Journal of Digital Imaging (2023) 36:964–972
https://doi.org/10.1007/s10278-023-00788-y

A Deep Learning Approach for Automated Bone Removal from Computed Tomography Angiography of the Brain

Masis Isikbay[1] · M. Travis Caton[2] · Evan Calabrese[1,3,4,5]

Received: 21 November 2022 / Revised: 29 January 2023 / Accepted: 30 January 2023 / Published online: 13 February 2023

Abstract

Advanced visualization techniques such as maximum intensity projection (MIP) and volume rendering (VR) are useful for evaluating neurovascular anatomy on CT angiography (CTA) of the brain; however, interference from surrounding osseous anatomy is common. Existing methods for removing bone from CTA images are limited in scope and/or accuracy, particularly at the skull base. We present a new brain CTA bone removal tool, which addresses many of these limitations. A deep convolutional neural network was designed and trained for bone removal using 72 brain CTAs. The model was tested on 15 CTAs from the same data source and 17 CTAs from an independent external dataset. Bone removal accuracy was assessed quantitatively, by comparing automated segmentation results to manual segmentations, and qualitatively by evaluating VR visualization of the carotid siphons compared to an existing method for automated bone removal. Average Dice overlap between automated and manual segmentations from the internal and external test datasets were 0.986 and 0.979 respectively. This was superior compared to a publicly available noncontrast head CT bone removal algorithm which had a Dice overlap of 0.947 (internal dataset) and 0.938 (external dataset). Our algorithm yielded better VR visualization of the carotid siphons than the publicly available bone removal tool in 14 out of 15 CTAs (93%, chi-square statistic of 22.5, p-value of <0.00001) from the internal test dataset and 15 out of 17 CTAs (88%, chi-square statistic of 23.1, p-value of <0.00001) from the external test dataset. Bone removal allowed subjectively superior MIP and VR visualization of vascular anatomy/pathology. The proposed brain CTA bone removal algorithm is rapid, accurate, and allows superior visualization of vascular anatomy and pathology compared to other available techniques and was validated on an independent external dataset.

Keywords Machine learning · Deep learning · Artificial Intelligence · CT angiography · Brain · Neurovascular

Abbreviations

3D	Three dimensional
CTA	Computed tomography angiography
DICOM	Digital Imaging and Communications in Medicine
DSA	Digital subtraction angiography
MIP	Maximum intensity projection
MRI	Magnetic resonance imaging
VR	Volume rendering

✉ Masis Isikbay
masis.isikbay@ucsf.edu

[1] Department of Radiology and Biomedical Imaging, University of California San Francisco, 505 Parnassus Ave, M-396, San Francisco, CA 94143, USA

[2] Cerebrovascular Center, Department of Neurosurgery, Icahn School of Medicine at Mount Sinai, 1450 Madison Ave, New York, NY 10029, USA

[3] Department of Radiology, Division of Neuroradiology, Duke University Medical Center, Box 3808 DUMC, Durham, NC 27710, USA

[4] Duke Center for Artificial Intelligence in Radiology (DAIR), Duke University Medical Center, Durham, NC 27710, USA

[5] Center for Intelligent Imaging, University of California San Francisco, San Francisco, CA 94143, USA

Background and Purpose

Medical image post-processing techniques are a powerful tool in the radiologist's arsenal. In particular, maximum intensity projections (MIPs), volume rendering (VR), and cinematic rendering can be useful for evaluation of vasculature on cross-sectional angiographic imaging [1–5]. These techniques facilitate visualization of key pathology (such as intracranial aneurysms) and can be beneficial for pre-procedural planning (for example prior to endovascular aneurysm coiling) [6]. Despite this utility, one

common obstacle for post-processing of CT angiography (CTA) images of the brain is interference from surrounding osseous anatomy, particularly at the skull base [7–9]. The complexity of skull base anatomy its similar density to contrast opacified vessels/arterial calcification often precludes detailed vessel evaluation in post-processed images. These limitations of existing post-processing techniques are most pronounced at vessel-bone interfaces (e.g., clinoidal and ophthalmic segments of the internal carotid artery, dural venous sinuses) and can result in missed or delayed diagnosis of potentially life-threatening vascular pathologies [10].

There are several existing methods for removing osseous structures from CTA exams, each with their respective benefits and limitations. These methods can be broadly categorized as acquisition-based methods (i.e., using specific CTA acquisition strategies to isolate bone signal) and post-hoc methods (i.e., image processing methods using previously acquired standard CTA data). Acquisition-based methods include digital subtraction CTA [11], which first acquires a non-contrast CT and subtracts it from the subsequently acquired CTA, as well as photon energy-based techniques (i.e. dual-energy and photon counting CT) [12–14], which rely on spectral attenuation differences to isolate bone (calcium) signal. Acquisition-based methods can be very accurate but require specialized equipment. They often relying on proprietary acquisition/processing software, and some methods result in increased radiation dose to the patient. Post-hoc methods include manual (human) segmentation, and various automated or semi-automated approaches. These include intensity threshold-based segmentation methods [15–17] and, more recently, deep learning-based methods [18]. These techniques can be applied retrospectively to any standard CTA data and do not require any specialized equipment. Manual segmentation is the most straightforward approach but is time consuming and requires dedicated personnel with considerable anatomic expertise. Threshold-based automated approaches, such as the free tool provided by the DICOM viewer Horos [15], produce rapid results but are hindered by similarities in density between contrast opacified vessels and adjacent bones. Deep learning-based methods have the potential to address many of the issues facing other bone removal methods; however, limited prior studies have raised concerns about the ability to accurately segment both arterial and venous structures, especially in the setting of vascular pathology or prior intracranial inetervention [18].

