

Article

Advancing Brain Metastases Detection in T1-Weighted Contrast-Enhanced 3D MRI Using Noisy Student-Based Training

Engin Dikici ^{1,*}, Xuan V. Nguyen ¹, Matthew Bigelow ¹, John L. Ryu ² and Luciano M. Prevedello ¹¹ Department of Radiology, The Ohio State University College of Medicine, Columbus, OH 43210, USA² ProScan Imaging, Columbus, OH 43230, USA

* Correspondence: engin.dikici@osumc.edu

Abstract: The detection of brain metastases (BM) in their early stages could have a positive impact on the outcome of cancer patients. The authors previously developed a framework for detecting small BM (with diameters of <15 mm) in T1-weighted contrast-enhanced 3D magnetic resonance images (T1c). This study aimed to advance the framework with a noisy-student-based self-training strategy to use a large corpus of unlabeled T1c data. Accordingly, a sensitivity-based noisy-student learning approach was formulated to provide high BM detection sensitivity with a reduced count of false positives. This paper (1) proposes student/teacher convolutional neural network architectures, (2) presents data and model noising mechanisms, and (3) introduces a novel pseudo-labeling strategy factoring in the sensitivity constraint. The evaluation was performed using 217 labeled and 1247 unlabeled exams via two-fold cross-validation. The framework utilizing only the labeled exams produced 9.23 false positives for 90% BM detection sensitivity, whereas the one using the introduced learning strategy led to ~9% reduction in false detections (i.e., 8.44). Significant reductions in false positives (>10%) were also observed in reduced labeled data scenarios (using 50% and 75% of labeled data). The results suggest that the introduced strategy could be utilized in existing medical detection applications with access to unlabeled datasets to elevate their performances.

Keywords: brain metastases; noisy student; semi-supervised training

Citation: Dikici, E.; Nguyen, X.V.; Bigelow, M.; Ryu, J.L.; Prevedello, L.M. Advancing Brain Metastases Detection in T1-Weighted Contrast-Enhanced 3D MRI Using Noisy Student-Based Training. *Diagnostics* **2022**, *12*, 2023. <https://doi.org/10.3390/diagnostics12082023>

Academic Editor: Andreas Kjaer

Received: 23 July 2022

Accepted: 19 August 2022

Published: 21 August 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Brain metastases (BM) are cancerous lesions indicating an advanced and disseminated state of disease. The early detection of BM may enable a treatment utilizing targeted radiotherapy that (1) allows for a less-invasive and less-costly procedure when compared to surgery and (2) leads to fewer adverse neurologic symptoms when compared to whole-brain radiation [1]. Contrast-enhanced 3D magnetic resonance imaging is the key modality for the detection, characterization, and monitoring of BM. However, the task can become challenging when BM lesions are very small; their low contrast and structural similarities with surrounding structures (in some slice angles) may obstruct/limit their visual detection [2].

The automated detection/segmentation of BM in MRI data via machine learning (ML) was investigated in several studies [3–8]; Cho et al. [9] provided a literature review study on the topic comparing state-of-the-art (SOTA) approaches based on the Checklist for Artificial Intelligence in Medical Imaging (CLAIM) [10] and Quality Assessment of Diagnostic Accuracy Studies (QUADAS-2) criteria [11]. The authors previously introduced a framework for T1-weighted contrast-enhanced 3D MRI [4] (analyzed among other SOTA approaches in [9]) for the detection of BM with diameters less than 15 mm to assist early detection of disease.

The noisy student (NS)-based self-training approach was first introduced in [12] to enable the usage of large amounts of unlabeled (i.e., not segmented or annotated for ML

training) datasets to improve the accuracy of existing ML-based solutions. While the approach is still new as of the writing of this manuscript, it has already been utilized in various domains, including medical imaging. In [13], the NS method was employed for abdominal organ segmentation in 3D computed tomography (CT) datasets. The study utilized 3D nnU-Net [14] as both the teacher and student models, which were trained with 41 labeled and 363 unlabeled datasets. The study reported ~3% improvement in the Dice similarity coefficient (DSC) due to unlabeled data. Rajan et al. [15] used the approach for multi-label classification of chest X-ray images. In addition to classical image augmentations, the method also utilized mixup [16] and confidence-tempering regularizations for the noisy student model's training. Their results showed that a ResNet-18 [17] trained via their scheme with 12.5k labeled and 15k unlabeled samples could outperform a similar model trained using a 138k labeled set. In [18], Shak et al. proposed an NS-based lung cancer prediction framework for CT images, where a DeepSEED network [19] was trained using labeled Lung Image Database Consortium image collection (LIDC) [20] and unlabeled National Lung Screening Trial (NLST) [21] datasets. The applicability of the NS for the segmentation of intracranial hemorrhages in CT was presented in [22]. The study included 456 labeled and 25k unlabeled head CT exams for the semi-supervised training of a PatchFCN [23], producing a ROC area under the curve (AUC) of 0.939 for the CQ500 dataset [24]. Kim et al. [25] utilized the NS for the quantification of severity and localization of COVID-19 lesions on chest radiographs. Their framework consisted of a vision transformer-based [26] backbone, trained using ~1000 labeled and ~190k unlabeled images collected from public datasets (i.e., [27,28]), and provided results comparable to those of expert radiologists. Lastly, Ref. [29] proposed a cascaded learning approach for whole heart segmentation in CT angiography data. The method employed a V-net-based backbone [30] and advanced the NS with a shape-constrained training, producing a Dice coefficient of 0.917 with only 16 labeled and 64 unlabeled 3D images.

The goal of this study is to advance the BM detection framework [4] via an NS-based self-training strategy to utilize a large corpus of unlabeled T1c data in a semi-supervised fashion. The BM detection sensitivity in connection with the number of false detections has been the main criteria for various BM detection solutions [3,4,6]. Thus, the study introduces a sensitivity-based noisy-student training approach to factor in the BM detection sensitivity. To this end, the advanced framework uses (1) the CropNet-based architectures as the teacher and student models, (2) model and data noising procedures to enable knowledge expansion, and (3) a novel pseudo-labeling strategy to factor in the BM detection sensitivity during the models' training. The proposed models are trained in an iterative fashion using a smaller group of labeled and a larger group of pseudo-labeled data; the proposed adaptation of the detection framework is illustrated in Figure 1.

This manuscript first provides a brief overview of the BM detection framework. Then, it describes (1) the model architectures, (2) noising mechanisms, and (3) a novel NS-based semi-supervised training setup that introduces a pseudo-labeling strategy. The results section summarizes the experiments performed using two-fold cross-validation (CV) on 217 labeled and 1247 unlabeled T1c exams. It presents comparative analyses for (1) the frameworks trained with and without the given NS-based strategy and (2) hyperparameters regarding the NS-based training process. Next, the results are discussed, where the framework is compared with other SOTA approaches via metrics including the BM detection sensitivity and false BM detection count. The report concludes with a summary of the novelties of the introduced study and future work considerations.

