (GM) of the brain (Fischl and Dale, 2000; Han et al., 2004; Oguz and Sonka, 2014b,a; Reuter et al., 2012). Evaluating the accuracy of these methods is difficult due to the lack of ground truth data. Many have proposed synthetic atrophy methods as a solution to generate longitudinal data with known changes in cortical thickness (CT) (Bernal et al., 2021; Davatzikos et al., 2001; Freeborough and Fox, 1997; Karaçali and Davatzikos, 2006; Khanal et al., 2017; Xia et al., 2019; Larson and Oguz, 2022, 2021). Several of these techniques (Davatzikos et al., 2001; Khanal et al., 2017; Larson and Oguz, 2022, 2021; Karaçali and Davatzikos, 2006) operate using a deformable registration to move the outer GM boundary in towards the white matter (WM), effectively shrinking the GM while expanding the cerebrospinal fluid (CSF). However, these registration-based methods often induce GM/CSF boundary blurring in the atrophied images due to interpolation. These methods can also struggle within deep sulci where CSF is not originally visible due to partial volume effects, which prevents deformable models from generating CSF in the atrophied result despite GM thinning. To overcome these two problems, we propose the use of inpainting with a self-supervised deep-learning model to generate CSF in these atrophied regions and create images with sharper GM/CSF boundaries.

© 2023 Anonymous.

Inpainting is often used to replace missing or corrupted regions in a volume. While traditional methods (Barnes et al., 2009; Li et al., 2017; Fleishman et al., 2018) can produce high-quality images, they are limited in their ability to semantically reconstruct the missing parts of an image (Zhou et al., 2021). In contrast, deep learning-based methods (Pathak et al., 2016; Iizuka et al., 2017; Nazeri et al., 2019; Zhang et al., 2020; Song et al., 2018; Xu et al., 2021; Ren et al., 2019; Suvorov et al., 2022) have demonstrated improved results, especially in larger missing regions. However, these are primarily restricted to 2D domains, rendering them unable to maintain consistency across image slices when applied to 3D volumes. Many of these methods also require extensive training data, which can be prohibitive in new datasets.

In this paper, we present a 3D GAN model for CSF inpainting. First, we learn to synthesize MRIs given the desired tissue labels using a patch-based dropout scheme in real, i.e., not atrophied, images. We then apply this model to synthetically atrophied images obtained using a registration-based method. Our main contributions are:

1. The use of CSF inpainting for synthetic cortical atrophy induction is the first of its kind, and provides an interesting new challenge for the MIDL community.

2. We construct a novel 3D GAN model capable of inpainting multiple tissue types by extending a previous method (Pathak et al., 2016) to include multi-label tissue information within dropout regions rather than a binary mask.

3. We combine several existing 2D methods and extend them to 3D, including a self-supervised dropout scheme (Pathak et al., 2016), edge priors (Nazeri et al., 2019), and sinusoidal positional encoding (SPE) (Xu et al., 2021).

4. We show that our model can generate higher quality synthetic images compared to 2D methods, and that it can effectively adapt to unseen datasets.

2. Methods

Fig. 1 describes the structure of our 3D GAN model. We derived 5 input channels from each training sample (Sec. 2.2), and used a 3D GAN network architecture (Sec. 2.3) with adapted 3D SPE. We performed cross-validation training and fine-tuning for unseen data (Sec. 2.4) with publicly available datasets (Sec. 2.1).

2.1. Data

NITRC Kirby Dataset (Kirby). The Kirby dataset (n=21) (Landman et al., 2011) is a test-retest dataset of healthy controls. Each subject is associated with two sets of MPRAGE/T1w ($1.0\times1.0\times1.2$ mm^3) and FLAIR ($1.1\times1.1\times1.1$ mm^3) images procured on the same day; each set was used as an independent training sample.

Validation Data for Cortical Reconstruction Algorithms (VDCRA). The VDCRA (Shiee et al., 2014) (n=10) consists of 5 healthy controls (HC) and 5 multiple sclerosis (MS) subjects, each with a T1w and a FLAIR image, with the same resolutions as the Kirby dataset. Additionally, each subject is associated with landmarks denoting the WM and GM cortical surfaces in seven different regions on either hemisphere (1,680 per subject).

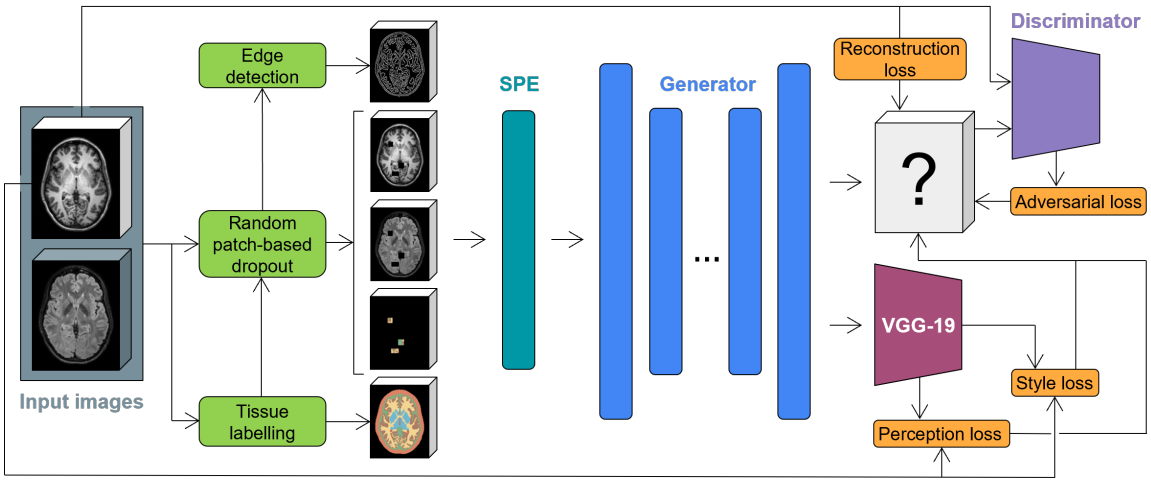

Figure 1: Our self-supervised training pipeline for a target T1w image. We pre-processed (green) both the T1w and FLAIR image (gray) to obtain 5 input channels. We used 3D SPE (teal) and adversarial training with a 3D generator (blue) and discriminator (purple). Loss functions (orange) were computed using the discriminator and a VGG-19 network (Simonyan and Zisserman, 2014) (magenta), supervised by the original T1w image.

2.2. Generation of Input Image Channels

We trained two separate models to inpaint T1w and FLAIR images, respectively. Each model took five input channels (Fig. 2), with $H \times W \times D$ voxels each, as described below.

