# Gradual modality dropout for segmenting ischemic stroke lesions in an unseen center with missing modalities

Sofia Vargas-Ibarra<sup>1</sup> Vincent Vigneron<sup>1</sup> Hichem Maaref<sup>1</sup> <sup>1</sup>Université Évry Paris-Saclay, IBISC, Évry, France.

SOFIA.VARGASIBARRA@UNIV-EVRY.FR VINCENT.VIGNERON@UNIV-EVRY.FR HICHEM.MAAREF@UNIV-EVRY.FR

Sonia Garcia-Salicetti<sup>2</sup> SONIA.GARCIA@TELECOM-SUDPARIS.EU <sup>2</sup>SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France.

Andreia Faria<sup>3</sup> <sup>3</sup>Johns Hopkins University, Baltimore, USA.

Editors: Accepted for publication at MIDL 2025

## Abstract

In clinical practice, imaging modalities may not always be available for every patient due to scheduling, cost, or patient-specific constraints. Additionally, multi-center imaging studies often face inconsistencies in protocols, machine settings, and artifacts, compromising data quality. We propose a 3D U-Net model for ischemic lesion segmentation using a novel training technique, gradual modality dropout, which progressively deactivates imaging modalities during training. This approach ensures robust performances when all modalities are present and improves segmentation accuracy in scenarios where one or more modalities are missing in unfamiliar contexts. The model demonstrates adaptability and reliability when trained on MRI scans of stroke patients across different phases (hyper-acute, sub-acute, acute, and post-treatment) and various hospital settings. Code available here: https://github.com/sofiavarib/Gradual-modality-dropout

Keywords: Stroke, segmentation, MRI modalities, lesion, missing modalities

# 1. Introduction

Stroke remains a leading cause of death globally, with millions of fatalities each year. It results from arterial blockages that reduce blood flow, causing irreversible brain damage. Accurate identification of the affected brain region (lesion) is vital for treatment decisions. While lesion segmentation is critical in improving survival and recovery, manual analysis is time-consuming and limits timely intervention.

To address this, various deep learning models have been developed to automate and accelerate lesion detection. These include 3D CNN-based models like 3D U-Net (Ronneberger et al., 2015; Omarov et al., 2022; Ashtari et al., 2023), hybrid 2D/3D architectures such as D-Unet (Zhou et al., 2021) and DFENet (Basak et al., 2021), patch-based sampling (Xue et al., 2020; Alquhayz et al., 2022), and attention-based methods like AABTS-Net (Tian et al., 2023) and PerfUnet (de Vries et al., 2023). However, many rely on the presence of specific imaging modalities. To overcome this, approaches like the Unified Representation Network (Lau et al., 2019) enable flexible input handling, while GANs have been used to generate missing modalities (Sharma and Hamarneh, 2019). Modality dropout strategies have also proven effective, with MultiUnet (Xu et al., 2024) and ModDrop+ (Liu et al., 2022) offering strong performance, the latter introducing a dynamic filter-scaling head.

#### afaria1@jhmi.edu

# 2. Datasets

This study uses two multicentric and multiequipement datasets: JHU (Liu et al., 2023), which consists of 1888 patients with acute ischemic strokes, and ISLES (Hernandez Petzsche et al., 2022), which has 150 patients. CHSF and MATAR datasets (Marnat et al., 2023) consist only of hyper-acute strokes from a single hospital and machine. CHSF has 65 patients with proximal occlusion and MATAR 125 patients with distal ones. All of them have Apparent diffusion coefficient (ADC), magnetic field strength (B0), and diffusion-weighted imaging (DWI), except for ISLE, S where B0 is missing.