To address the limitations of existing techniques we developed and evaluated a new deep-learning approach for automated removal of osseous structures from CTA exams using free, open source, and multi-platform software. This technique was designed to be rapid, accurate, fully automated, and applicable to patients with or without neurovascular pathology and/or prior intracranial intervention. Our technique does not require the use of any special scanning equipment or parameters, and the generalizability of our approach was tested on an independent, external dataset. In addition, the proposed algorithm is free and open source, which allows users to independently validate results.

Table 1 Patient cohort details for internal CTA dataset

Mean age (SD)	Patient sex	Intracranial pathology	Prior surgery
63 (7.3)	Male 57 (55%) Female: 43 (45%)	Yes: 51 (50%) No: 52: (50%)	Yes: 41 (40%) No: 62: (60%)

Methods

Study Cohort

We retrospectively identified an "internal cohort" of 103 adult patients (cohort details outlined in Table 1) who underwent CTA of the brain from 2021 to 2022 at a single medical center using enterprise radiology report search tools (mPower, Nuance, Burlington, MA) [19]. Patient data was de-identified at the time of download using the Automated Image Retrieval Portal [20] integrated into our PACS.

The cohort was manually selected to include a number of patients without significant pathology (approximately 50% of the cases) as well as patients with a variety of neurovascular abnormalities including severe intracranial vascular calcifications (15%), significant senescent white matter changes (25%),intracranial hemorrhage (approximately 10% of cases), prior aneurysm clipping with craniotomy (approximately 15% of cases), and some with a combination of post-surgical changes such as extra axial blood products, pneumocephalus, surgical hardware, craniotomy, and/or ventricular catheters (approximately 25% of cases). Patients ranging in age from 25–94 years old were included in our dataset.

Data Acquisition and Preprocessing

All CTA data in the internal cohort was acquired at a single institution on a variety of General Electric CT scanners (Eagle Revolution X, Revolution XRD, Light Speed 16 slice, VCT Light Speed 64 slice, and Discovery 750 HD 64 slice) using a standard helical acquisition with a reconstruction slice matrix of 512 × 512 (variable in plane resolution based on field of view) and slice thickness of 0.625 mm. Minimal automated data preprocessing steps included resampling to 0.5-mm isotropic resolution and voxel intensity scaling such that − 1000 Houndsfield units (HU) corresponded to an image intensity value of − 1.0 and 1000 HU corresponded to an image intensity value of 1.0.

Manual Segmentation of CTA Data

Initial brain segmentations for 20 CTA exams were achieved using simple thresholding at 300 Hounsfield units (HU) followed by binary morphologic erosion, isolation of the largest 26-connected three-dimensional component, and morphologic dilation using a 3-voxel spherical kernel. Segmentations were manually corrected by senior radiology trainees and approved by a fellowship-trained attending neuroradiologist utilizing the ITK-SNAP software [21]. Segmentation included the entire intracranial and visible upper cervical intrathecal contents (regardless of density), all major intracranial blood vessels (inclusive of any vascular calcifications/age related changes), and any intracranial surgical material (e.g. shunts, aneurysm clips). Segmentation was truncated in the plane of skull base foramina (except foramen magnum) and in the plane of any existing craniotomy defects. The entire intracranial internal carotid artery (including the petrous segment) was included in the segmentation to the level of the carotid foramen. The first 20 manually corrected exams were used to train a preliminary neural network as described in the next section. This preliminary network was then used to generate initial segmentations for the remaining internal and external cases, which were subsequently manually corrected and reviewed by an attending neuroradiologist as previously described.

Automated Segmentation Network Architecture

A 3D deep convolutional neural network (dCNN) based on the U-net architecture and adapted from a previously published method for brain extraction from 3D MRI data was implemented in Python 3.8 and Tensorflow 2.9.1 [22, 23]. The encoder limb of the U-net consisted of 3 sets (3, 3, and 4, respectively) of bottleneck residual convolution blocks with $3\times3\times3$ voxel kernels, which were separated by 2-strided convolution down-sampling blocks. The decoder limb was the reverse separated by 2-strided convolution transpose upsampling blocks. Long range concatenation skip connections were included before each down-sampling/up-sampling step. A base of 32 output filters was chosen for convolution layers, which was doubled with each down-sampling step and halved with each upsampling step. The output layer consisted of a 3D convolutional layer with a single output filter (logit), which was transformed into a probability using the sigmoid logistic function. Batch normalization and 40% feature dropout was implemented at each bottleneck residual convolution block. The final model had a total of 799,937 trainable parameters.