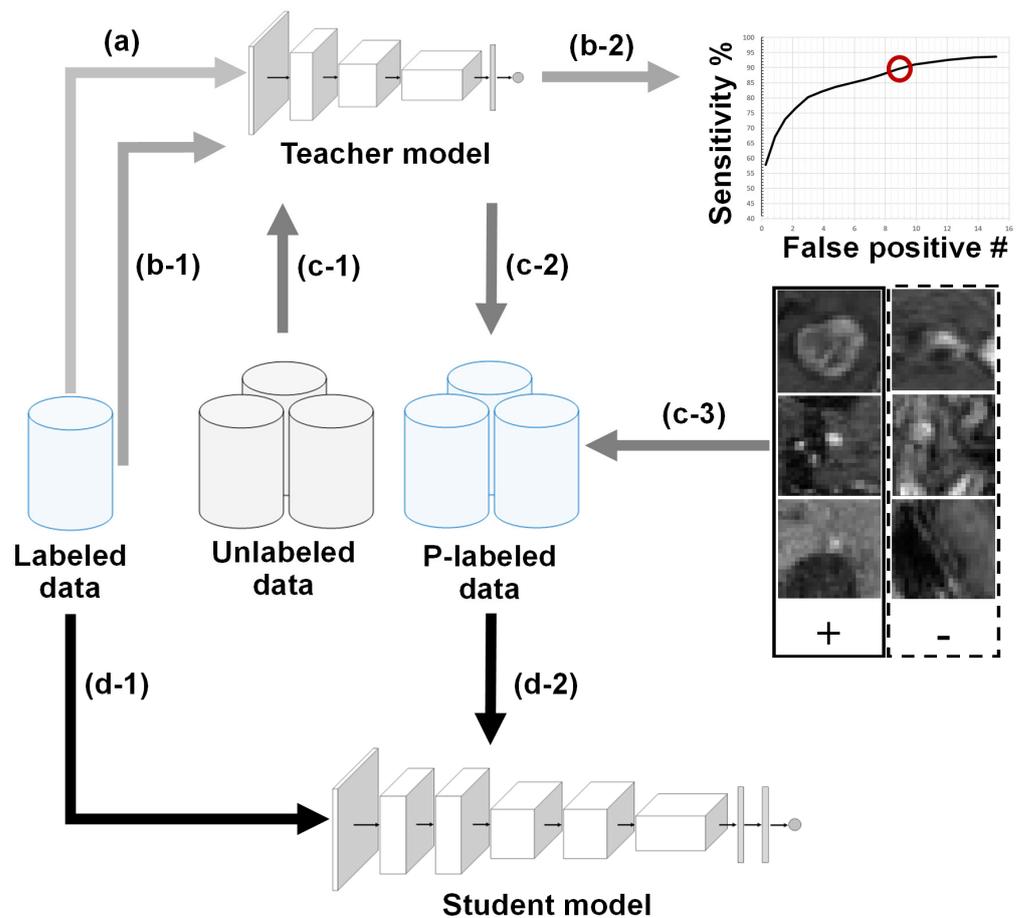


Figure 1. Adaptation of the NS-based self-training for the BM detection framework. The labeled T1c data is used for the training of a lower-capacity teacher model (a). Based on the labeled data (b-1) and a target sensitivity (b-2), the teacher model's threshold for pseudo-labeling is determined. Via this threshold and unlabeled data (c-1), the teacher model generates paired pseudo-labeled data (c-2). The higher-capacity student model is trained using labeled (d-1) and pseudo-labeled (d-2) T1c data. The training of teacher and student models could be performed iteratively, where the final student model is utilized as the deployment model.

2. Materials and Methods

2.1. BM Detection Framework Overview

The BM detection framework consists of two major stages: (1) candidate tumor detection and (2) classification of the candidates. The candidate selection procedure adapts the Laplacian of Gaussian (LoG) [31] approach for maximizing the BM detection sensitivity while minimizing the number of candidates [4]. The detected candidates are then used as centers of the volumetric region of interests (ROIs) for processing by the CropNet, a convolutional neural network, to classify a given ROI as BM or not.

The CropNet model architecture was first introduced in [4]. The model processes isotropically sampled cubic ROIs, with each voxel representing 1 mm^3 . It inherits a contracting network architecture, where (1) each resolution level is processed via a set of blocks consisting of convolution, rectified linear activation unit (i.e., ReLU), and dropout layers, and (2) the resolution downsampling is performed via max-pooling layers. The input volume's dimensions and the network's blocks per resolution level are denoted in the model's name; e.g., CropNet-bX-Ymm describes a network with X blocks per resolution level that processes cubic regions with the edge length of Y mm.

The LoG approach produces thousands of candidates (i.e., ~70k) for each 3D dataset; hence, the ROIs with BM are under-represented. Thus, the framework employs a paired

training strategy, where each training data batch has an equal number of positive and negative samples (i.e., ROIs with 1-BM and 0-non-BM). The binary cross-entropy loss for this classification problem is minimized during CropNet's training.

2.2. Teacher–Student Models and Noising Mechanisms

During the NS training, the student model capacity needs to meet or exceed that of the teacher model to enable knowledge expansion (e.g., EfficientNet-B7 and EfficientNet-L2 were used as the teacher–student pair in [12], presenting a seven-fold network capacity scaling). Accordingly, dedicated BM classification models CropNet-b2-16mm and CropNet-b4-16mm are utilized as the teacher–student pair in this study. The network architectures for these models are shown in Figure 2.

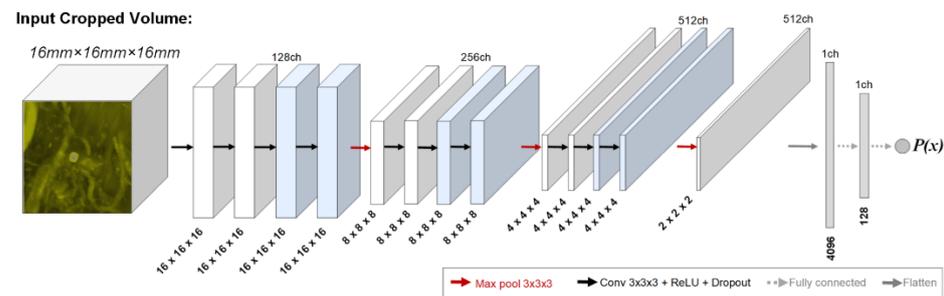


Figure 2. Framework-specific teacher–student models. The network architectures of CropNet-b4-16mm (all blocks) and CropNet-b2-16mm (blocks excluding the blue ones). The input is a 16 mm × 16 mm × 16 mm isotropic region of interest (ROI), where the output is a scalar in the range of [0, 1] giving the probability of input ROI including a BM lesion. The proposed solution utilizes CropNet-b2-16mm (lower-capacity) as the teacher and CropNet-b4-16mm (higher-capacity) as the student models.

The noising mechanisms are necessary for improving the generalizability of the neural networks, especially in limited-data scenarios [32]. They are critical for NS-based self-training strategies to enforce invariances in the decision function during the training of the student model. Both the model and data noising mechanisms are deployed: the data noising is applied via random Simard-type 3D deformations, gamma corrections, rotations, and image flips, whereas the model noising is provided via dropout layers (see Figure 3).

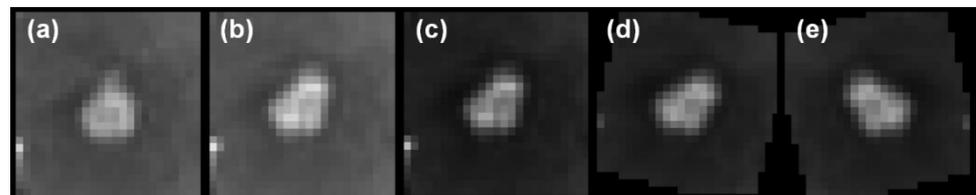


Figure 3. The data noising process. (a) Mid-axial slice of an original cropped sample. (b) Random elastic deformation, (c) random gamma correction, (d) random image rotation, and (e) random image flip are applied.

2.3. Technical Contribution: Sensitivity-Based NS Algorithm

The labeled data set is initially processed via the constrained LoG algorithm to generate BM candidates for the given data. The candidates and the manual BM annotations are then used for producing a set of ROI volumes X with their corresponding labels Y (1: BM, 0: non-BM), where X is a paired set with an equal number of positive and negative samples. The teacher model θ^t is trained with the data noised version of the extracted ROIs by minimizing the binary cross-entropy type loss function ℓ as:

$$\operatorname{argmin}_{\theta^t} (\ell(\theta^t(X_{\text{noised}}), Y)). \quad (1)$$

After the training of the teacher model, the unlabeled data is processed using the constrained LoG algorithm to generate BM candidates. As there are no annotations for the unlabeled data, the pseudo-labels need to be generated. The framework produces BM detections based on a model response threshold $\bar{\tau}$, determined using the BM detection sensitivity. Thus, the relationship between the threshold and the teacher model sensitivity (t_s) can be learned from the training data. Accordingly, the sensitivity in relation to the false BM detections (generated by θ^t on X) was computed, and a response threshold μ was set based on this value; for the unlabeled ROIs $\hat{X} = \{\hat{x}_1, \hat{x}_2, \dots, \hat{x}_N\}$, pseudo-labels \hat{Y} are determined by (see Figure 4),

$$\hat{y}_i = \begin{cases} 1 & \theta^t(\hat{x}_i) > \mu, \\ 0 & \text{else.} \end{cases} \quad (2)$$

