American Journal of Machine Learning

australiansciencejournals.com/ml

E-ISSN: 2689-0003

Volume 7 issue 1 2026

FastMRI-Guided Mamba UNet for Real-Time Emergency Brain Lesion Detection

Julien Martin
Department of Computer Science, École Polytechnique, 91128 Palaiseau, France
Claire Dubois
Department of Computer Science, École Polytechnique, 91128 Palaiseau, France
Mathieu Laurent
Faculty of Science and Engineering, Sorbonne University, 75005 Paris, France
Sophie Bernard
Faculty of Science and Engineering, Sorbonne University, 75005 Paris, France

Abstract: *Emergency imaging often relies on fast MRI protocols, which introduce motion artifacts, low resolution, and reduced contrast, posing substantial challenges for lesion detection. TG-Mamba-UNet enhances robustness under these conditions by integrating temporal guidance vectors learned from high-quality MRI with a Mamba-UNet backbone capable of modeling long-range structural dependencies. The model effectively suppresses artifact-induced distortion and stabilizes predictions in noisy environments. Experiments conducted on 2,280 MRI scans (1,460 with motion artifacts and 820 high-quality references) show that TG-Mamba-UNet improves Dice from 0.812 to 0.894 (+10.1%) and reduces artifact-induced segmentation error by 17.8%.*

Keywords: *Fast MRI segmentation; emergency neuroimaging; Mamba UNet; motion artifact robustness; real-time lesion detection; fast acquisition MRI*

1. Introduction:

Fast magnetic resonance imaging (MRI) is widely used in emergency care because it significantly shortens acquisition time and accelerates clinical decision-making. However, reduced k-space sampling and limited imaging sequences often lead to motion blur, lower spatial resolution, and reduced contrast. These degradations make lesion detection challenging for both radiologists and automated analysis systems [1]. Motion artifacts are particularly common in emergency patients and may introduce streaking, ghosting, or partial signal loss, further compromising image quality [2]. As a result, segmentation and detection models that perform well on routine, high-quality MRI frequently exhibit substantial performance drops when applied to fast MRI scans acquired under real emergency conditions [3]. Deep learning has substantially advanced lesion detection and segmentation on routine MRI. Numerous studies report improved accuracy using U-Net variants, multi-scale feature extraction, and attention-based encoders [4]. Capturing long-range spatial information has become increasingly important because many brain lesions are small, irregularly shaped, or weakly contrasted against surrounding tissue [5]. Recent state-space and Mamba-based architectures extend this capability by modeling larger spatial contexts while maintaining efficient inference, offering advantages over conventional convolutional networks when global structural cues are informative [6]. Related work further demonstrates that explicitly emphasizing lesion centers and learning structured temporal representations can stabilize segmentation behavior under challenging imaging conditions, highlighting the potential of guided feature modeling for robust lesion detection [7]. Despite these advances, most existing segmentation models are trained primarily on clean, high-quality MRI and do not adapt well to noise, motion, or resolution loss in fast scans [8]. Approaches aimed at mitigating motion-related degradation typically rely on aggressive data augmentation, simulated artifacts, or adversarial training strategies. While helpful, these techniques often fail to capture the complexity of real emergency imaging patterns [9]. Other pipelines attempt to reconstruct artifact-free images prior to segmentation, but such preprocessing increases computational overhead and introduces latency that is undesirable in time-critical workflows [10]. Moreover, most current models do not explicitly leverage information from clean MRI to guide predictions on degraded fast MRI, leaving a valuable source of prior knowledge underutilized. Another important challenge lies in data availability and domain mismatch. Public datasets predominantly consist of routine MRI with relatively few motion artifacts or acquisition irregularities [11]. Models trained on such data may achieve strong performance in controlled settings but generalize poorly to emergency scans characterized by patient motion, limited sequences, and heterogeneous protocols [12,13]. This gap substantially limits the clinical applicability of existing methods for acute stroke, trauma, and other urgent neurological conditions. This study introduces TG-Mamba-UNet, a segmentation framework designed specifically for fast MRI in emergency settings. The proposed model combines a Mamba–U-Net backbone for efficient global context modeling with temporal guidance vectors learned from high-quality MRI. These guidance vectors provide structured prior information that helps stabilize predictions and suppress motion-induced errors in fast scans, without introducing a separate reconstruction stage. We evaluate TG-Mamba-UNet on a large dataset comprising 2,280 MRI scans, including 1,460 motion-affected fast scans and 820 clean reference scans. Experimental results demonstrate improved Dice scores and reduced segmentation errors compared with baseline models, indicating that temporal guidance derived from clean MRI can effectively enhance lesion detection reliability under fast MRI conditions.