#### 3. Method

Inspired by the dropout technique (Srivastava et al., 2014), modality dropout (Lau et al., 2019) extends this concept to input modalities (Lau et al., 2019): instead of deactivating individual neurons, it zeroes out entire input images with a Bernoulli probability p, effectively simulating missing modalities and improving model robustness to incomplete data. We propose a gradual transition mechanism to replace the abrupt disruptions caused by completely dropping a modality, and evaluate its performance using the nnU-Net benchmark (Isensee et al., 2018). For a segmentation model F that processes input modalities  $\mathbf{x}_j, j \in \{1, \ldots, n\}$  and n the total number of modalities, the final output  $\mathbf{y}$  is computed as:

$$\mathbf{y} = F(\widetilde{\mathbf{r}} \odot \mathbf{x}) \quad \text{with} \quad \widetilde{\mathbf{r}}_j = \begin{cases} 1, & \text{if } \mathbf{r}_j = 1, \text{with probability } p \\ g_j(t), & \text{otherwise, with probability } 1 - p \end{cases}$$
(1)

with  $\mathbf{r}_j \sim \text{Bernouilli}(p)$  and ' $\odot$ ' the Hadamard product. When p = 1, no dropout is applied, and all modalities are used, whereas when p = 0, the corresponding modality is gradually dropout, controlled by the gradual function g(t). This decreasing function reaches zero at some time t and we propose the following g(t):

$$g_{j}(t) = \begin{cases} 0.75 + \epsilon, & \text{if } t < 0.25T \\ 0.5 + \epsilon, & \text{if } t < 0.5T \\ 0.25 + \epsilon, & \text{if } t < 0.75T \\ 0 + \epsilon, & \text{otherwise} \end{cases}$$
(2)

where T represents the total number of epochs, and a noise  $\epsilon$  is added from a normal distribution with a mean of 0 and a standard deviation of 0.01 to the function  $g_j(t)$ . Serving both as data augmentation, where  $\tilde{\mathbf{r}} \odot \mathbf{x}$  becomes an enhanced version of  $\mathbf{x}$ , and as regularization, it allows the model to adapt to missing modality data gradually.

### 4. Results

In clinical practice, a dataset from a new center can create a missing modality scenario. To mimic this case, all datasets are used during training except ISLES, where B0 is absent. Using ADC, DWI, and B0 as inputs for nnUnet, we apply modality dropout on B0 with varying probabilities and evaluate performance with all modalities available and under missing-data conditions. The results are shown in Table 1, where our proposed method

**Table 1:** The evaluation considers missing B0 on the unseen ISLES dataset, using modality dropout with probability p in gradual and non-gradual settings. All training configurations use ADC, DWI, and B0 as inputs. The upper bound (\*) includes ISLES in training, while the lower bound (†) excludes it, both without any method. The best Dice scores for unseen datasets and average results under missing-modality conditions are highlighted and the best without absent data are underlined. Evaluation symbols:  $\bullet$  for all modalities present, and  $\bigcirc$  when B0 is absent.