MRI preprocessing. The T1w and FLAIR images, denoted $\mathbf{I}_{\mathrm{T1}}$ and $\mathbf{I}_{\mathrm{FL}}$, were pre-processed using the FreeSurfer program (FS) (Dale et al., 1999) to perform isotropic resampling, intensity normalization, and bias field correction. Although we trained a separate model for each modality, both models used both $\mathbf{I}_{\mathrm{T1}}$ and $\mathbf{I}_{\mathrm{FL}}$ as inputs.

Random patch-based dropout masks. We generated an image $\mathbf{M}_{\mathrm{dr}}$ with a random number of randomly placed dropout patches. These patches were constrained such that their combined volume was less than 30% of the total image foreground volume (Tang et al., 2022). Patches containing background voxels were rejected. We then defined a function $\phi(\cdot)$ that replaced voxels inside patches with Gaussian noise $\mathcal{N}(0, 1)$. Fig. 2B/C depicts the dropout images $\phi(\mathbf{I}_{\mathrm{T1}}, \mathbf{M}_{\mathrm{dr}})$ and $\phi(\mathbf{I}_{\mathrm{FL}}, \mathbf{M}_{\mathrm{dr}})$, respectively.

Whole-brain tissue labels. We used FS to compute a skull-strip mask and a tissue label map (aseg.mgz) for each training sample. These were converted into a custom labelling, denoted $\mathbf{I}_{\mathrm{tis}}$, with the following classes: cortical GM, WM, non-ventricle CSF, ventricles, non-cortical brain structures, head (remaining foreground voxels outside the skull-strip mask), and background (Fig. 2D). We also created an auxiliary label map $\mathbf{M}_{\mathrm{tis}}$ that retains only the tissue labels within the dropout patches, i.e., $\mathbf{M}_{\mathrm{tis}} = \mathbf{I}_{\mathrm{tis}} * \mathbf{M}_{\mathrm{dr}}$ (Fig. 2E).

Edge detection. The last input channel is the edge map of the target image (i.e., $\phi(\mathbf{I}_*, \mathbf{M}_{\mathrm{dr}})$), where $\mathbf{I}_*$ refers to $\mathbf{I}_{\mathrm{T1}}$ or $\mathbf{I}_{\mathrm{FL}}$). A 2D Canny edge detector (Canny, 1986) was applied to each axial slice; these edge slices were stacked into a 3D volume $\mathbf{E}(\phi(\mathbf{I}_*, \mathbf{M}_{\mathrm{dr}}))$ (Fig. 2F).

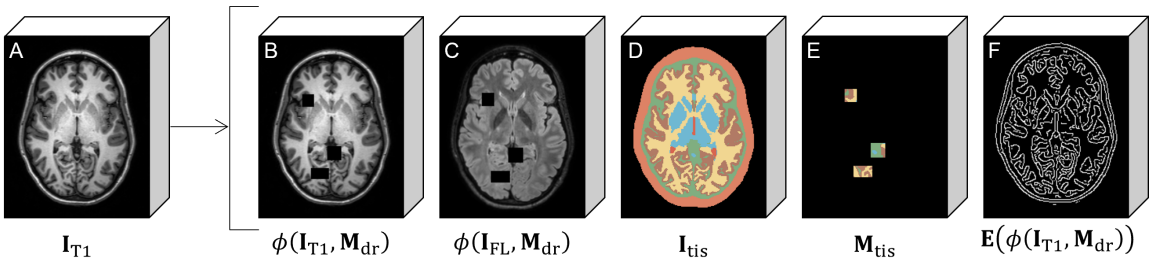

Figure 2: Example input channels for a target T1w image.

2.3. 3D GAN Model

Network architecture. Our proposed 3D GAN-style network consisted of two parts: a generator $\mathbf{G}(\cdot)$ and a discriminator $\mathbf{D}(\cdot)$. It possessed a typical (Nazeri et al., 2019; Zhang et al., 2020) encoder-decoder structure, but with 3D convolutions rather than 2D. The generator input was the concatenation of the five input channels, with the output $\mathbf{I}_{\mathrm{pred},*}$:

$$\mathbf{I}_{\mathrm{pred},*} = \mathbf{G}\left[\phi(\mathbf{I}_{T1}, \mathbf{M}_{\mathrm{dr}}), \phi(\mathbf{I}_{FL}, \mathbf{M}_{\mathrm{dr}}), \mathbf{I}_{\mathrm{tis}}, \mathbf{M}_{\mathrm{tis}}, \mathbf{E}(\phi(\mathbf{I}_*, \mathbf{M}_{\mathrm{dr}}))\right] \tag{1}$$

At each iteration of adversarial training, the generator (Apdx. B.1) synthesizes the whole image. We then extract the voxels within $\mathbf{M}_{\mathrm{dr}}$ to replace the corresponding voxels in $\phi(\mathbf{I}_*, \mathbf{M}_{\mathrm{dr}})$. The discriminator (Apdx. B.2) then classifies $\mathbf{I}_{\mathrm{pred},*}$ as either "real" or "fake".

Sinusoidal Positional Encoding. Recent research (Xu et al., 2021) demonstrated that the use of explicit positional encoding in generative models is crucial for maintaining a spatial inductive bias balanced across the entire image and consistent transformations between positions. We extended SPE to 3D with the following structure:

$$\underbrace{[sin(\omega_0 i), cos(\omega_0 i), \ldots]}_{\text{sagittal direction}}, \underbrace{[sin(\omega_0 j), cos(\omega_0 j), \ldots]}_{\text{coronal direction}}, \underbrace{[sin(\omega_0 k), cos(\omega_0 k), \ldots]}_{\text{axial direction}} \tag{2}$$

Here, $\omega_0 = 1/10000^{6/N}, i, j, k \in \mathbb{Z}$ in $[0, \lceil N/6 \rceil]$, for N input channels ($N = 5$ in our case).

Loss Function. Inspired by previous work (Nazeri et al., 2019), our proposed model was trained using a loss function $\mathcal{L}_{train} = \lambda_1 \mathcal{L}_{\mathrm{recon}} + \lambda_2 \mathcal{L}_{\mathrm{adv}} + \lambda_3 \mathcal{L}_{\mathrm{p\bar{e}rc}} + \lambda_4 \mathcal{L}_{\mathrm{s\bar{t}y}}$, given a reconstruction loss $\mathcal{L}_{\mathrm{recon}}$, adversarial loss $\mathcal{L}_{\mathrm{adv}}$, averaged perceptual loss $\mathcal{L}_{\mathrm{p\bar{e}rc}}$, and averaged style loss $\mathcal{L}_{\mathrm{s\bar{t}y}}$. Because both $\mathcal{L}_{\mathrm{p\bar{e}rc}}$ and $\mathcal{L}_{\mathrm{s\bar{t}y}}$ were calculated with the pre-trained 2D VGG-19 network, we computed these within axial slices $s \in S^{\mathrm{axial}}$ rather than in 3D.