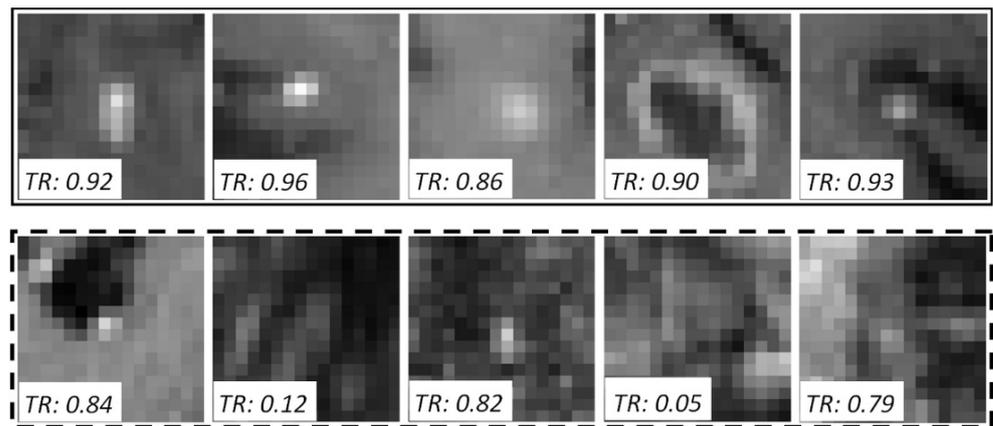


Figure 4. Mid-axial slices of an example pseudo-labeled batch of five pairs. The upper and lower rows show the positive and negative samples respectively. The response threshold (μ) is 0.85, and the teacher model responses ($TR = \theta^t(\hat{x}_i)$) are given with the images.

The student model is optimized using both labeled and pseudo-labeled data as

$$\operatorname{argmin}_{\theta^s} (\ell(\theta^s(X_{\text{noised}}), Y) + \lambda \ell(\theta^s(\hat{X}_{\text{noised}}), \hat{Y})), \quad (3)$$

where λ determines the weight of the unlabeled data on the final loss. In [12], their loss equally weighted the labeled and pseudo-labeled parts of the equation by normalizing these with the corresponding sample counts. Hence, their solution could be considered as a special case of our formulation with $\lambda = 1$.

After the generation of the student model, it replaces the teacher model. A new student model(s) could be trained iteratively following the same procedures. The suggested algorithm has three hyperparameters; μ , λ , and student model iterations. As μ is determined based on the teacher model's sensitivity for the training data, the usage of a value derived for the model's peak sensitivity is suggested. The motivation behind this choice is to allow the detection of most of the BM; even this may lead to more false-positive pseudo-labels. The alternative choice of setting μ for a lower sensitivity would lead to highly accurate pseudo-labels, whereas the challenging BM detections with low $\theta^t(\hat{x}_i)$ are excluded from the student model's training.

The adoption of λ value of 1.0 and performance of a single student model iteration are suggested, as each iteration requires an extensive amount of computational resources, where the resulting performance boost decays significantly after the first iteration. The effects of changing these parameters are presented via experimental studies.

2.4. Database

The data were collected retrospectively following Institutional Review Board Approval with a waiver of informed consent (institutional IRB ID: 2016H0084). The labeled data included 217 post-gadolinium T1-weighted 3D MRI exams (contrast agent: gadoterate

meglumine—0.1 mmol/kg of body weight) and their corresponding BM segmentation masks. The exams were acquired between January 2015 and February 2016 from 158 patients; 113 patients had a single, 33 patients had two, 10 patients had three, and two patients had four imaging exams. The group had (1) a mean age of 62 ($\sigma = 10.5$) and (2) a sex ratio (males:females) of 0.89. They were selected based on finalized brain MRI reports generated by a neuroradiologist, where the exams with reports detailing primarily parenchymal enhancing metastases, ideally those undergoing surveillance for radiation therapy/radiation treatment planning, were selected. Patients with metastases larger than 15 mm in diameter, primary brain neoplasms, central nervous system lymphoma, extra-axial disease, leptomeningeal disease, or equivocally enhancing foci were excluded. The BM segmentation masks were then generated by a fourth-year radiology resident, using the finalized examination report and/or annotated Picture Archiving and Communication System (PACS) images to ensure all BM were correctly delineated. The segmentation data included 932 annotated BM, where (1) the mean number of BM per patient was 4.29 ($\sigma = 5.52$), (2) the mean BM diameter was 5.45 mm ($\sigma = 2.67$ mm), and (3) the mean BM volume was 159.58 mm³ ($\sigma = 275.53$ mm³).

The unlabeled data included 1247 post-gadolinium T1-weighted 3D MRI exams, acquired between November 2016 and December 2019 from a non-overlapping group of 867 patients (i.e., no common patients between the labeled and unlabeled data). Of these, 579 patients had a single, 208 patients had two, 68 patients had three, and 12 patients had four MRI exams. The group had (1) a mean age of 56 ($\sigma = 14.5$) and (2) a sex ratio of 1.09. These patients were selected based on the diagnosis codes for malignant neoplasm(s) and/or secondary malignant brain neoplasm(s). A selected group of information of the labeled and unlabeled datasets is summarized in Table 1.

Table 1. Labeled and unlabeled data summary.

	Labeled Data	Unlabeled Data
Exam count	217	1247
Patient count	158	867
Acquisition date from	January 2015	November 2016
Acquisition date to	February 2016	December 2019
Mean patient age	62	56
Sex ratio	0.89	1.09

2.5. Validation Metric

The average number of false positives AFP (i.e., the count of falsely detected tumor locations) in connection with the detection sensitivity (i.e., the percentage of the actual BM detected) was used as the validation metric. A tumor was tagged as detected if the framework generated a detection up to 1.5 mm apart from the tumor's center. The metric provides a relevant measurement for the algorithm's applicability in real-life deployment scenarios, as (1) the sensitivity of a detection system is critical, and (2) the number of false positives needs to be low to ensure the system's feasibility for optimal clinical workflow. Accordingly, it was employed in various SOTA BM detection studies, including [3,4,6,33].

3. Results

3.1. Validation Study

The approach was validated using a two-fold CV. The CV bins were set based on patients (i.e., the data from each patient can only occur in either the training or testing bin), where each bin included the labeled data from 79 patients. For each CV-fold, with the LoG candidate detector tuned as in [4], the setup generated ~72K BM candidates and captured ~95% of the actual BM centers in the training bin. The teacher model was paired-trained using the labeled data in the training bin with its corresponding BM candidates. The teacher model's response threshold (μ) was determined by targeting the detection sensitivity (ts) of 90% for the training bin. Next, the student model was paired-trained using (1) the labeled

data in the training bin, (2) the unlabeled data that was pseudo-labeled by the teacher model, and (3) their corresponding BM candidates. The λ hyperparameter of the student model loss was set to 1.0. Finally, both the student and teacher models were tested for the testing bin (having no overlapping patients with the training bin and unlabeled data), generating the AFP versus sensitivity charts (see Figure 5). The teacher model produced AFPs of 2.95, 5.74, and 9.23 for the 80%, 85%, and 90% detection sensitivities, respectively. The student model produced AFPs of 2.91, 4.82, and 8.44 at the same sensitivity levels (please note that teacher model performance closely mimics the BM detector proposed in [4] as they both employed identical model architectures and were trained using similar labeled data). Accordingly, the AFP reduction for the 90% BM detection sensitivity was ~9% (from 9.23 to 8.44) for the introduced strategy. Figure 6 presents a set of example BM detections for the framework.

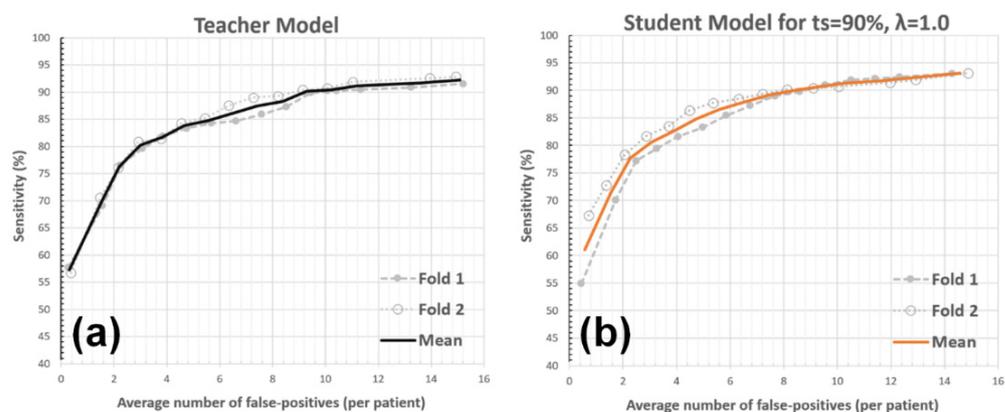


Figure 5. AFP versus detection sensitivity. The average number of false positives per patient (i.e., wrongly detected BM lesions for each patient) in relation to the sensitivity is illustrated for the teacher (a) and student (b) models. The mean curve (shown as the thick-line curve) represents the average of the CV folds.

