

A Non-contrast Head CT Foundation Model for Comprehensive Neuro-Trauma Triage

Youngjin Yoo¹, Bogdan Georgescu¹, Yanbo Zhang¹, Sasa Grbic¹, Han Liu¹,
Gabriela D. Aldea^{1,8,9}, Thomas J. Re¹, Jyotipriya Das¹, Poikavila
Ullaskrishnan¹, Eva Eibenberger², Andrei Chekkoury², Uttam K.
Bodanapally⁴, Savvas Nicolaou⁵, Pina C. Sanelli⁶, Thomas J. Schroepel⁷,
Yvonne W. Lui³, and Eli Gibson¹

¹ Digital Technology and Innovation, Siemens Healthineers, Princeton, NJ USA

² Department of Computed Tomography, Siemens Healthineers, Forchheim, Germany

³ Department of Radiology, New York University, New York, NY USA

⁴ Department of Radiology, University of Maryland Medical Center, Baltimore, MD USA

⁵ Department of Radiology, Vancouver General Hospital, Vancouver, BC Canada

⁶ Department of Radiology, Northwell Health, New York, NY USA

⁷ Department of Surgery, UCHealth Memorial Hospital, Colorado Springs, CO USA

⁸ Foundational Technologies, Siemens SRL, Braşov, Romania

⁹ Automation and Information Technology Department, Transilvania University of Braşov, Braşov, Romania

Abstract. Recent advancements in AI and medical imaging offer transformative potential in emergency head CT interpretation for reducing assessment times and improving accuracy in the face of an increasing request of such scans and a global shortage in radiologists. This study introduces a 3D foundation model for detecting diverse neuro-trauma findings with high accuracy and efficiency. Using large language models (LLMs) for automatic labeling, we generated comprehensive multi-label annotations for critical conditions. Our approach involved pretraining neural networks for hemorrhage subtype segmentation and brain anatomy parcellation, which were integrated into a pretrained comprehensive neuro-trauma detection network through multimodal fine-tuning. Performance evaluation against expert annotations and comparison with CT-CLIP demonstrated strong triage accuracy across major neuro-trauma findings, such as hemorrhage and midline shift, as well as less frequent critical conditions such as cerebral edema and arterial hyperdensity. The integration of neuro-specific features significantly enhanced diagnostic capabilities, achieving an average AUC of 0.861 for 16 neuro-trauma conditions. This work advances foundation models in medical imaging, serving as a benchmark for future AI-assisted neuro-trauma diagnostics in emergency radiology.

Keywords: Foundation Model · Neuro-Trauma Detection · Head CT.

1 Introduction

Head computed tomography (CT) is an essential diagnostic tool of emergency medicine, particularly for assessing acute neurological symptoms and head trauma [13, 20]. Its utilization is on the rise [26] while the availability of trained radiologists qualified to interpret its results is facing a worldwide shortage [16, 24]. AI-assisted interpretation of emergent head CT could help address this situation by increasing the efficiency and accuracy of available qualified radiologists and of less specialized clinicians in interpreting such imaging [19, 22]. AI-driven approaches have evolved from supervised methods [12, 21] to self-supervised and semi-supervised approaches [7, 18], reducing dependence on extensive annotations. Recently foundation models trained on large datasets have shown remarkable success across domains [2] such as MedViT [14] and MIMIC-CXR [10]. Notably, CT-CLIP [6] has enabled supervised-level zero-shot detection of chest abnormalities and FM-CT [27], a newly introduced head CT foundation model for detecting various neuro conditions such as hemorrhages, tumors and other abnormalities, illustrating the transformative potential in medical imaging.

Despite this progress, applying foundation models to head CT remains challenging due to anatomical complexity and the broad spectrum of neuro-trauma conditions. Rapid and accurate diagnosis is critical in emergency settings [17], yet traditional interpretation is time-consuming and prone to variability [4, 15]. Existing foundation models may underperform in neuro-trauma detection due to domain-specific limitations. To address this, we developed a 3D foundation model specialized for head CT, trained on a large multi-site dataset. This model enables accurate, efficient few-shot detection of neuro-trauma conditions, potentially enhancing trauma triage and improving patient outcomes.