$$\mathcal{L}_{\mathrm{recon}} = \mathbb{E}_{\mathbf{M}_{\mathrm{dr}}} ||\mathbf{I}_* - \mathbf{I}_{\mathrm{pred},*}|| \tag{3}$$

$$\mathcal{L}_{\mathrm{adv}} = \mathbb{E}_{(\mathbf{I}_*, \mathbf{I}_{\mathrm{tis}})}\left[\log \mathbf{D}(\mathbf{I}_*, \mathbf{I}_{\mathrm{tis}})\right] + \mathbb{E}_{(\mathbf{I}_{\mathrm{pred},*}, \mathbf{I}_{\mathrm{tis}})}\left[\log(1 - \mathbf{D}(\mathbf{I}_{\mathrm{pred},*}, \mathbf{I}_{\mathrm{tis}}))\right] \tag{4}$$

$$\mathcal{L}_{\mathrm{p\bar{e}rc}} = \sum_{s \in S^{\mathrm{axial}}} \mathbb{E}_i \left[\sum_i ||\psi_i(\mathbf{I}_*^s) - \psi_i(\mathbf{I}_{\mathrm{pred},*}^s)||\right] \tag{5}$$

$$\mathcal{L}_{\mathrm{s\bar{t}y}} = \sum_{s \in S^{\mathrm{axial}}} \mathbb{E}_j \left[||G_j^{\psi}(\mathbf{I}_{\mathrm{pred},*}^s) - G_j^{\psi}(\mathbf{I}_*^s)||\right] \tag{6}$$

Here, the function ψ corresponds to VGG-19 activation layers, and G_j^{ψ} is a $C_j \times C_j$ matrix constructed from activation maps ψ_j (at j^{th} layer) given the activation feature size of $C_j \times H_j \times W_j$. $\mathbb{E}_{(\mathbf{A},\mathbf{B})}$ denotes the combined mean of images $\mathbf{A}$ and $\mathbf{B}$. $\mathbf{I}^s \in \mathbb{R}^{H \times W}$ denotes the sth slice of the 3D image I.

2.4. Training and Fine-Tuning

Apdx. A details the training stages. We trained all models on an NVidia-2080 Ti GPU.
Cross-validation training strategy. We first trained our model using a five-fold, cross-validation strategy in the Kirby dataset. For each fold, 16 subjects were used for training and 5 for validation. Test-retest samples from the same subject were assigned to either both training or both validation to avoid leakage. For each fold, we used 210 $\mathbf{M}_{\text{dr}}$ masks per model (21 subjects $\times$ 2 samples $\times$ 5 dropout masks). For data augmentation (details in Apdx. C), we utilized MONAI's (Cardoso et al., 2022) foreground-cropping, rotating, and flipping methods in three axes. We also used the RandCropByPosNegLabeld method to crop the original images into smaller patches of size $48 \times 48 \times 48$ voxels centered around a random voxel in $\mathbf{M}_{\text{dr}}$. Loss function hyperparameters were $\lambda_1 = 1, \lambda_2 = 0.1, \lambda_3 = 0.1, \lambda_4 = 250$. We used a batch size of 4, the Adam optimizer with a momentum of 0.5, and an initial learning rate of 0.0002. After training for 80 epochs, we applied a learning rate decay strategy.
Fine-tuning in unseen data. We performed fine-tuning to generalize our model to the unseen VDCRA dataset. Subjects were separated into VDCRA-HC and VDCRA-MS cohorts ($n = 5$ each), and a separate network was fine-tuned for each. For each cohort, we generated 20 random $\mathbf{M}_{\text{dr}}$ masks per sample. Each model was fine-tuned using a single subject and validated with the remaining four using the same augmentations as before.

3. Validation of Inpainted Images

Since we had access to the original images without dropout, we used them for validation of the inpainted images. We calculated the L1 error, peak signal-to-noise ratio (PSNR), and structural similarity image metric (SSIM) between the inpainted images and the original ground truth. We evaluated the contributions of the input channels, the SPE, and various generator models in ablation studies. Table 1 shows these results for one validation fold.
Multi-modal input data. We first investigated the effects of each input channel. We found that as each input was introduced, all three metrics improved (Table 1-pink).
Sinusoidal positional encoding. We found that implementing our network with SPE also improved all three metrics compared to the model without SPE (Table 1-yellow).
Generator Architectures. We compared our proposed model to several 2D and 2.5D generators developed for inpainting tasks (Zhang et al., 2020; Suvorov et al., 2022). Table 1-blue shows that our model had the best PSNR and SSIM, while 2.5D LaMA (Suvorov et al., 2022) had the best L1.
Fine-tuning. We also assessed the affect of fine-tuning in the two VDCRA cohorts by comparing voxel intensity distributions in GM, WM, and CSF labels for the inpainted images before and after fine-tuning (Fig. 3). We found that performing fine-tuning led to a synthesized image histogram closer to the original image across all three tissue types, indicating a reduction in intensity variation between datasets.

Validation Experiment							T1w Model		
		Input Channels					L1 ↓	PSNR ↑	SSIM ↑
		$\phi(\mathbf{I}_{T1}, \mathbf{M}_{\text{dr}})$	$\mathbf{M}_{\text{tis}}$	$\mathbf{I}_{\text{tis}}$	$\phi(\mathbf{I}_{FL}, \mathbf{M}_{\text{dr}})$	$\mathbf{E}(\phi(\mathbf{I}_{T1}, \mathbf{M}_{\text{dr}}))$			
Input channel ablation	3D CNN	✓	✓				11.49	21.4	0.76
	3D CNN	✓	✓	✓			7.44	23.13	0.77
	3D CNN	✓	✓	✓	✓		7.29	23.46	0.83
	3D CNN	✓	✓	✓		✓	6.75	23.17	0.85
No SPE	3D CNN	✓	✓	✓	✓	✓	7.94	25.27	0.84
Generator architectures	3D CNN	✓	✓	✓	✓	✓	4.57	23.77	0.81
	2.5D CNN	✓	✓	✓	✓	✓	4.04	24.41	0.77
	2D LaMA	✓	✓	✓	✓	✓	3.47	25.67	0.84
	2.5D LaMA	✓	✓	✓	✓	✓	**2.97**	27.91	0.86
Ours	3D CNN	✓	✓	✓	✓	✓	5.37	**31.35**	**0.86**

Table 1: Quantitative results for validation of inpainted T1w images. This ablation study shows the contributions of the input channels, SPE as well as various generators.

4. Validation of the CSF Inpainting Task for Synthetic Cortical Atrophy

We next applied our model to the CSF inpainting task in images obtained using a registration-based synthetic atrophy (RBSA) method (Larson and Oguz, 2022). We compared FS cortical surfaces from the atrophied images before and after inpainting.