Network Training and Validation

The internal cohort dataset was randomly divided according to a 70%/15%/15% training/validation/testing split. Training comprised 20 epochs iterating through the entire training dataset (72 CTAs). Training inputs consisted of 80×80x80 voxel overlapping 3D patches from the training data. Random dimension swaps and 3-axis rotations of ±30 degrees were employed for data augmentation. Model weights were initialized using the Glorot and Bengio [24] initializer at the beginning of training. The training loss function was the product of the voxel wise Dice coefficient and voxel wise binary cross entropy and was minimized using the Adam optimizer. The learning rate was set to 0.001 for the first 12 epochs, and then halved prior to each subsequent epoch. Training was accomplished using two Nvidia A6000 GPUs with a mirrored distribution strategy and a batch size of 12 samples per GPU (24 samples per training step). Training validation was accomplished using the 15% validation split after each training epoch. The model weights yielding the lowest validation loss (epoch 20/20) were used for subsequent testing.

Network Testing

The trained model was evaluated on the held out 15% test split from the internal dataset. Quantitative model performance was evaluated using the Dice coefficient between the model generated segmentations and manual segmentations. The Dice of our algorithm was then compared against a publicly available non-contrast head CT bone removal algorithm [25] utilizing the same patient datasets.

External Testing

An independent external testing dataset consisting of 17 CTAs was selected from the UCLH Stroke EIT Dataset [26] This dataset consisted of a total of 22 patients who had stroke related imaging, however only 17 of the patients had CTA imaging of the head as part of their workup. Scanner make and model and details of acquisition technique were not available for this dataset. Data preprocessing methods were identical to the internal dataset. All external dataset manual segmentations were reviewed for accuracy and approved by both a fellowship-trained neuroradiology attending and a fellowship-trained neurointerventional radiology attending. The trained segmentation model was evaluated on the external dataset using identical methods to the internal testing.

Data Visualization and Volume Rendering

Automatically segmented data from the internal and external test datasets were visualized with MIP and VR techniques using publicly available software (Horos[15]) which could be used to all analyze internal and external data regardless of acquisition technique. Qualitative segmentation accuracy was assessed by determining which bone removal method yielded subjectively superior three-dimensional VR visualization of the carotid siphons (with