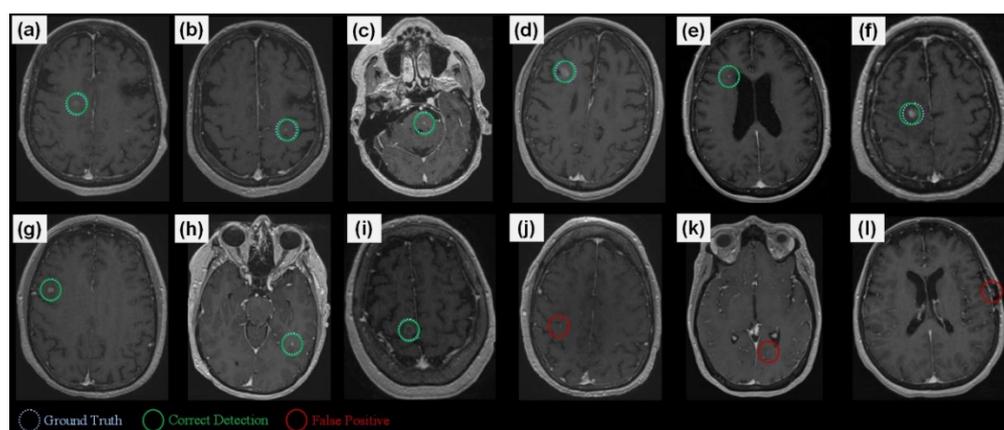


Figure 6. A set of detection examples. (a–i) Correct BM detections, (j,k) small vessels are wrongly detected as BMs, and (l) a formation in a surgical region is detected as a BM. The framework processes a given exam by first generating candidates and then processes these via CropNet, where the high response candidates are presented as the detection results.

The teacher (CropNet-b2-16mm) and student (CropNet-b4-16mm) networks consisted of ~14M and ~32M trainable parameters, respectively. The networks had architectures described in Section 2.2, processing cubic ROIs with 16 mm edges. The models' dropout rates were set at 0.15. The model optimizations were performed using the Adam algorithm [34] with (1) a learning rate of 0.00005 and (2) exponential decay rates for the first-

and second-moment estimates of 0.9 and 0.999, respectively. Both models were trained using 12,000 batch-training iterations, without an early stopping condition. Each batch included (1) 150 positive and 150 negative samples for the teacher models and (2) 75 positive, 75 pseudo-labeled positive, 75 negative, and 75 pseudo-labeled negative samples for the student models. The implementation was performed using the Python programming language (v3.6.10), and the neural networks were created using the TensorFlow (v1.12.0) library. The training times for the teacher and student models were ~ 7.25 and ~ 10.5 h, respectively, using an NVIDIA Tesla V100 graphics card with 32 GB RAM.

3.2. Experiments with System Parameters

The introduced noisy-student-based self-training strategy has three input parameters (i.e., μ , λ , and student model iterations); the reasoning behind their default values was provided previously. Three analyses were performed to illustrate the effects of these parameters, and an additional analysis was performed to present the impact of labeled data amount in our application.

3.2.1. Teacher Model Sensitivity

The response threshold (μ) for setting the pseudo-labels is determined based on the teacher model's detection sensitivity (t_s) for the training data (see Equation (2)). In our proposed solution, t_s was set as 90%, as the value is close to the maximum BM detection sensitivity the teacher model can achieve. In this analysis, μ values were determined by setting the teacher model's detection sensitivity to 80% and 85% for the training bin; hence, the impact of more-specific yet less-sensitive pseudo-labeling was evaluated. The other parameters (i.e., $\lambda = 1$ and a single student model iteration) were kept unchanged during the experiment. Figure 7 and Table 2 present the findings of this analysis.

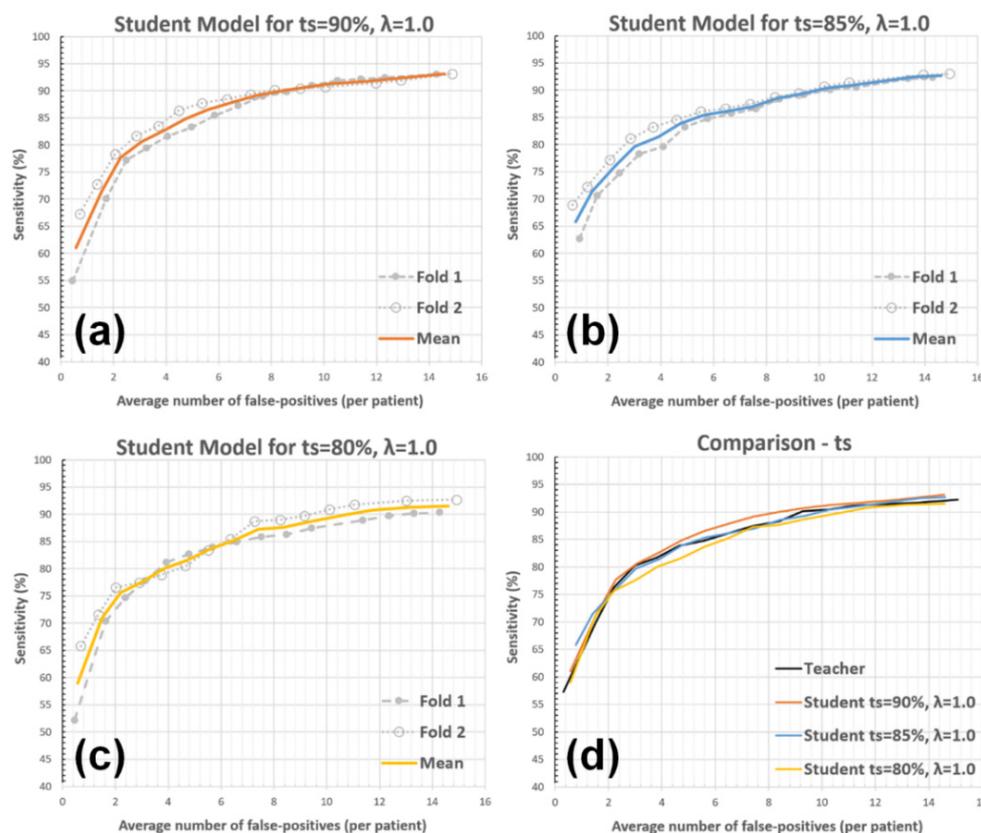


Figure 7. The impact of response threshold (μ) on the AFP metric. The student model performances for (a) $t_s = 90\%$ —(default setup), (b) $t_s = 85\%$, (c) $t_s = 80\%$, and (d) comparisons between the teacher and student models with these setups are presented.

Table 2. AFP vs. sensitivity for different ts values.

Detection Sensitivity	Teacher	Student ts = 90%	Student ts = 85%	Student ts = 80%
80%	2.94	2.91	3.18	3.82
85%	5.74	4.82	5.41	6.37
90%	9.23	8.44	9.97	10.84

3.2.2. Unlabeled Data Weights (λ)

The analysis aimed to represent the impact of reducing pseudo-labeled data weight during the student model’s training. This was achieved by evaluating the student models derived with $\lambda = 0.8$ and $\lambda = 0.6$ ($\lambda = 1.0$ is the default value, see Equation (3)). The other parameters were kept at their default values during the experiment (i.e., ts = 90% and a single student model iteration). Please see Figure 8 and Table 3 for the results.