Exp.	Mod. drop		B0	Unseen	Test set $(T_0)$ (Dice)				
num	Type	1-p		ISLES	CHSF	MATAR	ISLES (unseen)	JHU	Average
Exp1	-	-		1	$0.790 \pm 0.029$	$0.768 {\pm} 0.061$	-	$0.731{\pm}0.012$	$0.763 {\pm} 0.059$
Exp1	-	-	$\circ$	1	$0.787 \pm 0.027$	$0.767 {\pm} 0.065$	$0.622^{\dagger} \pm 0.040$	$0.723{\pm}0.014$	$0.725 {\pm} 0.055$
Exp2	MultiUnet	0.2		1	$0.787 \pm 0.027$	$0.762{\pm}0.056$	-	$0.750{\pm}0.012$	$0.767 {\pm} 0.061$
Exp2	MultiUnet	0.2	$\circ$	1	$0.788 \pm 0.028$	$0.765 {\pm} 0.056$	$0.672 {\pm} 0.039$	$0.748 {\pm} 0.013$	$0.743{\pm}0.051$
Exp3	Gradual	0.2	۲	1	$0.789 \pm 0.028$	$0.757 {\pm} 0.061$	-	$0.744{\pm}0.013$	$0.763 {\pm} 0.061$
Exp3	Gradual	0.2	$\circ$	1	$0.789 \pm 0.028$	$0.758 {\pm} 0.061$	$0.650 {\pm} 0.040$	$0.743 {\pm} 0.013$	$0.735{\pm}0.052$
Exp4	ModDrop+	0.2		1	$0.787 \pm 0.028$	$0.733 {\pm} 0.066$	-	$0.740{\pm}0.014$	$0.753 {\pm} 0.066$
Exp4	ModDrop+	0.2	0	1	$0.787 \pm 0.029$	$0.732{\pm}0.065$	$0.617{\pm}0.041$	$0.740{\pm}0.014$	$0.719{\pm}0.056$
Exp5	MultiUnet	0.5		1	$0.789 \pm 0.028$	$0.775 {\pm} 0.056$	-	$0.748 {\pm} 0.013$	$0.771 {\pm} 0.058$
Exp5	MultiUnet	0.5	0	1	$0.789 \pm 0.028$	$0.777 {\pm} 0.055$	$0.637 {\pm} 0.040$	$0.747{\pm}0.013$	$0.738 {\pm} 0.051$
Exp6	Gradual	0.5		1	$0.786 \pm 0.029$	$0.766 {\pm} 0.055$	-	$0.747 {\pm} 0.013$	$0.766 {\pm} 0.058$
Exp6	Gradual	0.5	0	1	$0.786 \pm 0.029$	$0.763 {\pm} 0.056$	$0.655 {\pm} 0.039$	$0.748 {\pm} 0.012$	$0.738 {\pm} 0.050$
Exp7	ModDrop+	0.5		1	$0.779 \pm 0.029$	$0.758 {\pm} 0.070$	-	$0.744{\pm}0.013$	$0.760{\pm}0.066$
Exp7	ModDrop+	0.5	0	1	$0.778 \pm 0.029$	$0.761 {\pm} 0.070$	$0.630 {\pm} 0.040$	$0.743 {\pm} 0.013$	$0.728 {\pm} 0.056$
Exp8	MultiUnet	0.8		1	$0.790 \pm 0.026$	$0.743 {\pm} 0.068$	-	$0.751 {\pm} 0.013$	$0.761 {\pm} 0.065$
Exp8	MultiUnet	0.8	0	1	$0.790 \pm 0.027$	$0.741{\pm}0.068$	$0.682{\pm}0.032$	$0.751 {\pm} 0.013$	$0.741{\pm}0.050$
Exp9	Gradual	0.8		1	$0.810 \pm 0.025$	$0.750{\pm}0.055$	-	$0.743 {\pm} 0.013$	$0.768 {\pm} 0.057$
Exp9	Gradual	0.8	0	1	$0.810{\pm}0.026$	$0.752{\pm}0.056$	$0.701{\pm}0.038$	$0.743 {\pm} 0.013$	$0.749{\pm}0.048$
Exp10	ModDrop+	0.8		1	$0.788 \pm 0.026$	$0.763 {\pm} 0.060$	-	$0.753 {\pm} 0.012$	$0.768 {\pm} 0.058$
Exp10	ModDrop+	0.8	0	1	$0.788 \pm 0.027$	$0.761 {\pm} 0.060$	$0.635 {\pm} 0.038$	$0.753{\pm}0.012$	$0.734{\pm}0.051$
Exp11	-	-			$0.773 \pm 0.031$	$0.753 {\pm} 0.056$	-	$0.756 {\pm} 0.012$	$0.761 {\pm} 0.057$
Exp11	-	-	0		$0.775 \pm 0.025$	$0.680 {\pm} 0.058$	$0.782^* {\pm} 0.029$	$0.756 {\pm} 0.012$	$0.747{\pm}0.041$

(Gradual) is compared with MultiUnet, where the modality dropout is done without a gradual function, zeroing abruptly the modalities and with ModDrop+, where a dynamic filter is included. Some visual predictions are shown in Appendix A.