Our contributions include setting benchmark performance for neuro-trauma detection, demonstrating robust generalization across common and rare critical findings, and emphasizing the importance of domain-specific pretraining [25, 27]. By integrating neuro-specific pathological and anatomical features, we highlight the advantages of specialized foundation models over CT-CLIP and the broader coverage of neuro-trauma findings compared to FM-CT.

2 The comprehensive neuro-trauma detection foundation model

To develop a head CT foundation model, a neuro-radiologist curated a comprehensive list of neuro-trauma findings requiring urgent clinical attention [23, 28], including Hemorrhage, Infarct, Mass Lesion, Mass Effect, Hydrocephalus, Midline Shift, Skull Fracture, Cerebral hemorrhagic Contusion, Diffuse Cerebral Edema, Microhemorrhage, Diffuse Axonal Injury, Generalized Cerebral Edema, Pneumocephalus, Brain Herniation, Arterial Hyper-density and Venous Sinus Hyper-density. Using this list, we automatically generated labels for a large-scale dataset through an Large Language Models (LLMs) pipeline. We independently

pretrained two task-specific vision networks and integrated these pretrained networks into a foundation model via multimodal finetuning with LLM-generated labels. Instead of training directly on image-report pairs [6], we utilized image-LLM multi-label pairs to streamline task-specific pretraining and directly integrate anatomical and pathological features.

2.1 Automatic comprehensive labeling

Leveraging recent advancements of LLMs in generating medical content [1], we automatically generated multi-labels for neuro-trauma findings for each radiological report in our dataset. The labels were generated using a private GPT4-o model on our private network. To efficiently label a large-scale dataset, we investigated prompts that requested 16 multiple labels, minimizing the need for repetitive label processing. The optimized prompt was “Given this radiology report, extract POS or NEG value for these concepts {0}. POS means the concept is present in patient as per report. NEG means not. Return as json format with keys being the name of the concepts and value being either POS or NEG. Report:{*report content*}”, and for {0} we provided the entire 16 neuro-trauma findings as a list. Presenting labeling examples to the LLM before labeling the entire dataset improved labeling accuracy.

2.2 Task-specific pretraining

We pretrained two networks independently, each performing a specific task: brain bleeding (hemorrhage) subtype segmentation and brain anatomy parcellation.

The hemorrhage subtype segmentation network is based on a 3D Dense U-Net architecture [8], specifically designed for classifying five hemorrhage subtypes: intraparenchymal, subarachnoid, intraventricular, subdural, and epidural hemorrhages [5]. We adapted the specialized network architecture described in [5] for this task. In this work, the network additionally incorporates Squeeze-and-Excitation (SE) blocks [9] to enhance feature recalibration throughout the architecture. SE blocks are strategically placed before and after each DenseBlock to adaptively recalibrate feature responses.

The brain parcellation network is designed for segmenting 15 brain structures: left/right hemispheres, supratentorial/infratentorial regions, frontal lobe, parietal lobe, occipital lobe, temporal lobe, cerebellum, basal ganglia, medulla oblongata, pons, midbrain, falx and ventricles. The network employs a U-Net architecture with 15 stages, using ReLU activations, batch normalization, strided convolutions for downsampling and transposed convolutions for upsampling. The model employs a multi-head architecture with three output layers: one dedicated to segmenting left-side hemispheres, another for supratentorial-infratentorial regions, and a third handling the remaining brain structures.

2.3 Building a foundation model with multimodal finetuning

Our foundation model is based on the 3D densely connected network specifically designed for brain hemorrhage classification [5], which we trained to perform

comprehensive neuro-trauma detection tasks. We refer to it as the Comprehensive Neuro Trauma Detection Network (CNTD-Net). To learn a wider variety of imaging features for handling heterogeneous pathologies beyond brain hemorrhage, we expanded the network capacity of the brain hemorrhage classification network by increasing the number of layers and feature channels, resulting in the DeepCNTD-Net. Specifically, we increased the convolutional channel growth rate of the 3D DenseBlock from 5 to 8, the initial number of feature maps from 16 to 64, the total number of 3D dense layers from 15 to 20, and the final feature vector dimension from 1638 to 4032. We pretrained both CNTD-Net and DeepCNTD-Net by performing the comprehensive neuro-trauma detection task using the LLM-generated multi-labels. To integrate the task-specific pretrained networks, we employed an encoder that extracts features with dimensions similar to those of DeepCNTD-Net by collapsing the segmentation features. The features from DeepCNTD-Net, the hemorrhage subtype segmentation network, and the brain parcellation network were then fused using linear layers, which were subsequently used for multi-label classification. The feature fusion process was performed after freezing the pretrained networks. Pre-training and fine-tuning were performed using the Adam optimizer [11] to adapt the learning rates for each parameter. Additionally, binary cross-entropy with logits loss was employed, incorporating adjusted class weights to effectively address the class imbalances in the dataset. The overall procedure is illustrated in Fig. 1.