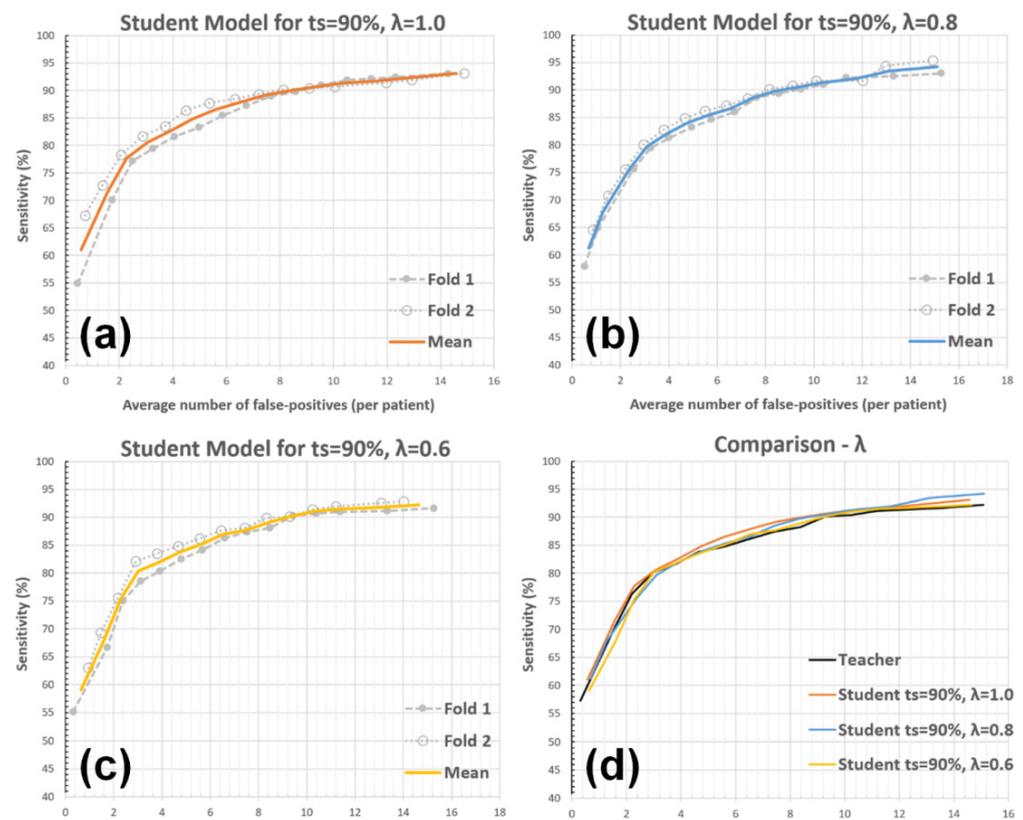


Figure 8. The impact of λ on the AFP metric. The student model performances for (a) $\lambda = 1.0$ (default setup), (b) $\lambda = 0.8$, (c) $\lambda = 0.6$, and (d) comparisons between the teacher and student models with these setups are presented.

Table 3. AFP vs. sensitivity for different λ values.

Detection Sensitivity	Teacher	Student $\lambda = 1.0$	Student $\lambda = 0.8$	Student $\lambda = 0.6$
80%	2.94	2.91	3.19	2.97
85%	5.74	4.82	5.36	5.55
90%	9.23	8.44	8.71	9.12

Results with ts = 90% and a single student model iteration. The gray column corresponds to the default setup.

3.2.3. Student Model Iterations

The experiment was targeted to assess the effect of further student training iterations, which are not part of our default solution due to low expected performance gains with a high additional computational cost. More explicitly, the student model from iteration-n (i.e., θ_n^s) was used as the teacher model in the next iteration; θ_n^s becomes θ_{n+1}^t and generates

pseudo-labeled data for the training of θ_{n+1}^s . Two additional training and testing iterations were performed for the experiment (i.e., the student model iterations 2 and 3), where the other default parameters were kept unchanged ($ts = 90\%$ and $\lambda = 1.0$). Please see Figure 9 and Table 4 for the corresponding results.

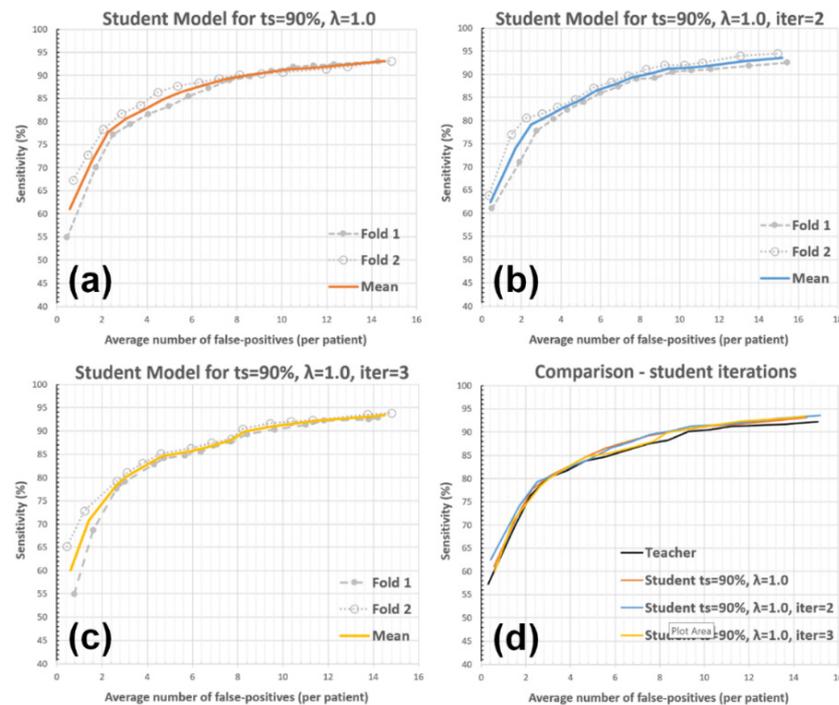


Figure 9. The impact of student model training iterations on the AFP metric. The student model performances for (a) single training iteration (default setup), (b) two iterations, (c) three iterations, and (d) comparisons between the teacher and student models with these setups are presented.

Table 4. AFP vs. sensitivity for different student model iterations.

Detection Sensitivity	Teacher	Student Iter = 1	Student Iter = 2	Student Iter = 3
80%	2.94	2.91	2.89	3.03
85%	5.74	4.82	5.20	5.09
90%	9.23	8.44	8.35	8.46

Results with $ts = 90\%$ and $\lambda = 1.0$. The gray column corresponds to the default setup.

3.2.4. Labeled-Data Utilization

Finally, the teacher and student models with a reduced number of labeled training data were evaluated. First, a random selection of 75% of patients were kept in the training bins (i.e., 60 patients), and the teacher and student models were trained using the labeled data collected from these patients and the complete set of unlabeled data. The test bins were kept unchanged, and the default noisy-student parameters were used during the analysis (i.e., $\lambda = 1$, $ts = 90\%$, and a single student model iteration). Next, the experiment was repeated by randomly keeping 50% of patients in the training bins (i.e., 40 patients). Figure 10 and Table 5 present the findings of this experiment.