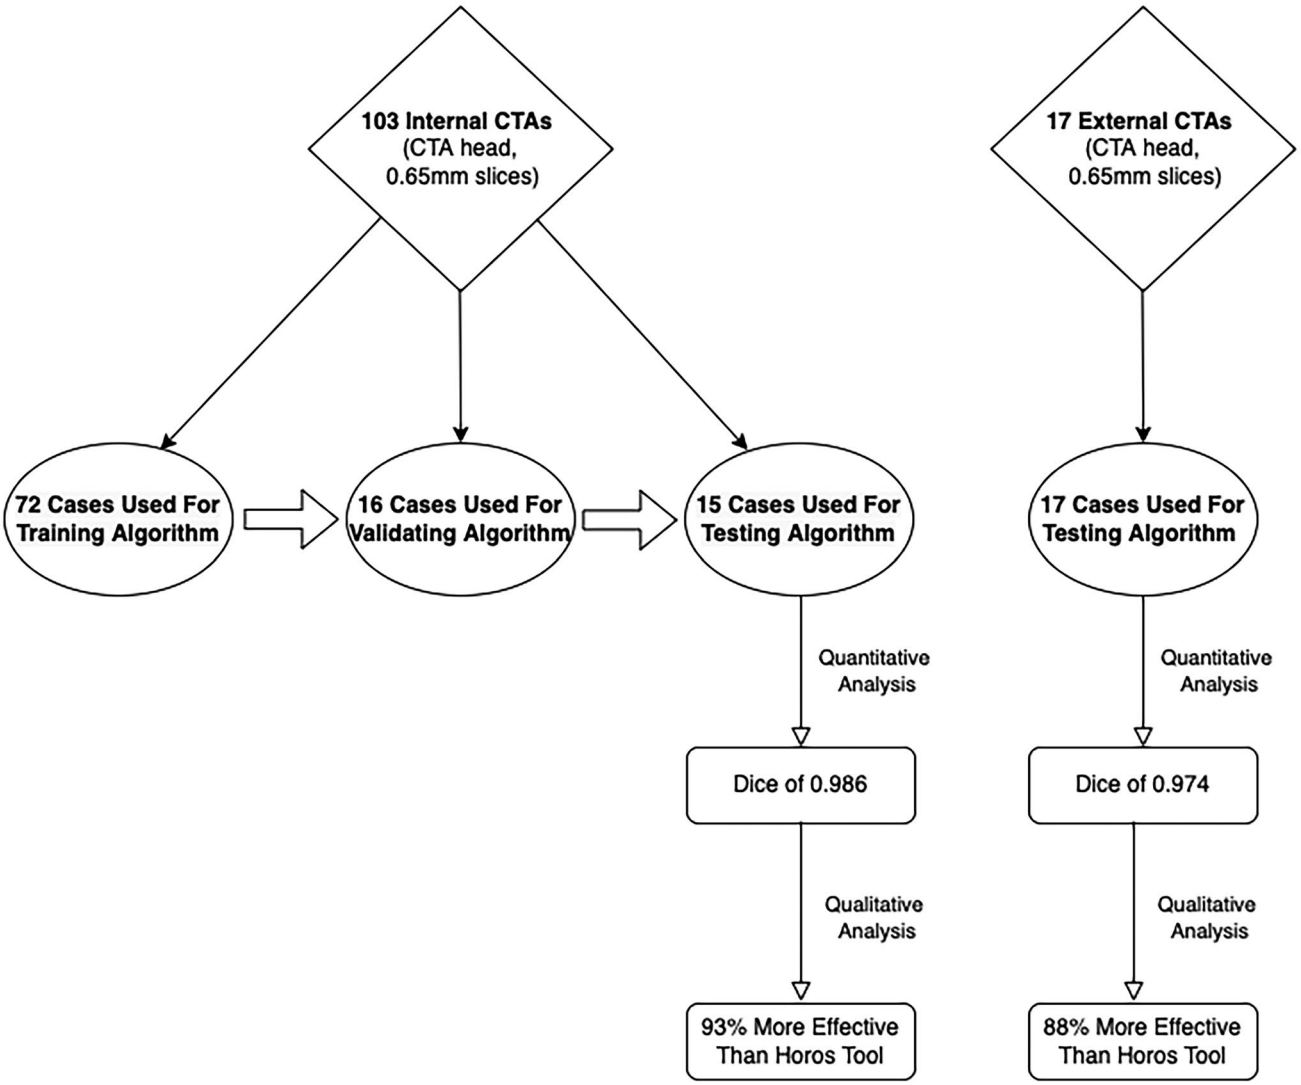

Fig. 1 Summary of study design including the number of patients included in the internal training, validation, and both internal/external test sets. Segmentation accuracy was compared quantitatively to manual segmentations (Dice coefficient) and qualitatively to an existing publicly available automated bone removal algorithm

an emphasis on identifying cases which were missing significant portions of one or both carotid arteries after bone removal). Upon identifying these cases, chi-square analysis was performed comparing the frequency of successful carotid siphon segmentation (those cases with both carotids segmented without gaps) between the Horos bone removal tool and our algorithm.

Results

Study Design

A diagrammatic flow chart of the study design is presented as Fig. 1.

Network Training

Training comprised 20 epochs, each with > 300,000 individual training examples and lasting approximately 135 mi for a total training time of approximately 45 h. A graphical schematic of the model is provided as Supplementary Fig. 1. The loss function decremented appropriately throughout training with a total loss reduction of ~ 130 × between the first and last epoch (Supplementary Fig. 2). Validation loss also decremented appropriately throughout training and reached its lowest value at the final (20th) epoch; however, some validation loss oscillation was noted in the last 10 epochs indicating that additional training epochs would be unlikely to provided substantial improvement in model performance [27, 28].

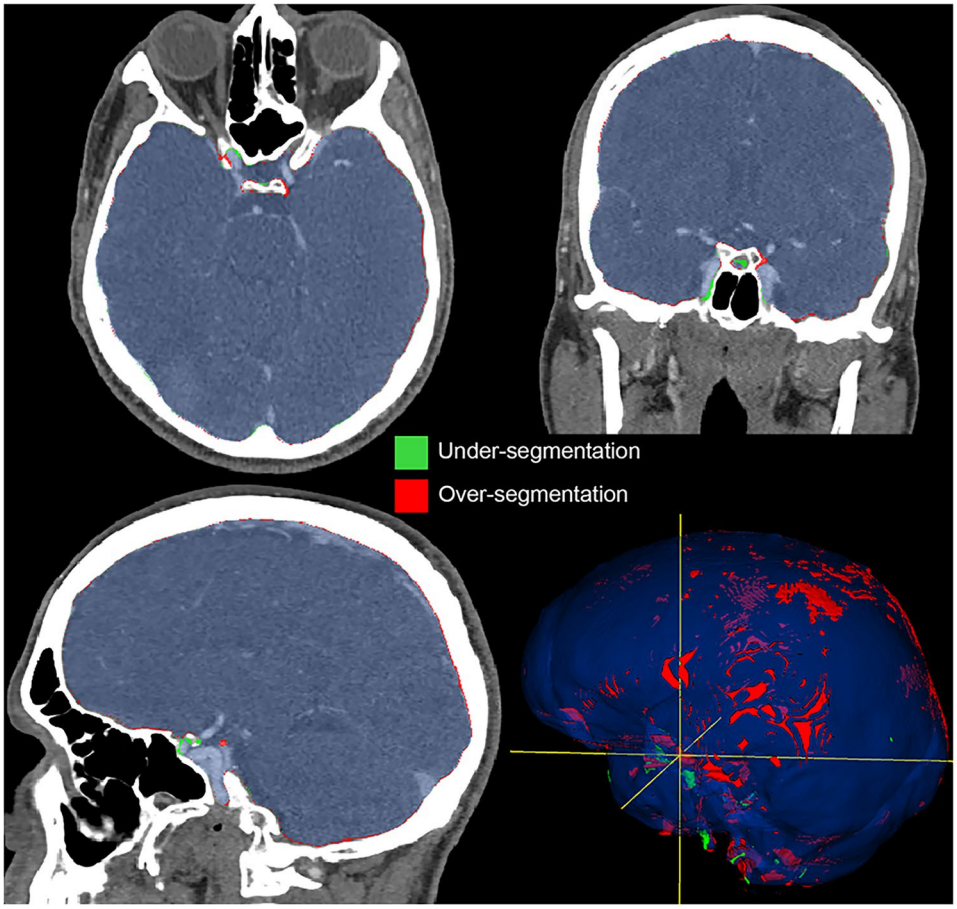

Fig. 2 Regions of over- and under-segmentation for a representative automated CTA brain segmentation from the external test dataset. Regions of over-segmentation are shown as red pixels while regions of under-segmentation are shown in green. Accurately segmented regions are shown in blue

