

HHS Public Access

Author manuscript *Proc SPIE Int Soc Opt Eng.* Author manuscript; available in PMC 2019 October 10.

Published in final edited form as: *Proc SPIE Int Soc Opt Eng.* 2019 March ; 10949: . doi:10.1117/12.2511997.

Distributed deep learning for robust multi-site segmentation of CT imaging after traumatic brain injury

Samuel Remedios^{1,2,3,4}, Snehashis Roy^{1,2}, Justin Blaber⁴, Camilo Bermudez⁵, Vishwesh Nath⁶, Mayur B. Patel⁷, John A. Butman², Bennett A. Landman^{4,5,6}, Dzung L. Pham^{1,2} ¹Center for Neuroscience and Regenerative Medicine, Henry Jackson Foundation

²Radiology and Imaging Sciences, Clinical Center, National Institute of Health

³Department of Computer Science, Middle Tennessee State University

⁴Department of Electrical Engineering, Vanderbilt University

⁵Department of Biomedical Engineering, Vanderbilt University

⁶Department of Computer Science, Vanderbilt University

⁷Departments of Surgery, Neurosurgery, Hearing & Speech Sciences; Center for Health Services Research, Vanderbilt Brain Institute; Critical Illness, Brain Dysfunction, and Survivorship Center, Vanderbilt University Medical Center; VA Tennessee Valley Healthcare System, Department of Veterans Affairs Medical Center

Abstract

Machine learning models are becoming commonplace in the domain of medical imaging, and with these methods comes an ever-increasing need for more data. However, to preserve patient anonymity it is frequently impractical or prohibited to transfer protected health information (PHI) between institutions. Additionally, due to the nature of some studies, there may not be a large public dataset available on which to train models. To address this conundrum, we analyze the efficacy of transferring the model itself in lieu of data between different sites. By doing so we accomplish two goals: 1) the model gains access to training on a larger dataset that it could not normally obtain and 2) the model better generalizes, having trained on data from separate locations. In this paper, we implement multi-site learning with disparate datasets from the National Institutes of Health (NIH) and Vanderbilt University Medical Center (VUMC) without compromising PHI. Three neural networks are trained to convergence on a computed tomography (CT) brain hematoma segmentation task: one only with NIH data, one only with VUMC data, and one multi-site model alternating between NIH and VUMC data. Resultant lesion masks with the multi-site model attain an average Dice similarity coefficient of 0.64 and the automatically segmented hematoma volumes correlate to those done manually with a Pearson correlation coefficient of 0.87, corresponding to an 8% and 5% improvement, respectively, over the single-site model counterparts.

Keywords

multi-site; distributed; deep learning; neural network; computed tomography (CT); hematoma; lesion; segmentation

1. INTRODUCTION

Deep learning has recently become a key approach for computer vision and medical imaging problems. Neural networks have been used to skull-strip CT scans,¹ segment magnetic resonance images,² locate and segment blood vessels,³ as well as segment brain regions⁴ and lesions.⁵ A wide variety of models and architectures have been implemented to solve these tasks, and there also exist pre-trained models prepared for general use cases.⁶ Regardless of the particular task for which a model is designed or selected, machine learning methods generally benefit from the inclusion of more data for training and validating the model.⁷ Traditionally, acquisition of multi-site data involves data transfer to a centralized location on which the desired model trains; however, it is frequently prohibited or difficult to acquire HIPAA-compliant health data transfer permits.⁸ These data restrictions are vital, though, as protected health information (PHI) policies enforce respect for patient privacy and anonymity.^{9,10,11} Herein lies a contradiction: machine learning models benefit greatly from a wealth of data, yet datasets related to healthcare cannot be shared between sites easily.

To address this problem, we propose to transfer the models themselves between sites in lieu of a dataset transfer. The concept of distributed learning is not new to machine learning, with one such example coming from Google's implementation of Federated Learning, through which models are averaged between mobile phones.¹² However, this approach does not have the goal of gaining accuracy or generalizability, and instead is a decentralized framework geared towards mobile devices and their limited computing power. Another distributed learning technique is transfer learning, ¹³ which aims to apply useful features learned from one task towards a kick-started learning for some other task. Different still is the concept of asynchronous stochastic gradient descent, ¹⁴ wherein a model is copied for some number of splits of training data, and their learned weights are aggregated once training is complete.