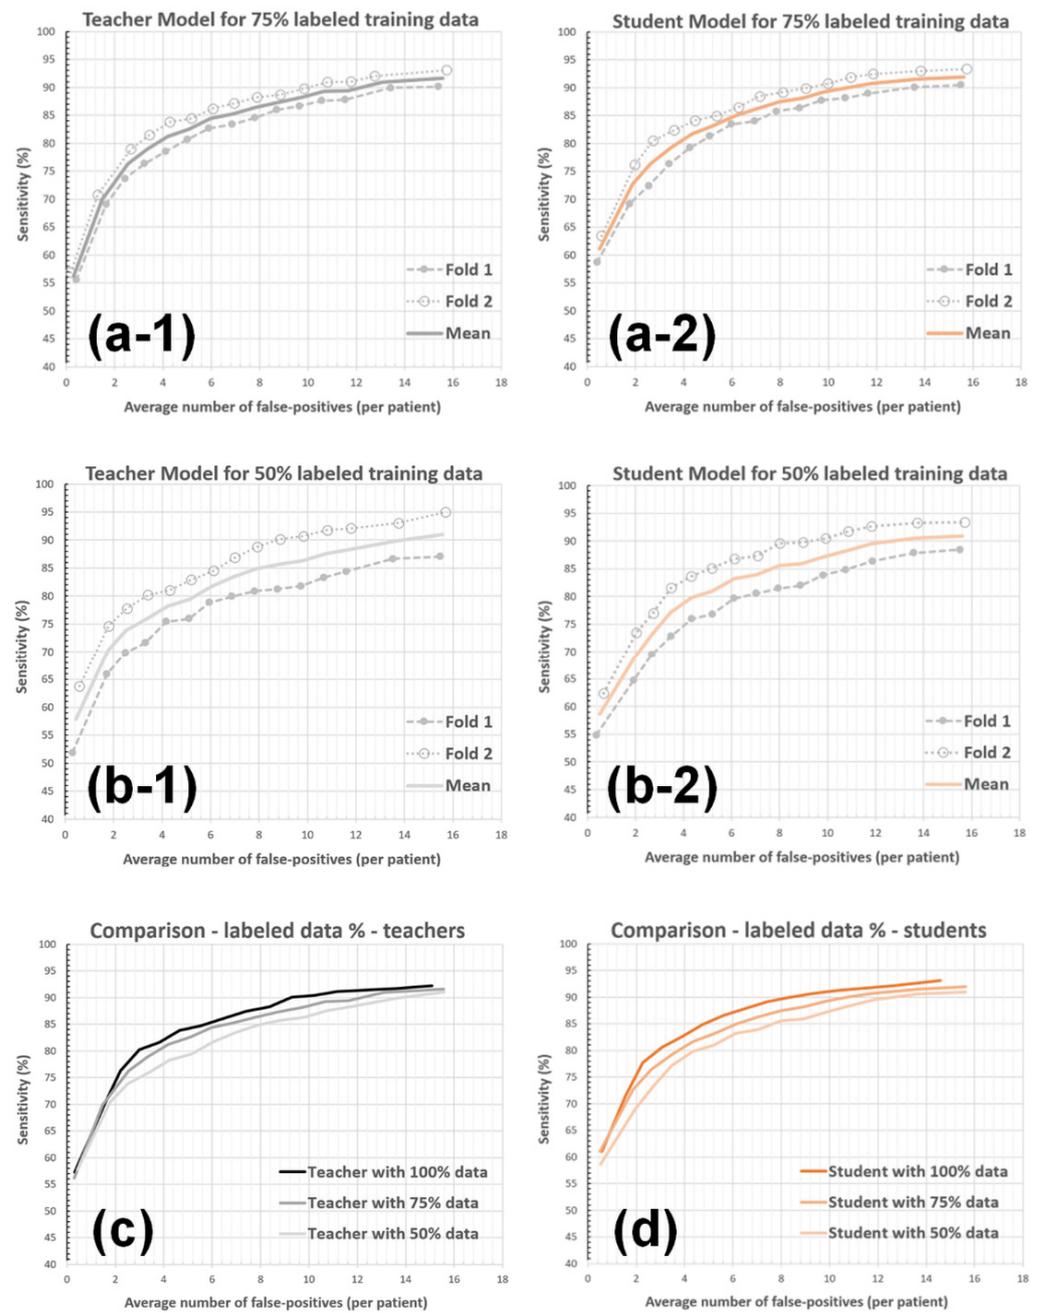


Figure 10. The impact of the amount of labeled data on the AFP metric. The teacher (a-1) and student (a-2) models were trained using 75% of the labeled data. The teacher (b-1) and student (b-2) models were trained using 50% of the labeled data. The teacher (c) and student (d) models were compared for different labeled-data utilization percentages.

Table 5. AFP vs. sensitivity for different amount of labeled-data (LD) utilization percentages.

Detection Sensitivity	Teacher LD:100%	Student LD:100%	Teacher LD:75%	Student LD:75%	Teacher LD:50%	Student LD:50%
80%	2.94	2.91	3.74	3.73	5.36	4.46
85%	5.74	4.82	6.60	6.15	8.05	7.65
90%	9.23	8.44	12.19	10.79	13.89	12.37

Results with $t_s = 90\%$, $\lambda = 1.0$, a single student model iteration. The gray column corresponds to the default setup.

4. Discussion

The validation study showed that the sensitivity-based NS training improves the framework performance by reducing the AFP count (i.e., 9.23 to 8.44 at 90% and 5.74 to 4.82 at 85% BM detection sensitivities). The algorithm's parameters need to be set properly to exploit its potential: (1) using the peak teacher model sensitivity (t_s) for the training data to determine μ , (2) adopting $\lambda = 1$, and (3) performing a single student training iteration led to satisfactory results in our experiments; the motivations for each of these choices and relevant experiments are provided.

The first experiment showed that reduced teacher model sensitivity during pseudo-labeling (i.e., yielding more-accurate yet less-sensitive pseudo-labeled data) leads to higher AFP counts. As shown in Table 2, setting t_s value to 80% caused the framework to produce the AFP value of 10.84 at 90% detection sensitivity (AFP values for 80% and 85% detection sensitivities for this setup were also relatively higher compared with the results acquired for $t_s = 90\%$). The NS-based training's effectiveness is partly due to the implicit data augmentation/enrichment it introduces. Accordingly, one may argue that keeping the teacher model highly sensitive allowed student model training with pseudo-labeled data, representing high variability. Thus, this higher level of implicit data augmentation led to an improved BM detection performance.

The second experiment aimed to find the effect of reducing the weight of pseudo-labeled data during the student model training. Two alternative outcomes for this analysis were hypothesized before the experiment: (1) λ value behaves as an interpolator parameter between the teacher and student models, and (2) a reduced λ (<1.0) may cause the student model to outperform the default setup (using $\lambda = 1.0$) by limiting the drift towards pseudo-samples, which have limited accuracy. The analysis' outcome supported the first of these (i.e., interpolation hypothesis). Table 3 shows that the students generated with $\lambda = 0.6$ and $\lambda = 0.8$ had performances bounded by the teacher and default student models, where the $\lambda = 0.6$ performed closer to the teacher model and $\lambda = 0.8$ performed closer to the default student model, respectively. Thus, the experiment's result also complements the first experiment by supporting the notion that a higher amount of novel information contributes to a better student model performance.

The third experiment examined the student model's performance with further student training iterations. The highest framework performances were observed after the first and second student model iterations, whereas the performance did not improve significantly after the third iteration (see Table 4). Our finding is consistent with [12], showing that the student model performance converges to a value at an arbitrary student model iteration, beyond which the accuracy does not change noticeably.

The final experiment aimed to quantify the impact of labeled training data amount on the framework's accuracy. Accordingly, the teacher/student model training and validation procedures were repeated using reduced percentages of the labeled data (i.e., 75% and 50%). The results show that the amount of labeled data is critical to the final performance; the student models produced 8.44, 10.79, and 12.37 false BM detections by utilizing 100%, 75%, and 50% of the labeled training data. Furthermore, a valuable performance improvement with the use of proposed NS-based training was observed; it led to ~9% (from 9.23 to 8.44), ~12% (from 12.19 to 10.79), and ~11% (from 13.89 to 12.37) reduction in AFP counts in these experiments (see Table 5).

Table 6 presents comparative information between a group of SOTA BM detection frameworks (i.e., [3–8,33]) and the proposed one. It provides (1) study patient counts, (2) neural network architectures, (3) data acquisition types, (4) dimensional information on detected tumors, (5) data partitioning during validation studies, and (6) test performances with regards to AFP and sensitivity. Please note that the table includes results from [4] that used a similar dataset during its validation as this study. The data show that the potential advantages of our framework include (1) the detection of particularly small BM lesions with relatively low AFP and high sensitivity and (2) the capability to utilize unlabeled datasets during its model training via NS-based strategy.

Table 6. Overview of selected group of SOTA BM segmentation/detection approaches.

Study	Patient #	Network	Acq.	BM Volume (mm ³)	Validation Type	Sens %	AFP
Charron et al. [3]	182-labeled	DeepMedic	Multi seq. ^a	Mean: 2400 Median: 500	Fixed train/test	93	7.8
Liu et al. [5]	490-labeled	En-DeepMedic	Multi seq. ^b	Mean: 672	5-fold CV	NA	NA
Bousabarah et al. [8]	509-labeled	U-Net	Multi seq. ^c	Mean: 1920 Median: 470	Fixed train/test	77–82	<1
Grøvik et al. [6]	156-labeled	GoogleNet	Multi seq. ^d	NA	Fixed train/test	83	8.3
Cao et al. [7]	195-labeled	Asym-UNet	T1cMRI	NA ^g	Fixed train/test	81–100 ^k	NA
Zhou et al. [33]	266-labeled	Custom + SSD	T1c MRI ^e	NA ^h	Fixed train/test	81	6
Dikici et al. [4]	158-labeled ⁱ	CropNet	T1c MRI	Mean: 160 Median: 50	5-fold CV	90	9.12
This study	158-labeled 867-unlabeled	CropNet+NS	T1c MRI ^f	Mean: 160 ^j Median: 50	2-fold CV	90	8.44

^a T1-weighted 3D MRI with gadolinium injection, T2-weighted 2D FLAIR, and T1-weighted 2D MRI sequences.