Fig. 1. Overview of the head CT foundation model for neuro-trauma triage.

3 Results and Discussion

3.1 Dataset and Preprocessing

Anonymized non-contrast CT (NCCT) head volumes were retrospectively collected from nine centers across the U.S., Canada, China, and India, with ethics committee approvals waiving informed consent. Data were sourced from pre-established cohorts or retrospective selections. NCCT volumes were acquired

using Siemens, GE, and Toshiba scanners. Exclusion criteria included age under 18 or absence of axial reconstruction. A total of 29,395 NCCT studies met inclusion criteria: (a) 26,514 studies were used for model development—23,592 for training and 2,922 for system optimization (architecture selection, parameter tuning, and classifier calibration); (b) 2,881 studies were reserved for independent performance evaluation. Patient-level random splitting was performed before development to prevent data contamination. The prevalence of trauma findings, based on LLM labels, is shown in Fig. 2. Preprocessing involved automatic alignment of axial NCCT volumes to a standard reference frame, resampling to a 1-mm in-plane and 4-mm out-of-plane resolution, and normalization using Hounsfield Unit (HU) windows: 0–80 HU (bleeding), -20–180 HU (brain), and -800–2000 HU (bone), scaled to 0–1. To enhance robustness against translation and scanner noise, data augmentation included random in-plane translation (± 10 mm), in-plane flipping (50% probability), random CT windowing noise (± 10 HU), and random image noise (0.01 STD).




Fig. 2. A comprehensive list of neuro-trauma findings that encompass the critical conditions requiring immediate clinical attention in trauma emergency centers curated by a neuro-radiologist and their prevalence in our dataset.

3.2 LLM labeling accuracy against manual expert label

We evaluated the accuracy of our LLM pipeline labeling against expert-generated manual labels for six major neuro-trauma findings in head CT scans (Table 1). These six major labels were previously generated by expert users. LLMs achieved 92–99% accuracy, except for ischemia / infarction at 79%, likely due to the lower

sensitivity of NCCT and variability in radiology report terminology. This highlights the LLMs’ robustness in detecting diverse pathologies. Table 1 compares LLM labels with neuro-radiologist annotations for 200 randomly selected cases for all 16 findings. The LLMs excelled in identifying microhemorrhages, diffuse axonal injuries, and venous sinus hyperdensity (accuracy: 1.0) and performed well on diffuse/generalized cerebral edema and skull fractures (0.99). These results demonstrate the LLMs’ effectiveness in accurately labeling complex head CT findings, supporting automated trauma triage. Residual ambiguities in subtle findings may be further addressed through clinically optimized prompting, LLM ensemble consensus, or uncertainty quantification.

Table 1. LLM labeling accuracy compared to expert labels. Accuracy is reported for six major findings across the entire dataset (values in parentheses) and for all 16 findings in 200 randomly selected cases.

Finding	Accuracy	Finding	Accuracy
Hemorrhage	0.95 (0.920)	Diffuse cerebral edema	0.99
Ischemia/Infarct	0.80 (0.786)	Microhemorrhage	1.0
Mass Lesion	0.95 (0.921)	Diffuse axonal injury	1.0
Mass Effect	0.94 (0.958)	Generalized cerebral edema	1.0
Hydrocephalus	0.95 (0.950)	Pneumocephalus	0.96
Midline Shift	0.99 (0.987)	Brain herniation	0.98
Skull fracture	0.99	Arterial Hyper-density	0.99
Cerebral hemorrhagic contusion	0.98	Venous Sinus Hyper-density	1.0