Except for ModDrop+ with (1-p) = 0.2, all methods allow to have an improvement on the unseen dataset ISLES with respect to the lower bound (which corresponds to training without ISLES and no dropout (Exp1)). The best result is obtained with our proposed method, arriving at a Dice of 0.701 on ISLES (Exp 9), but still staying far from the upper bound, which includes ISLES during training with a black image instead of B0 input (Exp11, Dice = 0.782). Seeing the average Dice scores under missing modalities conditions, the best average score (Exp9, Dice 0.749) gets close to the upper bound (Exp11, Dice 0.747), without a performance reduction under the non-missing scenario, showing the method's robustness.

#### 5. Conclusion

The proposed method includes modality dropout in a smoother way, where a gradual function reduces the input contrast over the training. Thanks to the regularizing and augmentation aspect of the method and without reducing the performance in the seen centers either under a missing or not missing scenario, it outperforms the other Dice score' methods on an unseen center, where a modality is absent. Gradual modality dropout is model-agnostic, being adaptable to any architecture under any training schedule.

# References

- Hani Alquhayz, Hafiz Zahid Tufail, and Basit Raza. The multi-level classification network (mcn) with modified residual u-net for ischemic stroke lesions segmentation from atlas. *Computers in Biology and Medicine*, 151:106332, 2022.
- Pooya Ashtari, Diana M. Sima, Lieven De Lathauwer, Dominique Sappey-Marinier, Frederik Maes, and Sabine Van Huffel. Factorizer: A scalable interpretable approach to context modeling for medical image segmentation. *Medical Image Analysis*, 84:102706, feb 2023. doi: 10.1016/j.media.2022.102706. URL https://doi.org/10.1016%2Fj.media.2022. 102706.
- Hritam Basak, Rukhshanda Hussain, and Ajay Rana. DFENet: A novel dimension fusion edge guided network for brain MRI segmentation. SN Computer Science, 2 (6), aug 2021. doi: 10.1007/s42979-021-00835-x. URL https://doi.org/10.1007% 2Fs42979-021-00835-x.
- Lucas de Vries, Bart J Emmer, Charles BLM Majoie, Henk A Marquering, and Efstratios Gavves. Perfu-net: Baseline infarct estimation from ct perfusion source data for acute ischemic stroke. *Medical Image Analysis*, 85:102749, 2023.
- Moritz R Hernandez Petzsche, Ezequiel de la Rosa, Uta Hanning, Roland Wiest, Waldo Valenzuela, Mauricio Reyes, Maria Meyer, Sook-Lei Liew, Florian Kofler, Ivan Ezhov, et al. Isles 2022: A multi-center magnetic resonance imaging stroke lesion segmentation dataset. *Scientific data*, 9(1):762, 2022.
- Fabian Isensee, Jens Petersen, Andre Klein, David Zimmerer, Paul F Jaeger, Simon Kohl, Jakob Wasserthal, Gregor Koehler, Tobias Norajitra, Sebastian Wirkert, et al. nnu-net: Self-adapting framework for u-net-based medical image segmentation. arXiv preprint arXiv:1809.10486, 2018.
- Kenneth Lau, Jonas Adler, and Jens Sjölund. A unified representation network for segmentation with missing modalities. arXiv preprint arXiv:1908.06683, 2019.
- Chin-Fu Liu, Richard Leigh, Brenda Johnson, Victor Urrutia, Johnny Hsu, Xin Xu, Xin Li, Susumu Mori, Argye E Hillis, and Andreia V Faria. A large public dataset of annotated clinical mris and metadata of patients with acute stroke. *Scientific Data*, 10(1):548, 2023.
- Han Liu, Yubo Fan, Hao Li, Jiacheng Wang, Dewei Hu, Can Cui, Ho Hin Lee, Huahong Zhang, and Ipek Oguz. Moddrop++: A dynamic filter network with intra-subject cotraining for multiple sclerosis lesion segmentation with missing modalities. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 444–453. Springer, 2022.
- Gaultier Marnat, Bertrand Lapergue, Benjamin Gory, Maeva Kyheng, Julien Labreuche, Guillaume Turc, Stephanze Olindo, Igor Sibon, Jildaz Caroff, Didier Smadja, et al. Intravenous thrombolysis with tenecteplase versus alteplase combined with endovascular treatment of anterior circulation tandem occlusions: A pooled analysis of etis and tetris. *European Stroke Journal*, page 23969873231206894, 2023.