Recently, a study has embarked to investigate whether a model can perform better if it accesses data from different sites,¹⁵ wherein the authors simulate a multi-site scenario by splitting an open-source dataset into groups and apply different transformations and noise to each group with the goal of making the data appear different. The authors investigate applying different multi-site training approaches, comparing transfer learning to different patterns of passing partially trained models.

In this paper, we expand upon this work by using empirical multi-site data, separately acquired from the NIH and VUMC. Because of differences in the acquisition at each site, as well as in delineation protocols, improved performance due to the combined training data gained by multi-site learning is not guaranteed. Thus, we employ the aforementioned paper's cyclic weight transfer as our training paradigm and forgo the uni-directional transfer learning approach.

Our specific contributions are the presentation of an extensible framework through which multiple sites can train the same model using private data and the validation of the efficacy of two different training schema on the segmentation of hematoma in traumatic brain injury (TBI) CT scans. In the latter contribution, we consider single-site learning at each of the two sites (NIH and VUMC), and multi-site learning between both sites.

Remedios et al.

Here, we target segmentation of hemorrhages and hematomas in patients with TBI (see Figure 1). Hemorrhages refer to active bleeding, while a hematoma is any collection or swelling of clotted blood outside of the blood vessels, the cause of which could be severe trauma or disease. The identification and segmentation of blood is an important consideration for diagnosis, prediction of patient recovery, and for examining correlations with long-term neurologic disabilities¹⁶ such as cognitive impairment.¹⁷ Improving the efficacy of hematoma segmentation will therefore assist developments in understanding and treating TBI.

2. METHOD

2.1 Data

CT images from 27 acute TBI patients presenting with intracranial hematomas were acquired as part of a research study by the Center for Neuroscience and Regenerative Medicine (CNRM) and NIH. At VUMC, 18 CT images of TBI patients were obtained in deidentified form. The resolutions of all scans from both sites were approximately $0.5 \times 0.5 \times$ 5.0 mm³. All scans were converted from DICOM to NIFTI and subsequently transformed into Hounsfield units. For training, 10 scans were used at the VUMC site while 17 were used at the NIH; the remaining 8 and 10, respectively, were set aside as the test dataset. Images from both the NIH and VUMC had a variety of hematoma types, sizes, and locations; however VUMC on average had a larger hematoma volume of 41,000 mm³ compared with 13,700 mm³ in the NIH dataset. For preprocessing, all CT image volumes underwent skullstripping by CT BET¹⁸ and were rigidly transformed to a common orientation. To address the low number of training images, we collected $1,000\ 255 \times 255\ 2D$ patches from each CT volume, 20% of which were used as a validation set for hyperparameter tuning. Since voxel intensities were in Hounsfield units, no normalization was applied and thus no intensities were scaled. Additionally, because the images have low through-plane resolution (5.0 mm) compared to the in-plane resolution (0.5 mm), only 2D segmentations were considered. Manual segmentations were performed by independent raters at the two sites and reviewed independently by a neuroradiologist; quantities are reflected in Table 1.

2.2 Model Architecture

Previously, an Inception Net-based architecture has performed well on hematoma segmentation from magnetic resonance images;²⁰ as such, we utilize a similar 2D architecture with arbitrary-sized inputs, permitting 2D patch-wise training and full slice automatic segmentation. This architecture is illustrated in Figure 2. Training continued to convergence, defined as no loss improvement of 1×10^{-4} in 10 epochs on the validation patch set. The learning rate was set at 1×10^{-4} with the Adam²¹ optimizer and the continuous Dice coefficient²² (cDC) as the loss. Resultant binary segmentation masks were generated by thresholding the probability masks at 0.5.

2.3 Framework Implementation

To implement multi-site learning using cyclic weight transfer, we established a server which both the NIH and VUMC could securely access. On this server we mounted a single directory where the neural network weights were kept. Identical Python scripts at both

institutions allowed the model to be loaded, trained, and saved via secure shell access to this tertiary server without opening up public connections to either institution's data.²³

Particularly, in our implementation, data at each site is never accessible to investigators outside that institution.