Quantitative Assessment of Segmentation Results

After training, the average Dice overlap between automated and manually corrected segmentations was 0.986 for the internal test dataset, and 0.979 for the external test dataset. Figure 2 shows areas of over- and under-segmentation for a representative CTA brain image from the external test dataset. The corresponding Dice overlap utilizing another publicly available bone removal algorithm was 0.947 for the internal test dataset, and 0.938 for the external test dataset. Example segmentation masks (for representative internal and external dataset cases) comparing algorithms are shown in Supplementary Fig. 3.

Qualitative Assessment of Segmentation Results

Qualitative analysis with the publicly available automated bone removal tool from Horos revealed that both carotid siphons were segmented without any gaps in 0/15 internal cases and 0/17 external cases. This was compared to our developed algorithm which segmented both carotid siphons without gaps in 14/15 (93%) of internal cases (chi-square statistic of 22.5, p-value of < 0.00001) and 15/17 CTAs (88%) in the external dataset (chi-square statistic of 23.1, p-value of < 0.00001). A representative visual (VR) comparison of the two bone removal methods is shown in Fig. 3.

Visualization of Neurovascular Anatomy and Pathology

MIP reformats were generated to allow visualization of key vascular anatomy such as the basilar artery (Fig. 4A) to demonstrate that bone removal allows for subjectively improved visualization of normal anatomy (Fig. 4B). Similarly, bone removal subjectively improved visualization of skull base vascular pathology such as an ophthalmic segment carotid artery aneurysm (Fig. 4C, D).

Data and Code Availability

The entire codebase for this project is open source and publicly available at https://github.com/ecalabr/brain_mask (specific commit ID: ee86a9678d07508c2f990f5d878350b3dae8ac51). The final trained model weights are also provided. External image data is available at https://doi.org/10.5281/zenodo.11993 9826. Internal cohort image data is not publicly available.

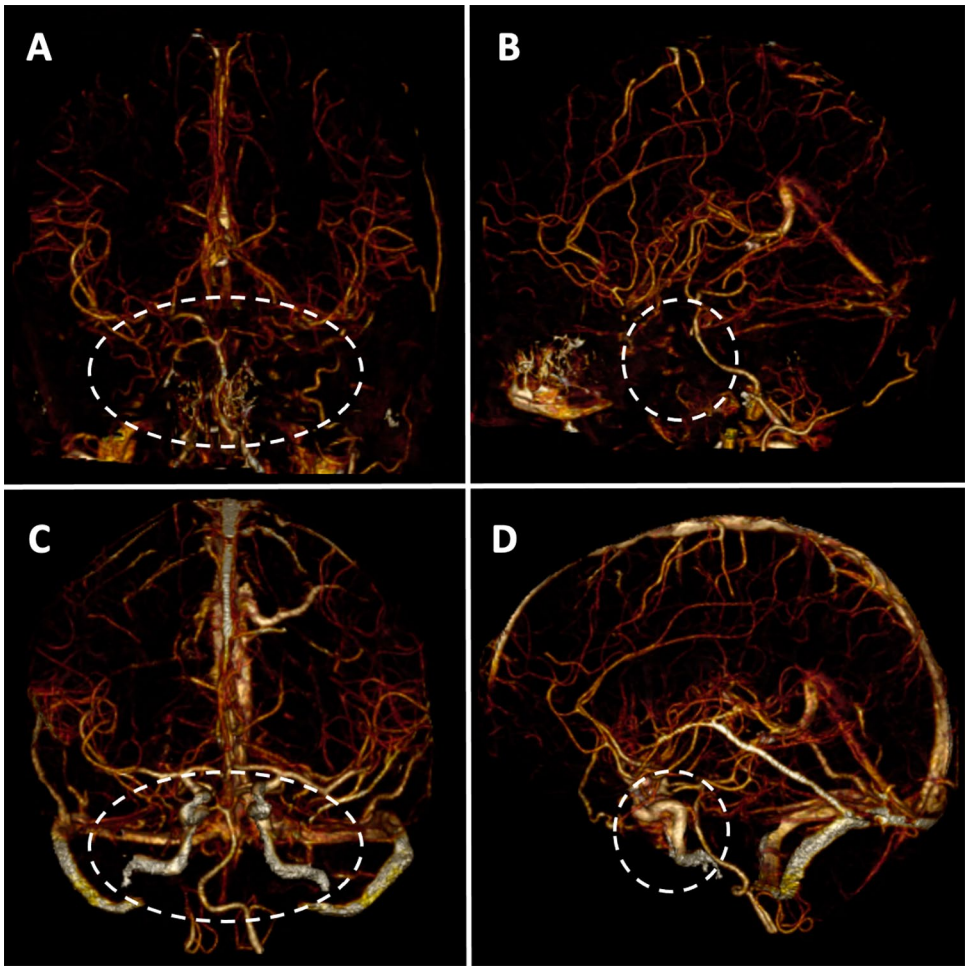

Fig. 3 Frontal **A**, **C** and lateral **B**, **D** three-dimensional volume rendered visualization with surface shaded display of the same CTA brain data after bone removal with a publicly available threshold-based bone removal method **A**, **B** and the proposed deep learning method **C**, **D**. Visualizations of both methods are shown with the same window/level and lighting parameters. Notably both arterial and venous anatomy is more clearly visualized with the deep learning method