^b T1c and T2-weighted FLAIR sequences. ^c T1c, T2-weighted, and T2-weighted FLAIR sequences. ^d Pre- and post-gadolinium CUBE, post-gadolinium T1-weighted 3D axial IR-prepped FSPGR (BRAVO), and 3D CUBE FLAIR sequences. ^e 3D T1-weighted contrast-enhanced spoiled gradient-echo MRI sequence. ^f The same dataset was used for training and validation in this study and [13]. ^g Tumor volumes were not reported. The testing was conducted for 72 smaller tumors (1–10 mm in diameter) and 17 larger tumors (11–26 mm in diameter) separately. ^h Tumor volumes were not reported. The average size of tumors in the study were 10 mm ± (standard deviation: 8 mm). ^j The volumetric stats are for the labeled data. ^k Sensitivity was 81 percent for smaller tumors and 100 percent for larger ones. ⁱ The labeled dataset used in the study is similar with the one used in this current study.

5. Conclusions

In this study, a novel formulation of the NS-based self-training strategy, which is applicable for detection systems prioritizing low false-positive counts for arbitrary detection sensitivities, was introduced. Next, the BM detection framework for 3D T1c data was extended with the introduced strategy. The analyses showed that the method reduced AFP by 9% (from 9.23 to 8.44) for 90% tumor-detection sensitivity by utilizing 1247 unlabeled T1c exams. Future studies may (1) further validate the approach with larger and multi-institutional unlabeled datasets and (2) investigate the algorithm's integration into other medical applications, such as the detection of primary brain tumors [35].

The noising procedure is essential in NS-based training strategies as it enables knowledge expansion. In data-limited clinical applications, its importance becomes even more prominent: the utilized highly parametric models could greatly benefit from the process, which provides data and model noise, to reduce the overfitting and hence improve their generalizability [36]. A wide range of noising mechanisms can be employed during the NS-based training, and the most appropriate selection may depend on (1) the data (i.e., type, amount, and characteristics) and (2) the deployed ML models. For instance, [12] utilized the stochastic depth approach for model noising [37], which is applicable for very deep neural networks (e.g., [38]), unlike the CropNet. The impact of a set of pertinent mechanisms such as (1) Mixup [16], (2) constrained GAN ensembles [39], and (3) stochastic dropouts [40] represent potential avenues for future research on the model noising aspect of the framework.

The applicability of artificial intelligence (AI) in medicine (among other fields) has increased significantly due to advancements in ML over the recent years. Accordingly, there have been various efforts to integrate AI in radiology workflows [41,42], where the AI-based medical imaging algorithms are deployed and updated periodically by utilizing a constant flow of annotated/labeled data. These workflows may also be adapted to use unlabeled data by adopting self-training methodologies, such as the one described in this study. Hence, the deployed medical imaging models are highly relevant to radiology AI systems that could benefit from the massive collection of unlabeled data sets stored in PACS systems.

Author Contributions: Conceptualization, E.D. and L.M.P.; methodology, E.D.; software, E.D.; validation, E.D. and X.V.N.; formal analysis, E.D. and X.V.N.; investigation, E.D.; resources, L.M.P.; data curation, E.D., J.L.R. and M.B.; writing—original draft preparation, E.D.; writing—review and editing, E.D., X.V.N., M.B., J.L.R. and L.M.P.; visualization, E.D.; supervision, L.M.P.; project administration, E.D.; funding acquisition, L.M.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Institutional Review Board Statement: The study was conducted in accordance with the Declaration of Helsinki and approved by the Institutional Review Board of the Ohio State University College of Medicine. The data were collected retrospectively following Institutional Review Board Approval with a waiver of informed consent (institutional IRB ID: 2016H0084, approved date: 31 August 2021).

Informed Consent Statement: Not applicable.

Data Availability Statement: The data is not available for public access.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Lester, S.C.; Taksler, G.B.; Kuremsky, J.G.; Lucas, J.T., Jr.; Ayala-Peacock, D.N.; Randolph, D.M.; Bourland, J.D.; Laxton, A.W.; Tatter, S.B.; Chan, M.D. Clinical and Economic Outcomes of Patients with Brain Metastases Based on Symptoms: An Argument for Routine Brain Screening of Those Treated with Upfront Radiosurgery. *Cancer* **2014**, *120*, 433–441. [[CrossRef](#)] [[PubMed](#)]
2. Tong, E.; McCullagh, K.L.; Iv, M. Advanced Imaging of Brain Metastases: From Augmenting Visualization and Improving Diagnosis to Evaluating Treatment Response. *Front. Neurol.* **2020**, *11*, 270. [[CrossRef](#)] [[PubMed](#)]
3. Charron, O.; Lallement, A.; Jarnet, D.; Noblet, V.; Clavier, J.-B.; Meyer, P. Automatic Detection and Segmentation of Brain Metastases on Multimodal MR Images with a Deep Convolutional Neural Network. *Comput. Biol. Med.* **2018**, *95*, 43–54. [[CrossRef](#)] [[PubMed](#)]
4. Dikici, E.; Ryu, J.L.; Demirer, M.; Bigelow, M.; White, R.D.; Slone, W.; Erdal, B.S.; Prevedello, L.M. Automated Brain Metastases Detection Framework for T1-Weighted Contrast-Enhanced 3D MRI. *IEEE J. Biomed. Health Inform.* **2020**, *24*, 2883–2893. [[CrossRef](#)]
5. Liu, Y.; Stojadinovic, S.; Hrycushko, B.; Wardak, Z.; Lau, S.; Lu, W.; Yan, Y.; Jiang, S.B.; Zhen, X.; Timmerman, R.; et al. A Deep Convolutional Neural Network-Based Automatic Delineation Strategy for Multiple Brain Metastases Stereotactic Radiosurgery. *PLoS ONE* **2017**, *12*, e0185844. [[CrossRef](#)]
6. Grovik, E.; Yi, D.; Iv, M.; Tong, E.; Rubin, D.; Zaharchuk, G. Deep Learning Enables Automatic Detection and Segmentation of Brain Metastases on Multisequence MRI. *J. Magn. Reson. Imaging* **2019**, *51*, 175–182. [[CrossRef](#)]
7. Cao, Y.; Vasantachart, A.; Jason, C.Y.; Yu, C.; Ruan, D.; Sheng, K.; Lao, Y.; Shen, Z.L.; Balik, S.; Bian, S.; et al. Automatic Detection and Segmentation of Multiple Brain Metastases on Magnetic Resonance Image Using Asymmetric UNet Architecture. *Phys. Med. Biol.* **2021**, *66*, 15003. [[CrossRef](#)]
8. Bousabarah, K.; Ruge, M.; Brand, J.-S.; Hoevens, M.; Rueß, D.; Borggrefe, J.; Große Hokamp, N.; Visser-Vandewalle, V.; Maintz, D.; Treuer, H.; et al. Deep Convolutional Neural Networks for Automated Segmentation of Brain Metastases Trained on Clinical Data. *Radiat. Oncol.* **2020**, *15*, 1–9. [[CrossRef](#)]
9. Cho, S.J.; Sunwoo, L.; Baik, S.H.; Bae, Y.J.; Choi, B.S.; Kim, J.H. Brain Metastasis Detection Using Machine Learning: A Systematic Review and Meta-Analysis. *Neuro-Oncol.* **2021**, *23*, 214–225. [[CrossRef](#)]
10. Mongan, J.; Moy, L.; Kahn, C.E., Jr. Checklist for Artificial Intelligence in Medical Imaging (CLAIM): A Guide for Authors and Reviewers. *Radiol. Artif. Intell.* **2020**, *2*, e200029. [[CrossRef](#)]
11. Whiting, P.F.; Rutjes, A.W.S.; Westwood, M.E.; Mallett, S.; Deeks, J.J.; Reitsma, J.B.; Leeflang, M.M.G.; Sterne, J.A.C.; Bossuyt, P.M.M. QUADAS-2: A Revised Tool for the Quality Assessment of Diagnostic Accuracy Studies. *Ann. Intern. Med.* **2011**, *155*, 529–536. [[CrossRef](#)] [[PubMed](#)]
12. Xie, Q.; Luong, M.-T.; Hovy, E.; Le, Q. V Self-Training with Noisy Student Improves Imagenet Classification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020; pp. 10687–10698.
13. Ma, J.; Zhang, Y.; Gu, S.; Zhang, Y.; Zhu, C.; Wang, Q.; Liu, X.; An, X.; Ge, C.; Cao, S.; et al. AbdomenCT-1K: Is Abdominal Organ Segmentation a Solved Problem? *arXiv* **2020**, arXiv:2010.14808. [[CrossRef](#)] [[PubMed](#)]
14. Isensee, F.; Jäger, P.F.; Kohl, S.A.A.; Petersen, J.; Maier-Hein, K.H. Automated Design of Deep Learning Methods for Biomedical Image Segmentation. *arXiv* **2019**, arXiv:1904.08128.
15. Rajan, D.; Thiagarajan, J.J.; Karargyris, A.; Kashyap, S. Self-Training with Improved Regularization for Sample-Efficient Chest X-ray Classification. In *Medical Imaging 2021: Computer-Aided Diagnosis, Proceedings of the SPIE Medical Imaging 2021, San Diego, CA, USA, 14–18 February 2021*; SPIE: Cergy, France, 2021; Volume 11597, p. 115971S.
16. Zhang, H.; Cisse, M.; Dauphin, Y.N.; Lopez-Paz, D. Mixup: Beyond Empirical Risk Minimization. *arXiv* **2017**, arXiv:1710.09412.