Synthetic atrophy induction. We created synthetically atrophied versions of both the T1w and FLAIR images using the RSBA pipeline (Larson and Oguz, 2022). We refer to the original and atrophied images as $\mathbf{I}_{*,\text{orig}}$ and $\mathbf{I}_{*,\text{atr}}$, where $*$ indicates T1w or FLAIR. For the Kirby data, we used the FS cortical parcellations (aparc+aseg.mgz) in one sample for each subject (the retest samples are not used in this experiment), and selected seven regions for atrophy induction. For VDCRA, similar to Larson et al., we created subject-specific label maps with ROIs surrounding each landmark cluster; this allowed us to also deform the landmarks. For all $\mathbf{I}_{*,\text{atr}}$, we generated tissue labellings $\mathbf{I}_{\text{tis,atr}}$ by taking the corresponding $\mathbf{I}_{\text{tis,orig}}$ and replacing the atrophied GM voxels with CSF.

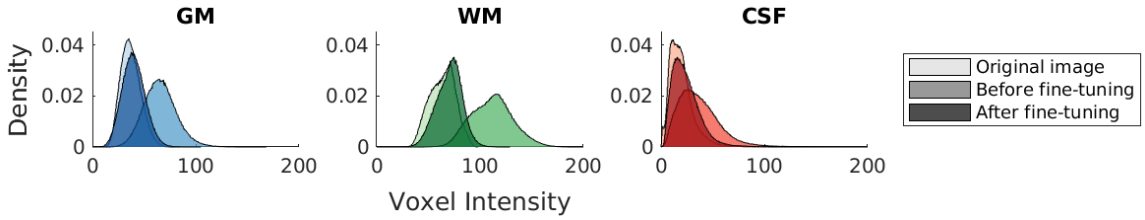

Figure 3: Voxel intensity distributions within GM, WM, and CSF in the VDCRA-HC cohort. The legend in grayscale applies to all three panels (lightest color: original images; medium: non-fine-tuned; darkest: fine-tuned).

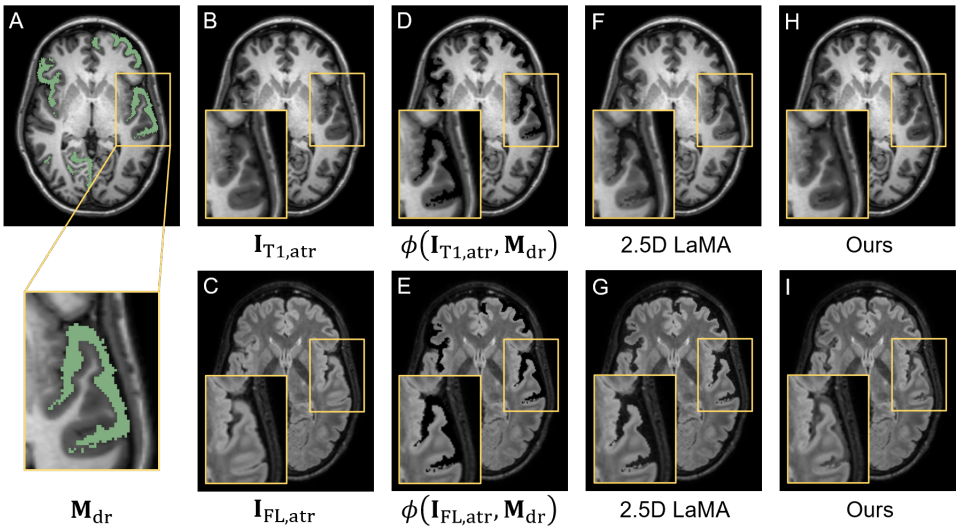

Figure 4: Representative CSF-inpainted T1w (top) and FLAIR (bottom) images. Our model reduces boundary blurring and creates realistic CSF even in the deep sulci.

Inpainting. Our model implementation for CSF inpainting differed slightly from training. First, rather than using random dropouts, we instead dropped out all CSF voxels affected by atrophy induction as our $\mathbf{M}_{\mathrm{dr}}$. This included (1) voxels that were classified as GM in $\mathbf{I}_{\mathrm{tis,orig}}$ but CSF in $\mathbf{I}_{\mathrm{tis,atr}}$, and (2) voxels surrounding the eroded GM. We performed inpainting using the model trained on the specific dataset and modality using the following five input channels: $\phi(\mathbf{I}_{\mathrm{T1,atr}}, \mathbf{M}_{\mathrm{dr}})$, $\phi(\mathbf{I}_{\mathrm{FL,atr}}, \mathbf{M}_{\mathrm{dr}})$, $\mathbf{I}_{\mathrm{tis,atr}}$, $\mathbf{M}_{\mathrm{dr}}$, and the edge map $\mathbf{E}(\phi(\mathbf{I}_{*,\mathrm{atr}}, \mathbf{M}_{\mathrm{dr}}))$. Apdx. A summarizes these changes.

To further improve the efficiency and accuracy of our model, we utilized a sliding window inference (details in Apdx. D) technique for this experiment. Each input channel was subdivided into overlapping patches of $96 \times 96 \times 96$ voxels, with a stride of 20 voxels between patches. Gaussian weights were applied to the patch edges to ensure smoothness of the generated result. This allowed us to generate high-quality, synthetic images while minimizing memory consumption compared to processing the entire image at once. Fig. 4 displays an example of the T1w and FLAIR images generated by our model alongside the non-inpainted images. As hypothesized, inpainting CSF alleviated the issue of blurring induced by RBSA methods around the GM/CSF interface, and also generated more realistic cortical thinning in narrow sulci. We also show images synthesized with the 2.5D LaMA generator, which did not create as realistic GM/CSF boundaries compared to our model, despite yielding better quantitative metrics in Sec. 3. A more detailed qualitative ablation study is available in Apdx. E.

Cortical thickness change. We evaluated the accuracy of the change in mean CT measured in the 7 atrophied regions in the Kirby dataset. Error was calculated as the absolute value of the difference between the measured thickness change and the true induced change. Fig. 5 displays these results, where we observe that our inpainting methods either significantly improved ($p < 0.05$, in paired t-tests) or had no significant effect on CT accuracy.

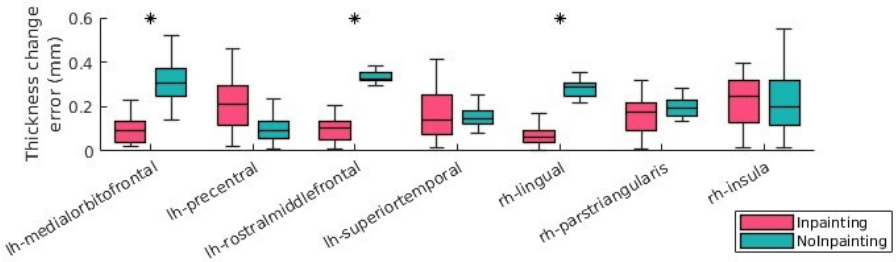

Figure 5: Mean error in measured CT change in each region selected for atrophy for the Kirby dataset. * indicates a significant difference ($p < 0.05$).