Discussion

Here we report a new deep learning-based technique for bone removal from standard CTA images of the brain, which addresses many of the issues that limit previously available methods. The proposed technique is generalizable in its applicability, as it was able to remove bone accurately and reliably from standard CTA brain data regardless of scanner type. This included cases from an external dataset where scanner model and technique was neither controlled nor available at the time of processing. We further demonstrate the potential utility of this method by using it to visualize the internal carotid arteries at the skull base, a location which has proved challenging for many existing bone removal techniques. The proposed bone removal algorithm yielded superior VR visualization of skull base vascular anatomy compared to other currently available techniques, and outperformed other previously created bone removal algorithms.

The advantages of our developed tool are that it can be more broadly applied than acquisition-based techniques that require special equipment (digital subtraction CTA, dual energy/photon counting CT), it is more accurate than currently available threshold-based approaches, and it is significantly faster than manual segmentation. In contrast to prior deep learning-based approaches, the proposed technique was trained and evaluated on patients with and without neurovascular pathology and/or prior intracranial intervention, and was designed to accurately segment both venous and arterial anatomy [18].

The broad applicability of the proposed technique opens the door for a variety of neurovascular visualization applications. For example, automated bone removal can simplify generation of MIP or VR visualizations of the circle of Willis, which are often created by CT technologists and/or radiologists for standard diagnostic CTAs of the brain. In addition, vascular pathology that closely opposes osseous structures (such as the relationship of the internal carotid arteries to the clivus) can be more clearly visualized after bone removal. This is demonstrated by the superior visualization of an ophthalmic segment internal carotid artery aneurysm following bone removal as shown in Fig. 4C and D. While we did not assess diagnostic accuracy in this study, it is possible that these improvements in visualization of neurovascular anatomy may improve detection and

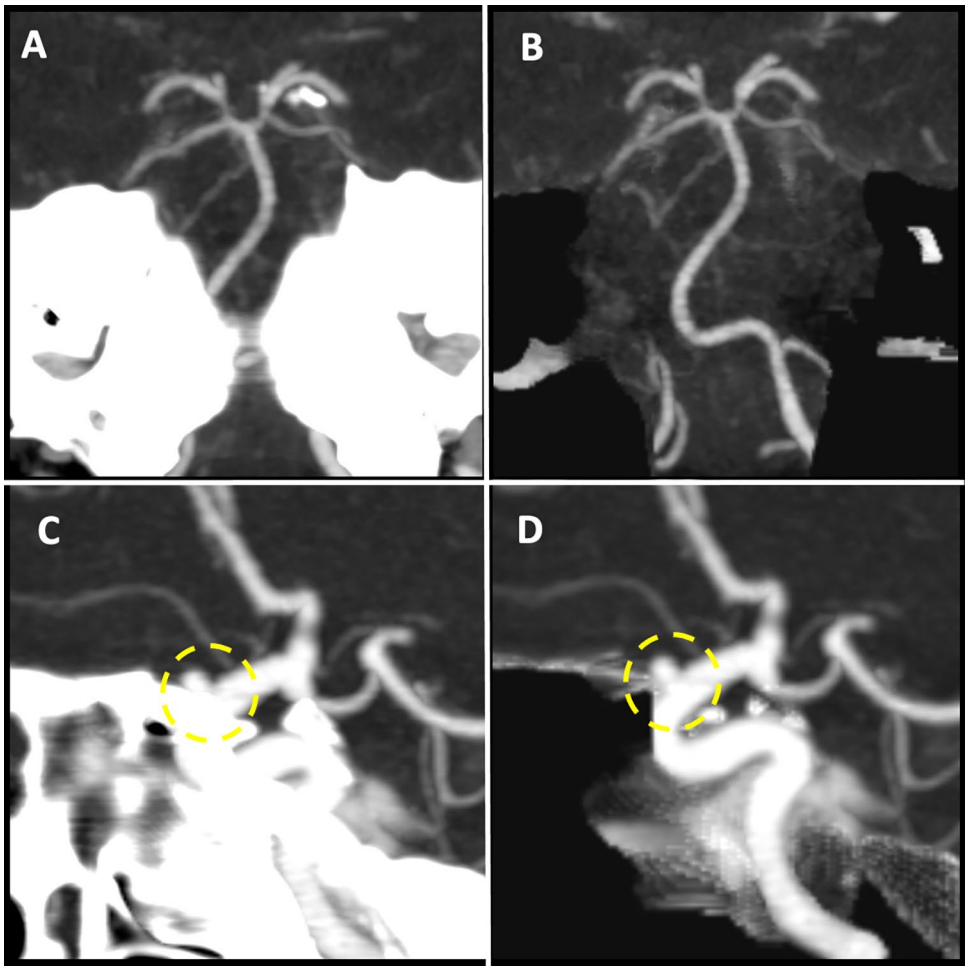

Fig. 4 Maximum intensity projection (MIP) images of the basilar artery before **A** and after bone removal **B** demonstrates subjectively superior visualization of normal anatomy with bone removal. Similarly, MIP images of the internal carotid artery before **C** and after bone removal **D** demonstrates subjectively superior visualization of an ophthalmic segment aneurysm after bone removal (yellow dashed circles)