2.4 Training Strategies

Single-Site Learning—As a baseline, each of the sites NIH and VUMC performed single-site learning (SSL) to convergence with their respective datasets. Once converged, each of the NIH SSL and VUMC SSL models were evaluated on the NIH test and VUMC test sets. Concretely, NIH SSL was trained on the NIH training dataset and tested on the NIH and VUMC testing datasets and VUMC SSL was trained on the VUMC train dataset and tested on both the NIH and VUMC testing datasets.

Multi-Site Learning—Multi-site learning (MSL) involved training the same model architecture from initialization (i.e.: no transfer learning), then passing the model to the next institution for the subsequent epoch. Thus, MSL would train for one epoch on the NIH train dataset, then one epoch on VUMC train dataset, then one epoch on NIH train dataset, and so on until convergence. As with the NIH SSL and VUMC SSL models, the MSL model was evaluated over both the NIH and VUMC testing datasets.

3. RESULTS

After training, we have three distinct sets of weights for our model: NIH SSL, VUMC SSL, and MSL. Each of these was evaluated over both the NIH and VUMC testing datasets. We validated all weight sets with two quantitative metrics: the Dice coefficient and hematoma volume correlation between the automatic and manual segmentations. Further explanation of these measurements follows.

3.1 Qualitative Evaluation

The automatic segmentations of test CT slices in Fig. 3 allow for qualitative comparisons between the different training sites. As expected, the model trained at its respective location shows fewer false positives than the model trained at the other location. However, in these scenarios we see the MSL model generally contains less predicted hematoma voxels. Yellow arrows indicate false positives not only near the blood-brain barrier, but also ones that are not present in the MSL segmentations.

3.2 Quantitative Evaluation

Separate from the cDC loss function, the traditional Dice coefficient was employed to judge the accuracy of the automatic masks. Figure 3 displays example segmentation results from four different patients while Table 2 shows the overall averages for all models.

To compare the efficacy of the MSL model against the two SSL models, we used the Wilcoxon signed-rank test over the corresponding Dice scores. Our findings, illustrated in Figures 4 and 5, show significant improvement between the MSL model and both the NIH SSL (p=0.009) and VUMC SSL (p=0.005) models with respect to the NIH test dataset, and a

Remedios et al.

significant improvement between the MSL model and the NIH SSL (p=0.01) over the VUMC test dataset. The VUMC SSL model outperformed the MSL model on the VUMC test data, but not significantly (p=0.337).

Two considerations are made regarding low Dice scores. First, specifically regarding the disparity of average Dice scores between the NIH and VUMC visible in Figures 4 and 5, data from the NIH had a lower average hematoma volume than VUMC data (13,700 mm³ for NIH data versus 41,000 mm³ for VUMC data), and Dice coefficients between two segmentations are known to be dependent on the volumes of the objects being considered. Second, regarding overall average Dice scores for both institutions, some 2D image slices near the top and bottom of the brain as well as along the blood-brain barrier suffer from increased false positives. These are shown in Figure 3, marked by yellow arrows.

As an alternate means to evaluate the accuracy of the automatic segmentation, we calculated the Pearson correlation coefficient between the total volume in the segmentation and manual masks, as provided in Table 2. Although the automatic segmentations contain some small false positives which reduce the Dice coefficient, overall, the volume correlations remain high.

4. DISCUSSION

To our best knowledge, this is the first application of multi-site distributed learning applied to clinical imaging data from different institutions. In this paper, we have presented and validated a technique to distributively train a convolutional neural network over disparate data housed at different institutions. While the multi-site model outperformed its single-site counterparts, our main contribution is a general framework to allow a neural model to train over more data than it would normally have access to while still preserving PHI. We show that for this task, multi-site learning did not detract from the network's ability to learn over tasks, and as expected, performance improved with more data availability. Additionally, our implementation to transfer the weights between sites automatically is straightforward, publicly available and can be generally applied to other epoch-based training scenarios. Future work includes exploring alternate neural architectures such as U-net and evaluating the generalizability of the MSL model compared with the SSL models using more than two sites.