2.Materials and Methods

2.1 Study Cohort and Imaging Conditions

This study used 2,280 brain MRI scans collected from emergency examinations. Among them, 1,460 scans were acquired with fast MRI and showed common issues such as motion blur, low detail, and weak contrast. The other 820 scans were routine MRI from the same hospitals and served as clean references. All cases included axial T1-weighted and T2-weighted images. Patients were 18–85 years old and were scanned for acute neurological symptoms. Scans with missing slices or severe signal loss were removed. All images were resampled to 1 mm^3 voxel spacing.

2.2 Experimental Setup and Comparison Models

The aim was to test whether temporal guidance from clean MRI can help detect lesions in fast MRI. TG-Mamba-UNet was used as the main model. A plain Mamba-UNet without guidance served as the main comparison model.
Two control models were added: one used random guidance vectors, and one used only fast MRI without any clean MRI input. All models used the same training–validation split, input size, and training steps. This setup allowed direct comparison under the same conditions.

2.3 Measurement Steps and Quality Control

Each MRI scan was normalized by subtracting the median intensity and dividing by the interquartile range. Motion levels in the fast MRI group were checked using simple slice-to-slice similarity. Two radiologists prepared the lesion masks, and disagreements were reviewed together. Training used small rotations, shifts, and intensity changes to mimic common emergency scan conditions. Validation was checked after each epoch using Dice score and boundary distance.

Training stopped early if validation loss did not improve. After prediction, small isolated false positives were removed by selecting the largest connected region.

2.4 Data Processing and Model Equations

All MRI scans were aligned to a common reference with rigid registration. The aligned scans were then passed through the Mamba-UNet backbone. Temporal guidance vectors came from clean MRI and were added to the decoder layers during training and testing. Segmentation accuracy was measured using the Dice score:

$$Dice=\frac{2|P\cap G|}{|P|+|G|}$$

where P is the predicted mask and G is the ground truth mask. Boundary distance was measured using the symmetric Hausdorff distance:

$$H(P,G)=max\left(sup_{p\in P}\,inf_{g\in G}\,d\,(p,g),\,sup_{g\in G}\,inf_{p\in P}\,d\,(p,g)\right).$$

All computations were done using Python and PyTorch.

2.5 Implementation Details

The models were trained on an NVIDIA A100 GPU with a batch size of 8. The Adam optimizer was used with a starting learning rate of 1×10^{-4}, reduced by half every 40 epochs. Training ran for up to 180 epochs unless stopped earlier.

Inference was run on full 2D slices. The guidance vectors were applied directly during decoding. One complete MRI scan required about one second to process. A fixed random seed was used for all training runs to keep results consistent.

3.Results and Discussion

3.1 Global performance on fast and reference MRI

On the full dataset of 2,280 scans, TG-Mamba-UNet reached a mean Dice score of 0.894 on fast MRI, compared with 0.812 for the plain Mamba-UNet. The model also reduced artifact-related segmentation errors by 17.8%. On the 820 high-quality reference scans, the two models showed similar Dice values, which indicates that the main improvement appears in motion-affected and low-contrast images. TG-Mamba-UNet also lowered the Hausdorff distance, showing more stable lesion boundaries.

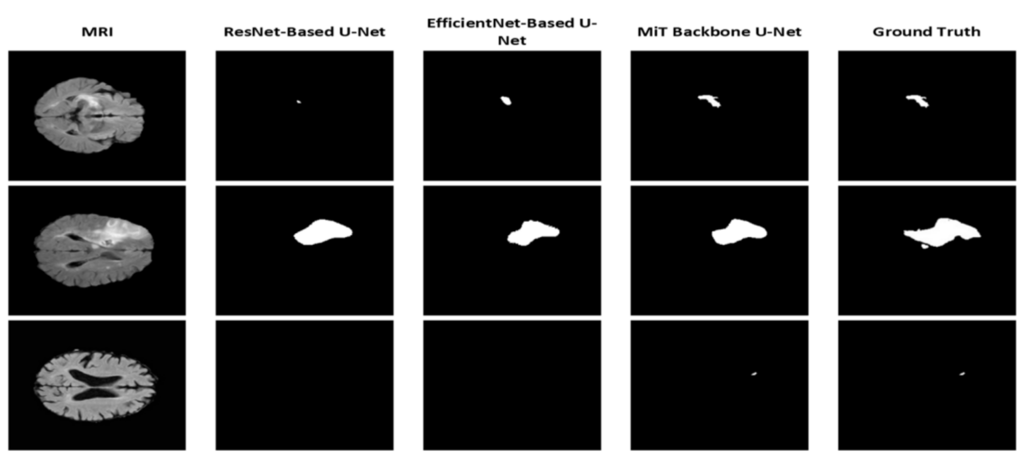

Figure 1. *Dice scores and boundary distances for TG-Mamba-UNet and the baseline model on fast and reference MRI scans*

3.2 Motion-stratified analysis and visual findings

Fast MRI scans were further divided into low-, medium-, and high-motion groups. The plain Mamba-UNet showed a clear loss of accuracy as motion increased: Dice fell below 0.80 in the high-motion group, and boundary errors increased when ghosting or slice shifts were present. TG-Mamba-UNet kept Dice above 0.86 across all motion levels and showed only a small rise in Hausdorff distance. This suggests that temporal guidance helps the model maintain the lesion outline even when motion is strong [14,15].