- Batyrkhan Omarov, Azhar Tursynova, Octavian Postolache, Khaled Gamry, Aidar Batyrbekov, Sapargali Aldeshov, Zhanar Azhibekova, Marat Nurtas, Akbayan Aliyeva, and Kadrzhan Shiyapov. Modified unet model for brain stroke lesion segmentation on computed tomography images. *Computers, Materials & Continua*, 71(3), 2022.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. CoRR, abs/1505.04597, 2015. URL http://arxiv.org/ abs/1505.04597.
- Anmol Sharma and Ghassan Hamarneh. Missing mri pulse sequence synthesis using multimodal generative adversarial network, 2019. URL https://arxiv.org/abs/1904.12200.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal* of machine learning research, 15(1):1929–1958, 2014.
- Weiwei Tian, Dengwang Li, Mengyu Lv, and Pu Huang. Axial attention convolutional neural network for brain tumor segmentation with multi-modality mri scans. *Brain Sciences*, 13(1), 2023. ISSN 2076-3425. doi: 10.3390/brainsci13010012. URL https://www.mdpi.com/2076-3425/13/1/12.
- Wentian Xu, Matthew Moffat, Thalia Seale, Ziyun Liang, Felix Wagner, Daniel Whitehouse, David Menon, Virginia Newcombe, Natalie Voets, Abhirup Banerjee, and Konstantinos Kamnitsas. Feasibility and benefits of joint learning from mri databases with different brain diseases and modalities for segmentation. In Ninon Burgos, Caroline Petitjean, Maria Vakalopoulou, Stergios Christodoulidis, Pierrick Coupe, Hervé Delingette, Carole Lartizien, and Diana Mateus, editors, Proceedings of The 7nd International Conference on Medical Imaging with Deep Learning, volume 250 of Proceedings of Machine Learning Research, pages 1771–1784. PMLR, 03–05 Jul 2024. URL https://proceedings.mlr. press/v250/xu24a.html.
- Yunzhe Xue, Fadi G Farhat, Olga Boukrina, AM Barrett, Jeffrey R Binder, Usman W Roshan, and William W Graves. A multi-path 2.5 dimensional convolutional neural network system for segmenting stroke lesions in brain mri images. *NeuroImage: Clinical*, 25:102118, 2020.
- Yongjin Zhou, Weijian Huang, Pei Dong, Yong Xia, and Shanshan Wang. D-unet: A dimension-fusion u shape network for chronic stroke lesion segmentation. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 18(3):940–950, 2021. doi: 10.1109/TCBB.2019.2939522.

# Appendix A. Out-domain evaluation with missing modalities - visual examples

In Fig. 1, the predictions of several models and the groundtruth are represented over the DWI image for ISLES dataset. All models have DWI, ADC, and missing B0 as input (black image). For the first column, the lower bound model is used, when ISLES is not seen during training, the second one is the best model obtained by adding the proposed method (1 - p=0.8), and the third column is when the model uses ISLES during training. Finally, in the last column, the groundtruth of the model is shown. In the first line, when the dataset is unseen and no method to manage the missing modality is used, some parts of the lesion are not detected. For the second case, false positives are obtained if the method is not applied. Finally, in the last case, we can see how seeing or not seeing ISLES, there are parts of the lesion that are still not detected.



Figure 1: Lesion segmentation from ISLES dataset using ADC, DWI and missing B0 (black image)