Cortical surface evaluation. We finally compared the quality of cortical surface reconstructions for the inpainted images generated using the VDCRA dataset with fine-tuning. We found that the surfaces generated using the inpainted images better matched the induced GM atrophy. Fig. 6 shows that FS often placed the surface for the atrophied image (green) outside that of the original image (pink); this is due to blurring of the GM/CSF boundaries caused by the RBSA method.

5. Discussion and Conclusion

We have presented a self-supervised, 3D GAN-based inpainting method that can realistically recreate multiple different tissue types in an MRI. We implemented 3D versions of several existing 2D methods including a patch-based dropout framework, edge detection for improved anatomical priors, and sinusoidal positional encoding. We showed that with fine-tuning, we can successfully apply our network to unseen data. Finally, we showed that our method improves the results of registration-based synthetic atrophy methods, and that the output can be successfully used for accuracy validation of cortical surface analyses.

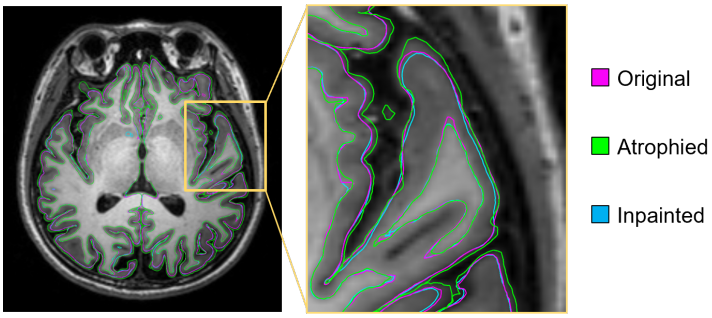

Figure 6: FS cortical surfaces overlaid on the original (unatrophied) image. Inpainting creates more realistic GM surfaces that match the atrophy (green GM surface is outside pink GM surface, but blue GM surface is correctly inside pink surface).

Acknowledgments

Acknowledgments withheld.

References

Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009.

Jose Bernal, Sergi Valverde, Kaisar Kushibar, Mariano Cabezas, Arnau Oliver, and Xavier Lladó. Generating Longitudinal Atrophy Evaluation Datasets on Brain Magnetic Resonance Images Using Convolutional Neural Networks and Segmentation Priors. *Neuroinformatics*, 19(3):477–492, 7 2021. ISSN 15590089. doi: 10.1007/S12021-020-09499-Z/FIGURES/6.

John Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6):679–698, 1986.

M Jorge Cardoso, Wenqi Li, Richard Brown, Nic Ma, Eric Kerfoot, Yiheng Wang, Benjamin Murrey, Andriy Myronenko, Can Zhao, Dong Yang, et al. Monai: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701*, 2022.

Anders M. Dale, Bruce Fischl, and Martin I. Sereno. Cortical Surface-Based Analysis I. Segmentation and Surface Reconstruction. *NeuroImage*, 9(2):195–207, 1999.

Christos Davatzikos, Ahmet Genc, Dongrong Xu, and Susan Resnick. Voxel-Based Morphometry Using the RAVENS Maps: Methods and Validation Using Simulated Longitudinal Atrophy. *NeuroImage*, 14(6):1361–1369, 2001. ISSN 1053-8119. doi: 10.1006/NIMG.2001.0937.

Bruce Fischl and Anders M. Dale. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proceedings of the National Academy of Sciences of the United States of America*, 97(20):11050–11055, 2000. ISSN 00278424. doi: 10.1073/pnas.200033797.

Greg M. Fleishman, Alessandra Valcarcel, Dzung L. Pham, Snehashis Roy, Peter A. Calabresi, Paul Yushkevich, Russell T. Shinohara, and Ipek Oguz. Joint intensity fusion image synthesis applied to multiple sclerosis lesion segmentation. In Alessandro Crimi, Spyridon Bakas, Hugo Kuijf, Bjoern Menze, and Mauricio Reyes, editors, *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, pages 43–54. Springer International Publishing, 2018. ISBN 978-3-319-75238-9.

Peter A. Freeborough and Nick C. Fox. The boundary shift integral: An accurate and robust measure of cerebral volume changes from registered repeat MRI. *IEEE Transactions on Medical Imaging*, 16(5):623–629, 1997. ISSN 02780062. doi: 10.1109/42.640753.

Xiao Han, Dzung L. Pham, Duygu Tosun, Maryam E. Rettmann, Chenyang Xu, and Jerry L. Prince. CRUISE: Cortical reconstruction using implicit surface evolution. *NeuroImage*, 23(3):997–1012, 2004. ISSN 10538119. doi: 10.1016/j.neuroimage.2004.06.043.

Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):1–14, 2017.

Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.

Bilge Karaçali and Christos Davatzikos. Simulation of Tissue Atrophy Using a Topology Preserving Transformation Model. *IEEE TRANSACTIONS ON MEDICAL IMAGING*, 25(5), 2006. doi: 10.1109/TMI.2006.873221.

Bishesh Khanal, Nicholas Ayache, and Xavier Pennec. Simulating longitudinal brain MRIs with known volume changes and realistic variations in image intensity. *Frontiers in Neuroscience*, 11(MAR):132, 3 2017. ISSN 1662453X. doi: 10.3389/FNINS.2017.00132/BIBTEX.

Bennett A. Landman, Alan J. Huang, Aliya Gifford, Deepti S. Vikram, Issel Anne L. Lim, Jonathan A.D. Farrell, John A. Bogovic, Jun Hua, Min Chen, Samson Jarso, Seth A. Smith, Suresh Joel, Susumu Mori, James J. Pekar, Peter B. Barker, Jerry L. Prince, and Peter C.M. van Zijl. Multi-parametric neuroimaging reproducibility: A 3-T resource study. *NeuroImage*, 54(4):2854–2866, 2011. ISSN 10538119. doi: 10.1016/j.neuroimage.2010.11.047.

Kathleen E Larson and Ipek Oguz. Synthetic atrophy for longitudinal surface-based cortical thickness measurement. In Bennett A. Landman and Ivana Išgum, editors, *Medical Imaging 2021: Image Processing*, volume 11596, page 125. SPIE, 2 2021. ISBN 9781510640214. doi: 10.1117/12.2580907.

Kathleen E. Larson and Ipek Oguz. Synthetic Atrophy for Longitudinal Cortical Surface Analyses. *Frontiers in Neuroimaging*, 6 2022. ISSN 2813-1193. doi: 10.3389/FNIMG.2022.861687.

Haodong Li, Weiqi Luo, and Jiwu Huang. Localization of diffusion-based inpainting in digital images. *IEEE transactions on information forensics and security*, 12(12):3050–3064, 2017.

Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Z Qureshi, and Mehran Ebrahimi. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212*, 2019.

Ipek Oguz and Milan Sonka. LOGISMOS-B: Layered Optimal graph image segmentation of multiple objects and surfaces for the brain. *IEEE Transactions on Medical Imaging*, 2014a. ISSN 1558254X. doi: 10.1109/TMI.2014.2304499.

Ipek Oguz and Milan Sonka. Robust Cortical Thickness Measurement with LOGISMOS-B. In *Medical Image Computing and Computer-Assisted Intervention –MICCAI 2014*, pages 722–730. Springer International Publishing Switzerland, 2014b.

Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.

Yurui Ren, Xiaoming Yu, Ruonan Zhang, Thomas H Li, Shan Liu, and Ge Li. Structureflow: Image inpainting via structure-aware appearance flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 181–190, 2019.

Martin Reuter, Nicholas J Schmansky, H Diana Rosas, and Bruce Fischl. Within-subject template estimation for unbiased longitudinal image analysis. *NeuroImage*, 61(4):1402–1418, 2012. doi: 10.1016/j.neuroimage.2012.02.084.

Navid Shiee, Pierre-Louis Bazin, Jennifer L Cuzzocreo, Chuyang Ye, Bhaskar Kishore, Aaron Carass, Peter A Calabresi, Daniel S Reich, Jerry L Prince, and Dzung L Pham. Reconstruction of the Human Cerebral Cortex Robust to White Matter Lesions: Method and Validation. *Human Brain Mapping2*, 35:3385–3401, 2014. doi: 10.1002/hbm.22409.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2014. URL https://arxiv.org/abs/1409.1556.

Yuhang Song, Chao Yang, Yeji Shen, Peng Wang, Qin Huang, and C-C Jay Kuo. Spg-net: Segmentation prediction and guidance network for image inpainting. *arXiv preprint arXiv:1805.03356*, 2018.

Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2149–2159, 2022.

Yucheng Tang, Dong Yang, Wenqi Li, Holger R Roth, Bennett Landman, Daguang Xu, Vishwesh Nath, and Ali Hatamizadeh. Self-supervised pre-training of swin transformers for 3d medical image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20730–20740, 2022.

Tian Xia, Agisilaos Chartsias, Chengjia Wang, and Sotirios A. Tsaftaris. Learning to synthesise the ageing brain without longitudinal data. *Medical Image Analysis*, 73, 12 2019. doi: 10.1016/j.media.2021.102169.

Rui Xu, Xintao Wang, Kai Chen, Bolei Zhou, and Chen Change Loy. Positional encoding as spatial inductive bias in gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13569–13578, 2021.

Huahong Zhang, Rohit Bakshi, Francesca Bagnato, and Ipek Oguz. Robust multiple sclerosis lesion inpainting with edge prior. In *International Workshop on Machine Learning in Medical Imaging*, pages 120–129. Springer, 2020.

Yuqian Zhou, Connelly Barnes, Eli Shechtman, and Sohrab Amirghodsi. Transfill: Reference-guided image inpainting by merging multiple color and spatial transformations.

In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2266–2276, 2021.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.

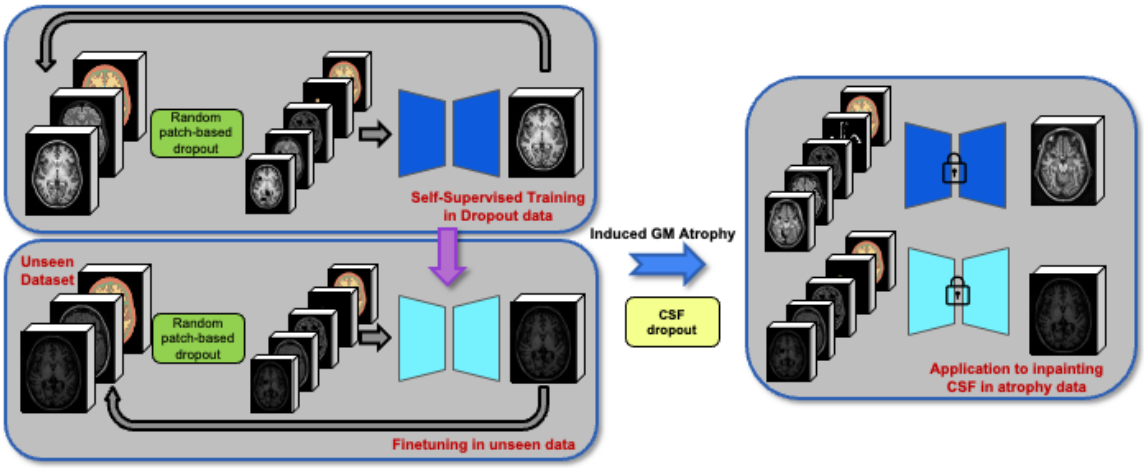

Figure A.1: Our training workflow consists of 3 main stages. In the first stage, we train a self-supervised model using random dropouts. In the second stage, we perform fine-tuning to unseen datasets. In the third stage, we freeze the model weights and apply to the CSF inpainting task.

Appendix A. Overall Pipeline: Training, Fine-tuning, and Application to CSF Inpainting

Our training pipeline consists of three distinct stages, as summarized in Figure A.1.

In the first stage, we train a generative model using a self-supervised approach. The model is capable of generating a variety of tissue types, given a dropout mask. We do this in a 5-fold cross-validation setting.

In the second stage, when dealing with an unseen dataset, we un-freeze the weights of the first phase and fine-tune the model on the new dataset. This is done by randomly sampling from the dataset and generating multiple different dropout masks.

In the final stage, using either of these models, the weights of the model are frozen and the model is applied to atrophied data with a cerebrospinal fluid (CSF) dropout, as described in Section 4. Here, we use the GM atrophy algorithm proposed in (Larson and Oguz, 2021) to induce GM atrophy in the synthetic tissue images.