3.3 Comprehensive neuro-trauma detection performance

Table 2 presents ablation study results, assessing the impact of different model components. The baseline CNTD-Net, designed for six major findings, achieved an AUC of 0.768. Expanding it to DeepCNTD-Net significantly improved performance (AUC: 0.858). Adding brain hemorrhage segmentation features (hemSegFeat) further increased AUC to 0.873, underscoring their importance. Incorporating brain anatomy features (brainAnatFeat) provided a slight boost to 0.875 (Sensitivity: 0.815 ± 0.105 , Specificity: 0.809 ± 0.070 , Precision: 0.449 ± 0.200 , F1: 0.558 ± 0.186). For detecting all 16 findings, DeepCNTD-Net reached an AUC of 0.849, improving to 0.859 with hemSegFeat and peaking at 0.861 with brainAnatFeat (Sensitivity: 0.827 ± 0.101 , Specificity: 0.800 ± 0.117 , Precision: 0.225 ± 0.222 , F1: 0.305 ± 0.249). Table 3 compares DeepCNTD-Net with CT-CLIP, which, despite being trained on the same LLM-generated labels and fine-tuned with CT-LiPro [6], achieved lower AUCs (0.822 for six major findings, 0.835 for all 16). In contrast, DeepCNTD-Net, leveraging hemSegFeat and brainAnatFeat, achieved superior scores of 0.875 and 0.861, highlighting its advantage in neuro-trauma detection through specialized feature integration.

Table 2. Average detection performance (AUC) results of ablation study for the six major neuro-trauma findings and all 16 findings. **hemSegFeat**: brain hemorrhage segmentation features; **brainAnatFeat**: brain anatomy segmentation features.

Model	Average AUC
CNTD-Net (6 major findings)	0.768 ± 0.064
DeepCNTD-Net (6 major findings)	0.858 ± 0.066
DeepCNTD-Net + hemSegFeat (6 major findings)	0.873 ± 0.068
DeepCNTD-Net + hemSegFeat + brainAnatFeat (6 major findings)	0.875 ± 0.065
DeepCNTD-Net (all 16 findings)	0.849 ± 0.090
DeepCNTD-Net + hemSegFeat (all 16 findings)	0.859 ± 0.085
DeepCNTD-Net + hemSegFeat + brainAnatFeat (all 16 findings)	0.861 ± 0.081

Table 3. Average detection performance of CT-CLIP and DeepCNTD-Net with hemSegFeat and brainAnatFeat.

Model	Average AUC
CT-CLIP (6 major findings)	0.822 ± 0.081
DeepCNTD-Net (6 major findings)	0.875 ± 0.065
CT-CLIP (all 16 findings)	0.835 ± 0.083
DeepCNTD-Net (all 16 findings)	0.861 ± 0.081

We analyzed individual AUC performance to identify key factors driving detection improvements (Fig. 3). Models incorporating segmentation and anatomical features consistently outperformed baselines across both major and rare findings, reinforcing the value of multimodal integration. DeepCNTD-Net, particularly with hemSegFeat and brainAnatFeat, delivered superior results, achieving an AUC of 0.92 for hemorrhage detection versus CT-CLIP’s 0.83 and FM-CT’s range of 0.835-0.929 (reported in [27]). For midline shift, DeepCNTD-Net variants reached up to 0.95, surpassing CT-CLIP’s 0.92. The enhanced models also improved mass effect detection (AUC: 0.89), while CT-CLIP excelled in generalized cerebral edema, suggesting complementary strengths that could be leveraged through model fusion. To assess generalizability, we evaluated DeepCNTD-Net on the CQ500 dataset [3], where it maintained strong performance in hemorrhage (AUC: 0.920) and midline shift (AUC: 0.965), outperforming FM-CT [27] (hemorrhage AUC: 0.776-0.850, midline shift AUC: 0.780), but exhibited lower accuracy for mass effect (AUC: 0.840) against FM-CT (AUC: 0.90). Average sensitivity was 0.850 ± 0.023 , specificity 0.857 ± 0.157 , precision 0.710 ± 0.226 , and F1 score 0.758 ± 0.132 . Additionally, FM-CT outperformed in detecting edema (AUC: 0.827-0.923) and hydrocephalus (AUC: 0.910-0.944), whereas DeepCNTD-Net achieved an AUC of 0.80-0.90 for edema and 0.90 for hydrocephalus. These findings highlight the model’s reliability while pointing to areas for further refinement in neuro-trauma detection. Fig. 4 showcases multi-finding detection, illustrating precise trauma identification alongside occasional false positives due to confounding pathologies.