accurate characterization of key vascular pathology. This is especially relevant given the various factors that contribute to diagnostic error [29], commonly reported "blind spots" on head/neck imaging [30], and that 33% of missed cerebral aneurysms involve the internal carotid artery [31]. Interestingly, other bone removal techniques have been reported to contribute to the misdiagnosis of certain aneurysms [32], which can have a significant impact on patient outcomes [10]. For this reason, thorough clinical assessment of this, or any similar postprocessing technique, will be required prior to clinical implementation.

More creative applications of this technique are also now a possibility. Given the ease by which 3D reformats can be manipulated after bone removal, processed CTA data could be used for endovascular neurosurgical planning such as determining the optimal projections for visualizing the pathology of interest (for example an aneurysm to be coiled) prior to the intervention. The deep learning method described here may also be applicable to other similar neuroimaging tasks given that bone removal will not impact standard 3D processing techniques of the source data such as surface shaded display (SSD) and the application of Laplacian filters. The network architecture that this work is based on (U-net) has proven to be extremely adaptable for a variety of semantic segmentation tasks, and we have taken steps to optimize this architecture for standard CT data. Further, we have made the codebase for this project freely available, including the network training code, to allow this algorithm to be retrained for other tasks. Possible future applications include extending this algorithm for bone removal of CTA neck and adapting for bone removal of non-contrast head CT to aid in the detection of extra axial hemorrhage.

This study has several important limitations, which should be carefully considered when interpreting the presented results. First, this algorithm was only evaluated on CTA brain data from the skull vertex to the skull base, and it was neither designed nor evaluated for segmentation below this level. In addition, our algorithm was trained on CTA images reconstructed at 0.65 mm slice thickness (before resampling to 0.5 mm isotropic). While this type of data is now routinely acquired and can be reconstructed from most current generation CT scanners, older scans acquired with thicker slices (such as 2–5-mm slices) are unlikely to

perform as well with our algorithm. Another limitation is the relatively small training set of < 100 patients used to build the model. While we specifically curated the internal cohort to include a variety of neurovascular pathologies and interventions, there are many others (e.g., ECA to MCA bypass) that were not included and could lead to unexpected results during inference. This could be solved by training the algorithm on a much larger more varied cohort, which will be an important area of future work. Finally, we employed a "human in the loop" semiautomated AI approach for generating gold-standard segmentations [33], which has some degree of inherent bias and could artificially boost quantitative evaluation results. While this bias is very likely to be present, it is also likely to be small and should be further mitigated by the fact that segmentations were manually corrected and reviewed by attending radiologists prior to use.

Conclusion

The proposed technique is a rapid and accurate method for automated removal of the skull from standard CTA of the brain. It outperforms existing publicly available methods for VR visualization of skull base vasculature, and closely mimics expert reviewed manual segmentations. This technique facilitates post-processing and three-dimensional visualization of both arterial and venous neurovascular anatomy in normal, diseased, and postoperative states, and may ultimately be helpful for increasing efficiency and accuracy for diagnosing neurovascular pathology.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s10278-023-00788-y.

Acknowledgements None.

Funding Dr. M. Isikbay was supported by a National Institutes of Health (NIBIB) T32 Training Grant, T32EB001631.

Declarations

Presentation Information/Sponsoring Societies Related work was presented at the 15th Annual Meeting of the American Society of Functional Neuroradiology (ASNFR 2022). The results of this work have been accepted to be presented at the annual meeting of the American Society of Neuroradiology (ASNR 2023).

Conflict of Interest N/A.