Appendix B. Network Architecture

B.1. Generator

Our generator structure is shown in Fig. A.2. We use 5 different configurations of blocks, each represented by a different color. The basic building block of our generator is a 3D convolutional layer, followed by Spectral Normalization, Instance Normalization, and ReLU activation, as proposed in (Johnson et al., 2016). The kernel size of the convolutional layers varies for each block, with the blue block utilizing a kernel size of $7 \times 7 \times 7$ with stride 1, and the orange block utilizing a kernel size of $4 \times 4 \times 4$ with stride 2 for downsampling. Additionally, we include 8 residual blocks of the first layer, each utilizing a $3 \times 3 \times 3$

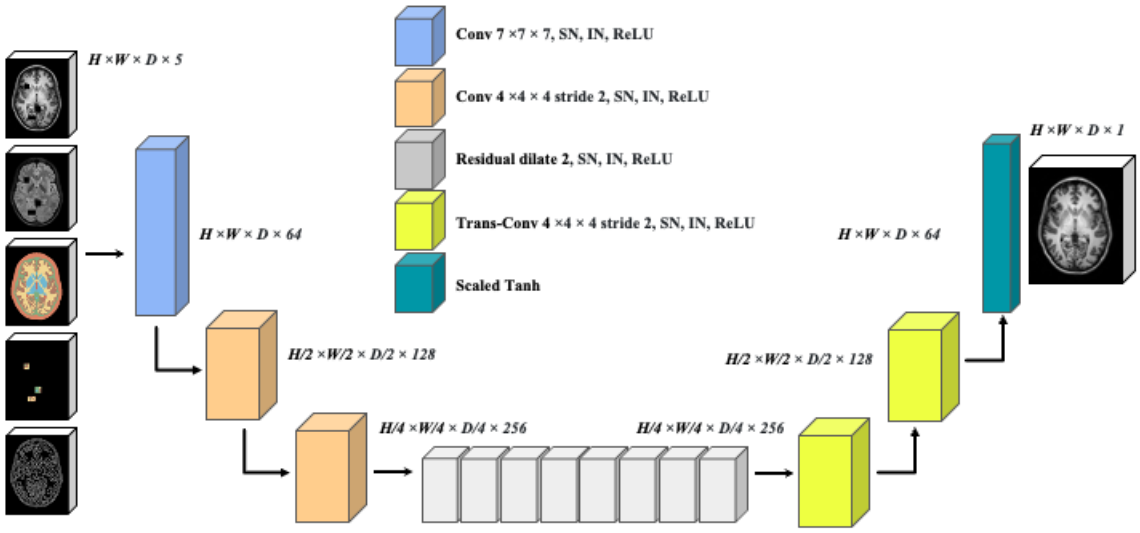

Figure A.2: Generator structure.

convolutional layer with dilation of 2. The yellow block is the transposed version of the orange block, used for upsampling. Our generator output is produced via an activation layer, shown in teal.

B.2. Discriminator

The architecture of our proposed discriminator is shown in Fig. A.3 and is based on the 70×70 PatchGAN (Zhu et al., 2017). The discriminator consists of two distinct block configurations. The violet blocks are comprised of a 4x4x4 convolutional layer, followed by Spectral Normalization and LeakyReLU activation, with a stride of 2. The navy blocks, on the other hand, are also comprised of a $4 \times 4 \times 4$ convolutional layer, Spectral Normalization and LeakyReLU activation, but with a stride of 1. Each block is stacked on top of each other to form the overall discriminator architecture.

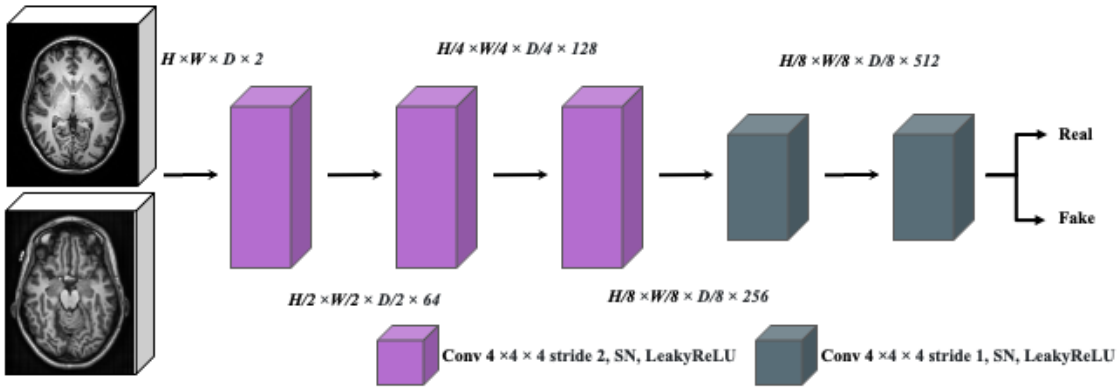

Figure A.3: Discriminator structure.

Transform	Parameters	p
Intensity Normalization	[0-1]	1
Foreground Crop	N/A	1
Random Crop By Pos/Neg Label	dim=(48,48,48), #=4,neg=0	1
Random Flip	axis 0	0.1
Random Flip	axis 1	0.1
Random Flip	axis 2	0.1
Random Rotate 90	$k = 3$	0.1

Table A.1: Data augmentation during our training process. **p** indicates the probability of applying an augmentation.

Appendix C. Data Augmentation Strategy Details

In our proposed model, we use various data augmentation techniques during the training process to increase the robustness of our model. The specific configurations for these augmentations are listed in Table A.1, which details the parameters and probabilities used. It is important to note that these augmentations are excluded during the validation phase. We apply the intensity normalization first by mapping the intensity range $[0, 255]$ to $[0, 1]$. We also include Foreground Crop, Random Crop, Random Flip, and Random Rotate 90, and we followed the implementation of these augmentations provided in MONAI library (Cardoso et al., 2022).

Appendix D. Sliding Window Inference

In this work, we implement a sliding window inference technique to address the limitation of GPU memory when processing large 3D medical image data. The implementation is depicted in Fig. A.4.

Specifically, we adopt the sliding window patch approach used in MONAI (Cardoso et al., 2022), using a patch size of $96 \times 96 \times 96$ pixels. This configuration strikes a balance between GPU consumption and the field of view (FOV) of the input data. The gap between each patch is set at 20 pixels.

After generating the inpainted image of each patch, we concatenate adjacent patches by applying a Gaussian weighting to obtain the final output to avoid seam artifacts. This approach allows us to process large 3D medical images while maintaining a high level of detail in the inpainted regions.

Appendix E. Detailed Qualitative Comparison Results for CSF Inpainting Task

In this section, we present detailed qualitative comparison results for each model from the ablation study detailed in Sec. 3 when applied to the CSF inpainting task. Figures A.5 and A.6 display example Kirby T1w and FLAIR image results, respectively. These images show that inpainting decreases the blurring affect induced by RBSA methods, as hypothesized. We found that although the 2D and 2.5D generators were able to produce CSF in atrophied