17. 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, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
18. Shak, K.; Al-Shabi, M.; Liew, A.; Lan, B.L.; Chan, W.Y.; Ng, K.H.; Tan, M. A New Semi-Supervised Self-Training Method for Lung Cancer Prediction. *arXiv* **2020**, arXiv:2012.09472.
19. Li, Y.; Fan, Y. DeepSEED: 3D Squeeze-and-Excitation Encoder-Decoder Convolutional Neural Networks for Pulmonary Nodule Detection. In Proceedings of the 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI), Iowa City, IA, USA, 3–7 April 2020; pp. 1866–1869.
20. Armato III, S.G.; McLennan, G.; Bidaut, L.; McNitt-Gray, M.F.; Meyer, C.R.; Reeves, A.P.; Zhao, B.; Aberle, D.R.; Henschke, C.I.; Hoffman, E.A.; et al. The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): A Completed Reference Database of Lung Nodules on CT Scans. *Med. Phys.* **2011**, *38*, 915–931. [[CrossRef](#)] [[PubMed](#)]
21. Team, N.L.S.T.R. The National Lung Screening Trial: Overview and Study Design. *Radiology* **2011**, *258*, 243–253.
22. Lin, E.; Kuo, W.; Yuh, E. Noisy Student Learning for Cross-Institution Brain Hemorrhage Detection. *arXiv* **2021**, arXiv:2105.00582.
23. Kuo, W.; Häne, C.; Yuh, E.; Mukherjee, P.; Malik, J. PatchFCN for Intracranial Hemorrhage Detection. *arXiv* **2018**, arXiv:1806.03265.
24. Chilamkurthy, S.; Ghosh, R.; Tanamala, S.; Biviji, M.; Campeau, N.G.; Venugopal, V.K.; Mahajan, V.; Rao, P.; Warier, P. Development and Validation of Deep Learning Algorithms for Detection of Critical Findings in Head CT Scans. *arXiv* **2018**, arXiv:1803.05854.
25. Kim, G.; Park, S.; Oh, Y.; Seo, J.B.; Lee, S.M.; Kim, J.H.; Moon, S.; Lim, J.-K.; Ye, J.C. Severity Quantification and Lesion Localization of COVID-19 on CXR Using Vision Transformer. *arXiv* **2021**, arXiv:2103.07062.
26. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv* **2020**, arXiv:2010.11929.
27. Irvin, J.; Rajpurkar, P.; Ko, M.; Yu, Y.; Ciurea-Ilcus, S.; Chute, C.; Marklund, H.; Haghighi, B.; Ball, R.; Shpanskaya, K.; et al. Chexpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. In Proceedings of the AAAI Conference on Artificial Intelligence, Honolulu, HI, USA, 27 January–1 February 2019; Volume 33, pp. 590–597.
28. Shiraishi, J.; Katsuragawa, S.; Ikezoe, J.; Matsumoto, T.; Kobayashi, T.; Komatsu, K.; Matsui, M.; Fujita, H.; Kodera, Y.; Doi, K. Development of a Digital Image Database for Chest Radiographs with and without a Lung Nodule: Receiver Operating Characteristic Analysis of Radiologists' Detection of Pulmonary Nodules. *Am. J. Roentgenol.* **2000**, *174*, 71–74. [[CrossRef](#)] [[PubMed](#)]
29. Wang, W.; Xia, Q.; Hu, Z.; Yan, Z.; Li, Z.; Wu, Y.; Huang, N.; Gao, Y.; Metaxas, D.; Zhang, S. Few-Shot Learning by a Cascaded Framework with Shape-Constrained Pseudo Label Assessment for Whole Heart Segmentation. *IEEE Trans. Med. Imaging* **2021**, *40*, 2629–2641. [[CrossRef](#)]
30. Milletari, F.; Navab, N.; Ahmadi, S.-A. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In Proceedings of the 2016 Fourth International Conference on 3D Vision (3DV), Stanford, CA, USA, 25–28 October 2016; pp. 565–571.
31. Lindeberg, T. Scale Selection Properties of Generalized Scale-Space Interest Point Detectors. *J. Math. Imaging Vis.* **2013**, *46*, 177–210. [[CrossRef](#)]
32. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Munich, Germany, 5–9 October 2015; Springer: Cham, Switzerland; pp. 234–241.
33. Zhou, Z.; Sanders, J.W.; Johnson, J.M.; Gule-Monroe, M.K.; Chen, M.M.; Briere, T.M.; Wang, Y.; Son, J.B.; Pagel, M.D.; Li, J.; et al. Computer-Aided Detection of Brain Metastases in T1-Weighted MRI for Stereotactic Radiosurgery Using Deep Learning Single-Shot Detectors. *Radiology* **2020**, *295*, 407–415. [[CrossRef](#)]
34. Kingma, D.; Ba, J. Adam: A Method for Stochastic Optimization. In Proceedings of the International Conference on Learning Representation, Banff, AB, Canada, 14–16 April 2014.
35. Louis, D.N.; Perry, A.; Reifenberger, G.; Von Deimling, A.; Figarella-Branger, D.; Cavenee, W.K.; Ohgaki, H.; Wiestler, O.D.; Kleihues, P.; Ellison, D.W. The 2016 World Health Organization Classification of Tumors of the Central Nervous System: A Summary. *Acta Neuropathol.* **2016**, *131*, 803–820. [[CrossRef](#)]
36. Mutasa, S.; Sun, S.; Ha, R. Understanding Artificial Intelligence Based Radiology Studies: What Is Overfitting? *Clin. Imaging* **2020**, *65*, 96–99. [[CrossRef](#)]
37. Huang, G.; Sun, Y.; Liu, Z.; Sedra, D.; Weinberger, K.Q. Deep Networks with Stochastic Depth. In Proceedings of the European Conference on Computer Vision, Amsterdam, The Netherlands, 8–16 October 2016; pp. 646–661.
38. Tan, M.; Le, Q. Efficientnet: Rethinking Model Scaling for Convolutional Neural Networks. In Proceedings of the International Conference on Machine Learning, Long Beach, CA, USA, 10–15 June 2019; pp. 6105–6114.
39. Dikici, E.; Bigelow, M.; White, R.D.; Erdal, B.S.; Prevedello, L.M. Constrained Generative Adversarial Network Ensembles for Sharable Synthetic Medical Images. *J. Med. Imaging* **2021**, *8*, 24004. [[CrossRef](#)]
40. Park, S.; Kwak, N. Analysis on the Dropout Effect in Convolutional Neural Networks. In Proceedings of the Asian Conference on Computer Vision, Taipei, Taiwan, 20–24 November 2016; pp. 189–204.
41. Cohen, R.Y.; Sodickson, A.D. An Orchestration Platform That Puts Radiologists in the Driver's Seat of AI Innovation: A Methodological Approach. *arXiv* **2021**, arXiv:2107.04409.
42. Dikici, E.; Bigelow, M.; Prevedello, L.M.; White, R.D.; Erdal, B.S. Integrating AI into Radiology Workflow: Levels of Research, Production, and Feedback Maturity. *J. Med. Imaging* **2020**, *7*, 16502. [[CrossRef](#)] [[PubMed](#)]