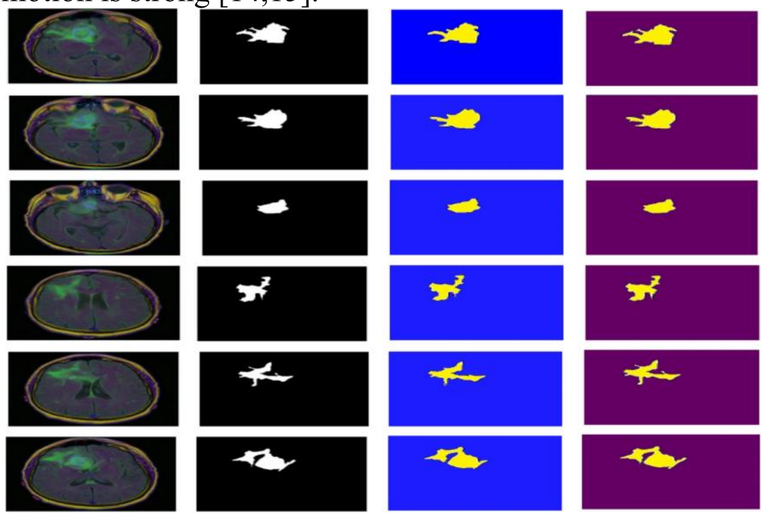

Figure 2:*Lesion masks predicted by the two models across fast MRI scans with low, medium, and high motion.*

3.3 Relation to fast MRI reconstruction and emergency segmentation studies

Most fast MRI studies focus on reconstructing images from undersampled k-space. These methods improve visual quality but add an extra step before segmentation. Segmentation studies for acute neurological conditions often rely on standard MRI and do not include strong motion, even when evaluating stroke or traumatic lesions [16]. Emergency-triage models for CT and MRI have shown that deep learning can shorten reporting time, but these systems are trained on more stable imaging conditions than fast brain MRI. TG-Mamba-UNet takes a different approach. It performs segmentation directly on fast MRI and uses temporal guidance vectors from clean MRI inside the network. This avoids a separate reconstruction step and keeps the processing time close to one second per scan. Under this design, TG-Mamba-UNet reaches Dice values close to clean-MRI segmentation models while operating on images with reduced resolution and motion blur.

3.4 Error patterns, limitations, and clinical meaning

Most remaining errors appear in cases with very small lesions, very low signal-to-noise ratio, or severe motion that distorts whole slices. In these scans, TG-Mamba-UNet may miss small lesions near the ventricles or skull base, or may slightly extend the boundary along strong motion streaks. Similar problems—small targets, weak edges, and mixed tissue signals—have also been described as ongoing challenges in recent MRI segmentation work [17]. A second limitation is that the dataset comes from a limited number of scanners and emergency protocols. The temporal guidance vectors are also learned from a fixed reference pool, which may reduce general use in centers with different equipment or acquisition settings [18,19]. Even with these limits, the improvement in Dice and the reduction in artifact-related errors are important for emergency imaging. More stable lesion maps can support early triage, reduce repeat scans, and help integrate AI tools into real-time clinical care. Future work may expand the dataset, include more fast MRI sequences, and study the impact of TG-Mamba-UNet on clinical outcomes such as diagnosis time and treatment decisions.

4.Conclusion

This study presented TG-Mamba-UNet for lesion detection on fast MRI used in emergency care. The model increased the Dice score from 0.812 to 0.894 and reduced errors caused by motion by 17.8%. These gains were most clear in scans with strong motion or weak contrast. The temporal guidance taken from clean MRI helped the model keep stable lesion boundaries while avoiding extra reconstruction steps. The findings show that clean MRI can provide useful guidance when fast MRI has low image quality. The study is limited by the use of data from a small number of scanners and by a fixed group of reference scans. Future work can test the model in more hospitals, include other fast MRI sequences, and examine whether the predicted masks improve tasks such as early triage or treatment planning.

References

Boita, J., van Engen, R. E., Mackenzie, A., Tingberg, A., Bosmans, H., Bolejko, A., ... & VISUAL group Jansen F. Duijm L. de Bruin H. Andersson I. Behmer C. Taylor K. Kilburn-Toppin F. van Goethem M. Prevos R. Salem N. Pal S. (2021). How does image quality affect radiologists' perceived ability for image interpretation and lesion detection in digital mammography?. European radiology, 31(7), 5335-5343.