Fig. 3. Individual detection performance (AUC) for six major critical findings and all 16 neuro-trauma conditions. CNTD-Net was evaluated exclusively for the six major findings.




Fig. 4. Examples of comprehensive neuro-trauma triage. "POS" indicates a positive case, while "NEG" denotes a negative case, both based on the LLM-generated multi-labels, and AI triage score is between 0.0 and 1.0.

Table 4. Individual detection performance of DeepCNTD-Net with hemSegFeat and brainAnatFeat on the CQ500 dataset for hemorrhage, mass effect and midline shift.

Finding	AUC (95% CI)
Hemorrhage	0.920 (0.876-0.956)
Mass Effect	0.840 (0.755-0.910)
Midline Shift	0.965 (0.922-0.996)

4 Conclusion

This study highlights the potential of a specialized 3D foundation model for head trauma triage, addressing the need for rapid, accurate diagnostics in emergency medicine. By integrating LLM-driven automated labeling with task-specific neural networks for hemorrhage segmentation and brain anatomy parcellation, we developed a robust framework for detecting a broad spectrum of neuro-trauma findings in CT scans. Our DeepCNTD-Net variant achieved high accuracy across both common and less frequent critical conditions, with multimodal feature integration enhancing performance. It reached an average AUC of 0.861 for 16 neuro-trauma findings, reinforcing the importance of domain-specific pretraining. This work contributes to the growing evidence supporting AI-driven foundation models in clinical practice, helping bridge the gap between increasing emergency head CT scans and the radiologist shortage. It also serves as a research reference point for future foundation AI applications in neuro-trauma triage. To maximize clinical impact, future efforts will focus on real-world validation and seamless clinical integration to improve patient care.

Disclaimer The concepts and information presented in this paper are based on research results that are not commercially available. Future availability cannot be guaranteed.

Disclosure of Interests The authors have no competing interests to declare.

References

1. Bicknell, B.T., Butler, D., Whalen, S., Ricks, J., Dixon, C.J., Clark, A.B., Spaedy, O., Skelton, A., Edupuganti, N., Dzubinski, L., et al.: ChatGPT-4 omni performance in usmle disciplines and clinical skills: Comparative analysis. *JMIR Medical Education* **10**(1), e63430 (2024)
2. Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., et al.: On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021)
3. Chilamkurthy, S., Ghosh, R., Tanamala, S., Biviji, M., Campeau, N.G., Venugopal, V.K., Mahajan, V., Rao, P., Warier, P.: Deep learning algorithms for detection of critical findings in head CT scans: a retrospective study. *The Lancet* **392**(10162), 2388–2396 (2018)
4. Erly, W.K., Berger, W.G., Krupinski, E., Seeger, J.F., Guisto, J.A.: Radiology resident evaluation of head CT scan orders in the emergency department. *American journal of neuroradiology* **23**(1), 103–107 (2002)
5. Gibson, E., Georgescu, B., Ceccaldi, P., Trigan, P.H., Yoo, Y., Das, J., Re, T.J., Rs, V., Balachandran, A., Eibenberger, E., et al.: Artificial intelligence with statistical confidence scores for detection of acute or subacute hemorrhage on noncontrast CT head scans. *Radiology: Artificial Intelligence* **4**(3), e210115 (2022)
6. Hamamci, I.E., Er, S., Almas, F., Simsek, A.G., Esirgun, S.N., Dogan, I., Dasdelen, M.F., Wittmann, B., Simsar, E., Simsar, M., et al.: A foundation model utilizing chest CT volumes and radiology reports for supervised-level zero-shot detection of abnormalities. *CoRR* (2024)