References

1. Caton Jr. MT, Wiggins WF, Nunez D. Three-Dimensional Cinematic Rendering to Optimize Visualization of Cerebrovascular Anatomy and Disease in CT Angiography. *Journal of Neuroimaging* 2020;30:286–96.
2. Korogi Y, Takahashi M, Katada K, et al. Intracranial aneurysms: detection with three-dimensional CT angiography with volume rendering--comparison with conventional angiographic and surgical findings. *Radiology* 1999;211:497–506.
3. White PM, Teasdale EM, Wardlaw JM, et al. Intracranial aneurysms: CT angiography and MR angiography for detection prospective blinded comparison in a large patient cohort. *Radiology* 2001;219:739–49.
4. Fishman EK, Ney DR, Heath DG, et al. Volume Rendering versus Maximum Intensity Projection in CT Angiography: What Works Best, When, and Why. *RadioGraphics* 2006;26:905–22.
5. Perandini S, Faccioli N, Zaccarella A, et al. The diagnostic contribution of CT volumetric rendering techniques in routine practice. *Indian J Radiol Imaging* 2010;20:92–7.
6. Chen W, Yang Y, Xing W, et al. Applications of multislice CT angiography in the surgical clipping and endovascular coiling of intracranial aneurysms. *J Biomed Res* 2010;24:467–73.
7. Broder J, Preston R. Chapter 1 - Imaging the Head and Brain. In: Broder J, ed. *Diagnostic Imaging for the Emergency Physician*. Saint Louis: W.B. Saunders; 2011:1–45.
8. Bello HR, Graves JA, Rohatgi S, et al. Skull Base–related Lesions at Routine Head CT from the Emergency Department: Pearls, Pitfalls, and Lessons Learned. *RadioGraphics* 2019;39:1161–82.
9. Anderson GB, Ashforth R, Steinke DE, et al. CT Angiography for the Detection and Characterization of Carotid Artery Bifurcation Disease. *Stroke* 2000;31:2168–74.
10. Mayer PL, Awad IA, Todor R, et al. Misdiagnosis of Symptomatic Cerebral Aneurysm. *Stroke* 1996;27:1558–63.
11. Lu L, Zhang LJ, Poon CS, et al. Digital subtraction CT angiography for detection of intracranial aneurysms: comparison with three-dimensional digital subtraction angiography. *Radiology* 2012;262:605–12.
12. Watanabe Y, Uotani K, Nakazawa T, et al. Dual-energy direct bone removal CT angiography for evaluation of intracranial aneurysm or stenosis: comparison with conventional digital subtraction angiography. *Eur Radiol* 2009;19:1019–24.
13. Postma AA, Das M, Stadler AAR, et al. Dual-Energy CT: What the Neuroradiologist Should Know. *Curr Radiol Rep* 2015;3:16.
14. Sommer WH, Johnson TR, Becker CR, et al. The value of dual-energy bone removal in maximum intensity projections of lower extremity computed tomography angiography. *Invest Radiol* 2009;44:285–92.
15. Nimble Co LLC. Horos Project. 2018 Feb 7. [Epub ahead of print].
16. Friedli L, Kloukos D, Kanavakis G, et al. The effect of threshold level on bone segmentation of cranial base structures from CT and CBCT images. *Sci Rep* 2020;10:7361.
17. van Straten M, Schaap M, Dijkshoorn ML, et al. Automated bone removal in CT angiography: Comparison of methods based on single energy and dual energy scans. *Medical Physics* 2011;38:6128–37.
18. Fu F, Wei J, Zhang M, et al. Rapid vessel segmentation and reconstruction of head and neck angiograms using 3D convolutional neural network. *Nat Commun* 2020;11:4829.
19. mPower Clinical Analytics for medical imaging | Nuance. *Nuance Communications*.
20. Automated Image Retrieval (AIR) - PACS. *UCSF Data Resources*.
21. ITK-SNAP. Paul A. Yushkevich, Joseph Piven, Heather Cody Hazlett, Rachel Gimpel Smith, Sean Ho, James C. Gee, and Guido Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage* 2006 Jul 1;31(3):1116–28.
22. Calabrese E, Villanueva-Meyer JE, Cha S. A fully automated artificial intelligence method for non-invasive, imaging-based identification of genetic alterations in glioblastomas. *Sci Rep* 2020;10:11852.
23. Calabrese E, Rudie JD, Rauschecker AM, et al. Feasibility of Simulated Postcontrast MRI of Glioblastomas and Lower-Grade Gliomas by Using Three-dimensional Fully Convolutional Neural Networks. *Radiology: Artificial Intelligence* 2021;3:e200276.

24. Glorot X, Bengio Y. Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings; 2010:249–56.
25. aqqush. CT_BET: Robust Brain Extraction Tool for CT Head Images. 2022 Oct 20. [Epub ahead of print].
26. Goren, Nir, Dowrick, Thomas, Avery, James, & Holder, David. UCLH Stroke EIT Dataset - Radiology Data | Zenodo. https://doi.org/10.5281/zenodo.1199398.
27. Caruana R, Lawrence S, Giles L. Overfitting in neural nets: 14th Annual Neural Information Processing Systems Conference, NIPS 2000. *Advances in Neural Information Processing Systems 13 - Proceedings of the 2000 Conference, NIPS 2000* 2001. [Epub ahead of print].
28. Prechelt L. Early Stopping — But When? In: Montavon G, Orr GB, Müller K-R, eds. *Neural Networks: Tricks of the Trade: Second Edition*. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer; 2012:53–67.
29. Itri JN, Tappouni RR, McEachern RO, et al. Fundamentals of Diagnostic Error in Imaging. *Radiographics* 2018;38:1845–65.
30. Bahrami S, Yim CM. Quality Initiatives: Blind Spots at Brain Imaging. *RadioGraphics* 2009;29:1877–96.
31. Biddle G, Assadsangabi R, Broadhead K, et al. Diagnostic Errors in Cerebrovascular Pathology: Retrospective Analysis of a Neuroradiology Database at a Large Tertiary Academic Medical Center. *AJNR Am J Neuroradiol* https://doi.org/10.3174/ajnr.A7596.
32. He L, Li Z. The aneurysm close to the skull base was wiped off by bone subtraction on 3D CTA images. *AIP Conference Proceedings* 2019;2079:020031.
33. Budd S, Robinson EC, Kainz B. A survey on active learning and human-in-the-loop deep learning for medical image analysis. *Medical Image Analysis* 2021;71:102062.