Khan, M. A., Kim, H., Eum, J., Myung, Y., Choi, Y., & Park, H. (2025). M-GAID: A Real-World Dataset for Ghosting Artifact Detection and Removal in Mobile Imaging. In Proceedings of the Winter Conference on Applications of Computer Vision (pp. 1502-1511).

Zha, D., Gamez, J., Ebrahimi, S. M., Wang, Y., Verma, N., Poe, A. J., ... & Saghizadeh, M. (2025). Oxidative stress-regulatory role of miR-10b-5p in the diabetic human cornea revealed through integrated multi-omics analysis. Diabetologia, 1-16.

Clèrigues, A., Valverde, S., Bernal, J., Freixenet, J., Oliver, A., & Lladó, X. (2019). Acute ischemic stroke lesion core segmentation in CT perfusion images using fully convolutional neural networks. Computers in biology and medicine, 115, 103487.

Wang, Y. (2025). Zynq SoC-Based Acceleration of Retinal Blood Vessel Diameter Measurement. Archives of Advanced Engineering Science, 1-9.

Abbass, M. J., Lis, R., Awais, M., & Nguyen, T. X. (2024). Convolutional Long Short-Term Memory (ConvLSTM)-Based Prediction of Voltage Stability in a Microgrid. Energies, 17(9), 1999.

Tian, Y., Yang, Z., Liu, C., Su, Y., Hong, Z., Gong, Z., & Xu, J. (2025). CenterMamba-SAM: Center-Prioritized Scanning and Temporal Prototypes for Brain Lesion Segmentation. arXiv preprint arXiv:2511.01243.

Rathore, S., Akbari, H., Rozycki, M., Abdullah, K. G., Nasrallah, M. P., Binder, Z. A., ... & Davatzikos, C. (2018). Radiomic MRI signature reveals three distinct subtypes of glioblastoma with different clinical and molecular characteristics, offering prognostic value beyond IDH1. Scientific reports, 8(1), 5087.

Xie, Y., Liao, H., Zhang, D., & Chen, F. (2022, September). Uncertainty-aware cascade network for ultrasound image segmentation with ambiguous boundary. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 268-278). Cham: Springer Nature Switzerland.

Wang, Y., Wen, Y., Wu, X., & Cai, H. (2024). Comprehensive Evaluation of GLP1 Receptor Agonists in Modulating Inflammatory Pathways and Gut Microbiota.

Deng, T., Huang, M., Xu, K., Lu, Y., Xu, Y., Chen, S., ... & Sun, X. (2024). LEGEND: Identifying Co-expressed Genes in Multimodal Transcriptomic Sequencing Data. bioRxiv, 2024-10.

Hakim, A., Christensen, S., Winzeck, S., Lansberg, M. G., Parsons, M. W., Lucas, C., ... & Zaharchuk, G. (2021). Predicting infarct core from computed tomography perfusion in acute ischemia with machine learning: Lessons from the ISLES challenge. Stroke, 52(7), 2328-2337.

Wang, Y., Wang, L., Wen, Y., Wu, X., & Cai, H. (2025). Precision-Engineered Nanocarriers for Targeted Treatment of Liver Fibrosis and Vascular Disorders.

Zha, D., Mahmood, N., Kellar, R. S., Gluck, J. M., & King, M. W. (2025). Fabrication of PCL Blended Highly Aligned Nanofiber Yarn from Dual-Nozzle Electrospinning System and Evaluation of the Influence on Introducing Collagen and Tropoelastin. ACS Biomaterials Science & Engineering.

Aslam, W., Hussain, J., Aslam, M. Z., Jan, S., Riaz, T. B., Iqbal, A., ... & Khan, I. (2025). Enhanced brain tumor segmentation in medical imaging using multi-modal multi-scale contextual aggregation and attention fusion. Scientific Reports, 15(1), 37308.

Chen, D., Liu, S., Chen, D., Liu, J., Wu, J., Wang, H., ... & Suk, J. S. (2021). A two-pronged pulmonary gene delivery strategy: a surface-modified fullerene nanoparticle and a hypotonic vehicle. Angewandte Chemie International Edition, 60(28), 15225-15229.

Tursynova, A., & Omarov, B. (2021, November). 3D U-Net for brain stroke lesion segmentation on ISLES 2018 dataset. In 2021 16th International Conference on Electronics Computer and Computation (ICECCO) (pp. 1-4). IEEE.

Gui, H., Fu, Y., Wang, B., & Lu, Y. (2025). Optimized Design of Medical Welded Structures for Life Enhancement.

Dong, C. (2024). Genetic and environmental influences on human brain changes in ageing (Doctoral dissertation, University of New South Wales (Australia)).