ACKNOWLEDGEMENTS

Support for this work included funding from the Intramural Research Program of the NIH Clinical Center and the Department of Defense in the Center for Neuroscience and Regenerative Medicine, and NIH grants 1R01EB017230-01A1 (Landman) and 1R01GM120484-01A1 (Patel), as well as NSF 1452485 (Landman). The VUMC dataset was obtained from ImageVU, a research resource supported by the VICTR C-TSA award (ULTR000445 from NCATS/NIH), Vanderbilt University Medical Center institutional funding and Patient-Centered Outcomes Research Institute (PCORI; contract CDRN-1306-04869). This work received support from the Advanced Computing Center for Research and Education (ACCRE) at the Vanderbilt University, Nashville, TN, as well as in part by ViSE/VICTR VR3029. We also extend gratitude to NVIDIA for their support by means of the NVIDIA hardware grant.

REFERENCES

- Akkus Z, Kostandy PM, Philbrick KA, and Erickson BJ, "Extraction of brain tissue from CT head images using fully convolutional neural networks," in [*Medical Imaging 2018: Image Processing*], 10574, 1057420, International Society for Optics and Photonics (2018).
- [2]. Pereira S, Pinto A, Alves V, and Silva CA, "Brain tumor segmentation using convolutional neural networks in MRI images," IEEE transactions on medical imaging 35(5), 1240–1251 (2016).
 [PubMed: 26960222]
- [3]. Liskowski P and Krawiec K, "Segmenting retinal blood vessels with deep neural networks," IEEE transactions on medical imaging 35(11), 2369–2380 (2016). [PubMed: 27046869]
- [4]. de Brebisson A and Montana G, "Deep neural networks for anatomical brain segmentation," in [Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops], 20–28 (2015).
- [5]. Kanmitsas K, Chen L, Ledig C, Rueckert D, and Glocker B, "Multi-scale 3d convolutional neural networks for lesion segmentation in brain MRI," Ischemic stroke lesion segmentation 13, 46 (2015).
- [6]. Lin M, Chen Q, and Yan S, "Network in network," arXiv preprint arXiv:1312.4400 (2013).
- [7]. Halevy A, Norvig P, and Pereira F, "The unreasonable effectiveness of data," IEEE Intelligent Systems 24(2), 8–12 (2009).
- [8]. "Nih data sharing policy and implementation guidance." https://grants.nih.gov/grants/policy/ data_sharing/data_sharing_guidance.htm (2003). Accessed: 2018-07-27.
- [9]. Luxton DD, Kayl RA, and Mishkind MC, "mhealth data security: The need for hipaa-compliant standardization," Telemedicine and e-Health. 18(4), 284–288 (2012). [PubMed: 22400974]
- [10]. Thompson LA, Black E, Duff WP, Black NP, Saliba H, and Dawson K, "Protected health information on social networking sites: ethical and legal considerations," Journal of medical Internet research 13(1) (2011).
- [11]. Fetzer DT and West OC, "The hipaa privacy rule and protected health information: implications in research involving dicom image databases," Academic radiology 15(3), 390–395 (2008).
 [PubMed: 18280936]
- [12]. Kone ný J, McMahan HB, Yu FX, Richtarik P, Suresh AT, and Bacon D, "Federated learning: Strategies for improving communication efficiency," arXiv preprint arXiv:1610.05492 (2016).
- [13]. Pan SJ, Yang Q, et al., "A survey on transfer learning," IEEE Transactions on knowledge and data engineering 22(10), 1345–1359 (2010).
- [14]. Zhang S, Zhang C, You Z, Zheng R, and Xu B, "Asynchronous stochastic gradient descent for dnn training," in [Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on], 6660–6663, IEEE (2013).
- [15]. Chang K, Balachandar N, Lam C, Yi D, Brown J, Beers A, Rosen B, Rubin DL, and Kalpathy-Cramer J, "Distributed deep learning networks among institutions for medical imaging," Journal of the American Medical Informatics Association (2018).
- [16]. Trifan G, Gattu R, Haaeke EM, Kou Z, and Benson RR, "Mr imaging findings in mild traumatic brain injury with persistent neurological impairment," Magnetic resonance imaging 37, 243–251 (2017). [PubMed: 27939436]
- [17]. Kinnunen KM, Greenwood R, Powell JH, Leech R, Hawkins PC, Bonnelle V, Patel MC, Counsell SJ, and Sharp DJ, "White matter damage and cognitive impairment after traumatic brain injury," Brain 134(2), 449–463 (2010). [PubMed: 21193486]
- [18]. Muschelli J, Ullman NL, Mould WA, Vespa P, Hanley DF, and Crainiceanu CM, "Validated automatic brain extraction of head CT images," Neuroimage 114, 379–385 (2015). [PubMed: 25862260]
- [19]. Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, and Rabinovich A, "Going deeper with convolutions," in [Proceedings of the IEEE conference on computer vision and pattern recognition], 1–9 (2015).
- [20]. Roy S, Butman JA, Chan L, and Pham DL, "Tbi contusion segmentation from MRI using convolutional neural networks," in [[Biomedical Imaging (ISBI 2018), 2018 IEEE 15th International Symposium on], 158–162, IEEE (2018).