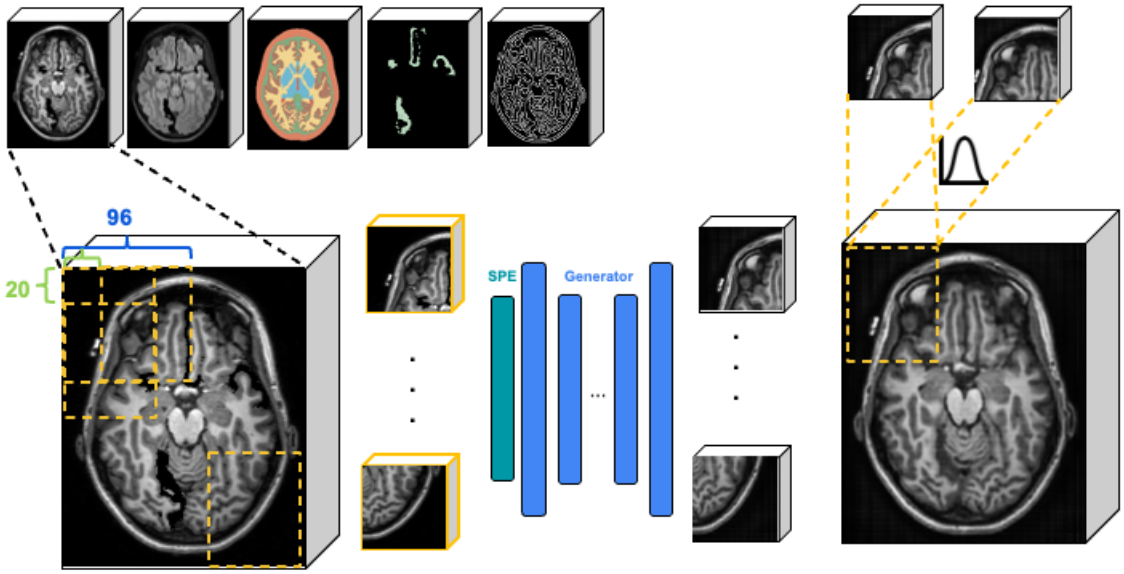

Figure A.4: An example of sliding window inference for generating a T1w image. The right top panel shows the application of Gaussian weights to generate the fused image output.

regions, they were not able to recreate realistic atrophy, and instead resulted in jagged edges and discontinuities along tissue boundaries as well as unrealistic tissue contrast. On the other hand, our method produced an image with cleaner GM/CSF boundaries, more similar to real data. We also found that the use of SPE contributed to the synthesis of these more realistic images.

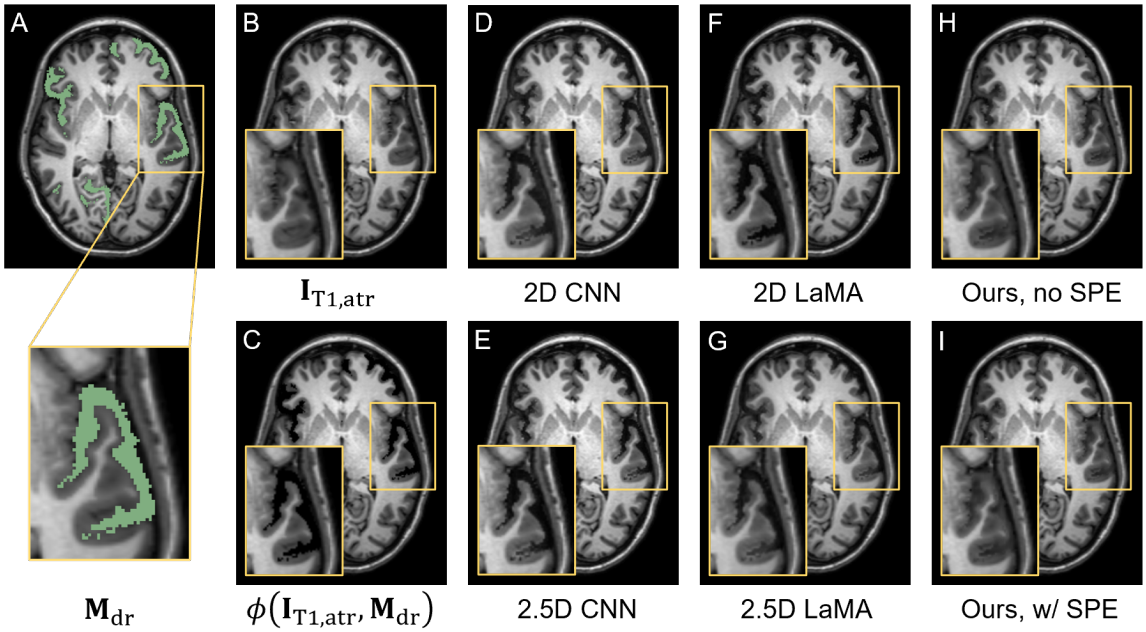

Figure A.5: T1w CSF inpainting results (Kirby). Our model synthesizes realistic images, while the other models suffer from jagged edges and discontinuities.

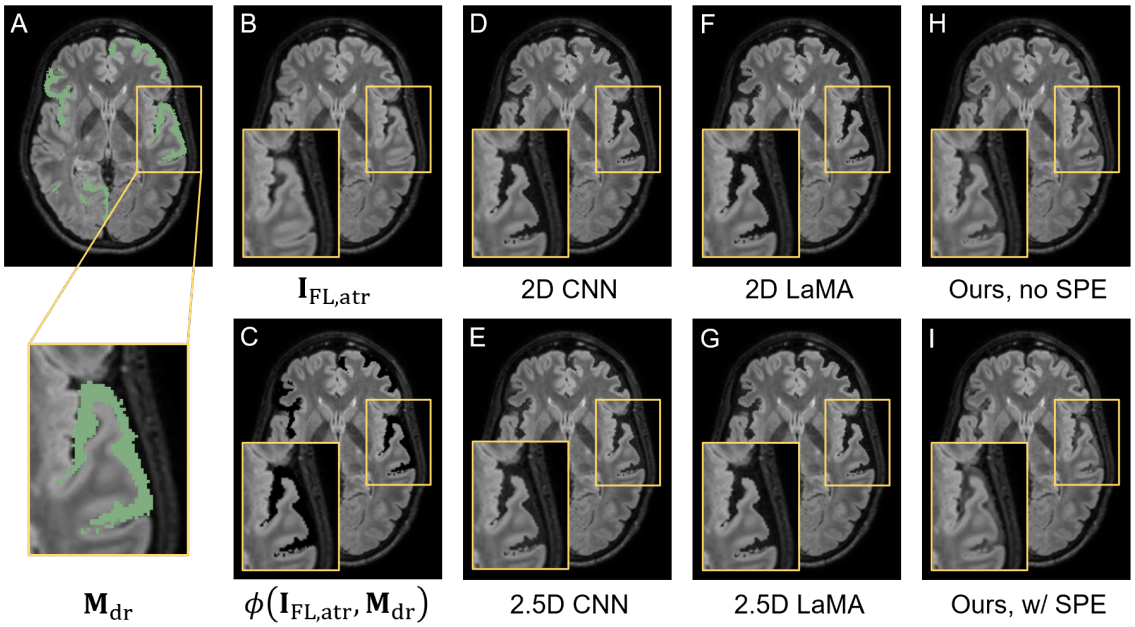

Figure A.6: FLAIR CSF inpainting results (Kirby). Our model synthesizes realistic images, while the other models suffer from jagged edges and discontinuities.