7. 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)
8. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)
9. Jie, H., Li, S., Gang, S., Albanie, S.: Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. vol. 5 (2018)
10. Johnson, A.E., Pollard, T.J., Berkowitz, S.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Mark, R.G., Horng, S.: MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data* **6**(1), 317 (2019)
11. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *CoRR* **abs/1412.6980** (2014), <https://arxiv.org/abs/1412.6980>
12. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* **25** (2012)
13. Lolli, V., Pezzullo, M., Delpierre, I., Sadeghi, N.: MDCT imaging of traumatic brain injury. *The British journal of radiology* **89**(1061), 20150849 (2016)
14. Manzari, O.N., Ahmadabadi, H., Kashiani, H., Shokouhi, S.B., Ayatollahi, A.: MedViT: a robust vision transformer for generalized medical image classification. *Computers in Biology and Medicine* **157**, 106791 (2023)
15. Parag, P., Hardcastle, T.C.: Interpretation of emergency ct scans of the head in trauma: neurosurgeon vs radiologist. *World Journal of Surgery* **46**(6), 1389–1395 (2022)
16. Parag, P., Hardcastle, T.C.: Shortage of radiologists in low to middle income countries in the interpretation of CT scans in trauma. *Bangladesh Journal of Medical Science* **21**(3), 489–491 (2022)
17. Rincon, S., Gupta, R., Ptak, T.: Imaging of head trauma. *Handbook of clinical neurology* **135**, 447–477 (2016)
18. 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)
19. Savage, C.H., Tanwar, M., Elkassem, A.A., Sturdivant, A., Hamki, O., Sotoudeh, H., Sirineni, G., Singhal, A., Milner, D., Jones, J., et al.: Prospective evaluation of artificial intelligence triage of intracranial hemorrhage on noncontrast head CT examinations. *American Journal of Roentgenology* **223**(5), e2431639 (2024)
20. Sharp, A.L., Nagaraj, G., Rippberger, E.J., Shen, E., Swap, C.J., Silver, M.A., McCormick, T., Vinson, D.R., Hoffman, J.R.: Computed tomography use for adults with head injury: describing likely avoidable emergency department imaging based on the canadian CT head rule. *Academic Emergency Medicine* **24**(1), 22–30 (2017)
21. Simonyan, K.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
22. Vimalaesar, K., Robert, D., Kumar, S., Kumar, A., Narbone, M., Dharmadhikari, R., Harrison, M., Ather, S., Novak, A., Grzeda, M., et al.: Assessing the effectiveness of artificial intelligence (AI) in prioritising ct head interpretation: study protocol for a stepped-wedge cluster randomised trial (ACCEPT-AI). *BMJ open* **14**(6), e078227 (2024)

23. Wintermark, M., Sanelli, P.C., Anzai, Y., Tsiouris, A.J., Whitlow, C.T., Druzgal, T.J., Gean, A.D., Lui, Y.W., Norbash, A.M., Raji, C., et al.: Imaging evidence and recommendations for traumatic brain injury: conventional neuroimaging techniques. *Journal of the American College of Radiology* **12**(2), e1–e14 (2015)
24. Yaghmai, V.: Editorial comment: Bridging the gap—international medical graduates and the radiologist shortage in the united states. *American Journal of Roentgenology* **222**(1), e2330540 (2024)
25. Yoo, Y., Zhao, G., Sandu, A.E., Re, T.J., Das, J., Wang, H., Kim, M., Shen, C., Lee, Y., Kondziolka, D., et al.: The importance of data domain on self-supervised learning for brain metastasis detection and segmentation. In: *Medical Imaging 2023: Computer-Aided Diagnosis*. vol. 12465, pp. 556–562. SPIE (2023)
26. Yun, B.J., Borczuk, P., Zachrison, K.S., Goldstein, J.N., Berlyand, Y., Raja, A.S.: Utilization of head CT during injury visits to united states emergency departments: 2012–2015. *The American journal of emergency medicine* **36**(8), 1463–1466 (2018)
27. Zhu, W., Huang, H., Tang, H., Musthyala, R., Yu, B., Chen, L., Vega, E., O'Donnell, T., Dehkharghani, S., Frontera, J.A., et al.: 3D foundation AI model for generalizable disease detection in head computed tomography. *arXiv preprint arXiv:2502.02779* (2025)
28. Zimmermann, L.L., Tran, D., Lovett, M., Mangat, H.: Emergency neurological life support, traumatic brain injury protocol (version 4.0). *Neurocritical Care Society* pp. 1–33 (2019)