- [21]. Kingma DP and Ba J, "Adam: A method for stochastic optimization," arXiv preprint arXiv: 1412.6980 (2014).
- [22]. Shamir RR, Duchin Y, Kim J, Sapiro G, and Harel N, "Continuous dice coefficient: a method for evaluating probabilistic segmentations," bioRxiv , 306977 (2018).
- [23]. "Multi-site TBI CT lesion segmentation implementation." https://github.com/MASILab/ tbi_ct_lesion_segmentation (2018).

Remedios et al.



Figure 1:

Representative 5.0 mm thick transverse CT sections through the head in 10 subjects with TBI. In-plane resolution is approximately 0.5×0.5 mm. In each case, the hemorrhagic lesion appears intermediate density between normal brain tissue and bone. Note the heterogeneity of size, location, density and configuration.

Remedios et al.



Figure 2:

Illustrated is the proposed model architecture. Convolution layers are indicated in yellow, with notation $N @ k^2$ representing N 2D kernels of size $k \times k$. The activation for all convolution layers is **ReLU**, except for the final 1 @ 1² convolution which uses a sigmoid activation. **Avg** @ k^2 and **Max** @ k^2 respectively correspond to average pooling and max pooling with strides $k \times k$. The modified Inception Module, shown to the right, is a variation of the original Inception Module presented by Google.¹⁹

Remedios et al.



Figure 3:

Examples of automatic segmentations are shown. From left to right, the images correspond to the CT, the manual "ground truth" (GT) segmentation, the NIF1 SSL, VUMC SSL, and MSL segmentations. Both the image volume and the specific image slice's Dice coefficient are overlaid on that segmentation. Yellow arrows specify examples of false positives near the blood-brain barrier which were not present in the MSL segmentations.

Remedios et al.



Figure 4:

Model Dice coefficient comparison over the NIH testing dataset. The MSL model performed significantly better than either SSL models, where the asterisk indicates a significant difference (p < 0.05, found via the Wilcoxon signed-rank test).

Remedios et al.



Figure 5:

Model Dice coefficient comparison over the VUMC testing dataset. The MSL model performed significantly better than the NIH SSL model, but performed worse than the VUMC SSL model. The asterisk indicates significance (p < 0.05, according to the Wilcoxon signed-rank test) where "ns" corresponds to "not significant" (p >= 0.05).

Table 1:

Distribution of CT image volumes between training and test sets for both sites.

Training Location	# Training Images	# Testing Images	
VUMC	10	8	
NIH	17	10	
Total	27	18	

_

Table 2:

Average Dice coefficients and Pearson correlation coefficients for the three training strategies over the NIH and VUMC datasets. The average result over both datasets is shown to illustrate each model's general ability. An asterisk indicates significant improvements in Dice coefficient (p < 0.05) between the MSL and each of the NIH SSL and VUMC SSL models as evaluated by the Wilcoxon signed-rank test, and bold text indicates the highest Pearson correlation coefficient between automatic and manual segmented hematoma volumes.

	NIH Data		VUMC Data		Average of NIH and VUMC data	
	Dice	Correlation	Dice	Correlation	Dice	Correlation
Inter-Rater	0.687	n/a	n/a	n/a	n/a	n/a
NIH SSL	0.512	0.913	0.690	0.752	0.601	0.832
VUMC SSL	0.407	0.859	0.745	0.754	0.576	0.807
MSL	0.552*	0.943	0.725	0.791	0.63*	0.867