3D Deep Learning Angiography (3D-DLA) from C-arm Conebeam CT

J.C. Montoya, Y. Li, C. Strother, and G.-H. Chen

ABSTRACT

BACKGROUND AND PURPOSE: Deep learning is a branch of artificial intelligence that has demonstrated unprecedented performance in many medical imaging applications. Our purpose was to develop a deep learning angiography method to generate 3D cerebral angiograms from a single contrast-enhanced C-arm conebeam CT acquisition in order to reduce image artifacts and radiation dose.

MATERIALS AND METHODS: A set of 105 3D rotational angiography examinations were randomly selected from an internal data base. All were acquired using a clinical system in conjunction with a standard injection protocol. More than 150 million labeled voxels from 35 subjects were used for training. A deep convolutional neural network was trained to classify each image voxel into 3 tissue types (vasculature, bone, and soft tissue). The trained deep learning angiography model was then applied for tissue classification into a validation cohort of 8 subjects and a final testing cohort of the remaining 62 subjects. The final vasculature tissue class was used to generate the 3D deep learning angiography images. To quantify the generalization error of the trained model, we calculated the accuracy, sensitivity, precision, and Dice similarity coefficients for vasculature classification in relevant anatomy. The 3D deep learning angiography and clinical 3D rotational angiography images were subjected to a qualitative assessment for the presence of intersweep motion artifacts.

RESULTS: Vasculature classification accuracy and 95% CI in the testing dataset were 98.7% (98.3%–99.1%). No residual signal from osseous structures was observed for any 3D deep learning angiography testing cases except for small regions in the otic capsule and nasal cavity compared with 37% (23/62) of the 3D rotational angiographies.

CONCLUSIONS: Deep learning angiography accurately recreated the vascular anatomy of the 3D rotational angiography reconstructions without a mask. Deep learning angiography reduced misregistration artifacts induced by intersweep motion, and it reduced radiation exposure required to obtain clinically useful 3D rotational angiography.

ABBREVIATIONS: CNN = convolutional neural network; 3DRA = 3D rotational angiography; DLA = deep learning angiography; GPU = graphics processing unit

Cerebrovascular diseases are common causes of morbidity and mortality in the adult population worldwide.^{1–3} Most cerebrovascular diseases are found during routine brain imaging with CT or MR imaging; however, 2D-DSA remains the criterion standard for their accurate angiographic evaluation and characterization, in particular for arteriovenous malformations,⁴ cerebral aneurysms,^{5,6} and dural arteriovenous fistulas.⁷ Additional 3D

rotational angiography (3DRA) is used to improve the visualization and spatial understanding of vascular structures during the diagnostic work-up of these conditions. Currently, with many angiographic systems, obtaining a 3DRA still requires 2 rotational acquisitions, one without injection of contrast (mask run) and one during injection of contrast (fill run). These 2 datasets are used to compute log-subtracted projections, which are then used to reconstruct a subtracted 3DRA volume.^{8,9}

Machine learning is a discipline within computer science, closely related to statistics and mathematic optimization, that aims to learn patterns directly from a large set of examples that demonstrate a desired outcome or behavior without the need of explicit instructions.¹⁰ In the context of medical imaging, machine learning methods have been investigated since the early 1990s, initially for computer-aided detection and diagnosis in mammography and pulmonary embolism^{11–14}; however, recent advances in deep learning¹⁵ (ie, a specific machine-learning tech-

Received September 1, 2017; accepted after revision January 26, 2018.

From the Departments of Medical Physics (J.C.M., Y.L., G.-H.C.) and Radiology (C.S., G.-H.C.), University of Wisconsin School of Medicine and Public Health, Madison, Wisconsin.

The work was partially supported by a National Institutes of Health grant U01EB021183.

Please address correspondence to Guang-Hong Chen, PhD, L1167, WIMR, 1111 Highland Ave, Madison, WI 53705; e-mail: gchen7@wisc.edu

Indicates open access to non-subscribers at www.ajnr.org

<http://dx.doi.org/10.3174/ajnr.A5597>

FIG 1. Overall study schema.

nique) have demonstrated unprecedented performance in many applications, including detection of diabetic retinopathy¹⁶ and breast cancer,^{17,18} quantitative analysis of brain tumors in MR imaging,^{19,20} computer-aided detection of cerebral aneurysms in MR angiography,²¹ and computer-aided detection and classification of thoracic diseases.^{22,23}

With recent advances in deep learning and the universal approximation properties of feedforward neural networks,^{24,25} it is hypothesized that a deep neural network is capable of computing cerebral angiograms with only the vascular information contained in the fill scan of a 3DRA examination acquired with a C-arm conebeam CT system. Potential benefits of eliminating the mask scan include the following: 1) reduction of intersweep patient motion artifacts caused by the misregistration of the mask and fill scans, and 2) radiation dose reduction by at least a factor of 2.

The purpose of this work was to develop and test the capability of a deep learning angiography (DLA) method based on convolutional neural networks (CNNs) to generate subtracted 3D cerebral angiograms from a single contrast-enhanced examination without the need for a mask acquisition.

MATERIALS AND METHODS

In the following sections, the patient inclusion criteria and image-acquisition protocols are first presented, followed by a description of the datasets and methods used to train the DLA model. Finally, the image analysis and statistical analysis are described. The overall study schema is shown in Fig 1.

Patient Cohort

All studies were Health Insurance Portability and Accountability Act-compliant and performed under an institutional review board-approved protocol. Clinically indicated rotational angiography examinations for the assessment of cerebrovascular abnormalities of 105 patients, scanned from August 2014 through April 2016, were retrospectively collected. Cases were selected in a random fashion to reduce the potential bias in patient selection. It was thought that the randomized selection during this period would result in a dataset that was representative of the varieties of conditions that are referred for angiographic studies.

Imaging Acquisition and Reconstruction

All subjects were imaged with a standard 3DRA data acquisition protocol using a C-arm conebeam CT system (Axiom Artis zee; Siemens, Erlangen, Germany). The protocol consists of 2

conebeam CT acquisitions (ie, mask and fill acquisitions) with 172 or 304 projection images for a 6- or 13-second rotation time, respectively. Angular coverage for all data acquisitions was 260°, with a tube potential of 70 kVp, detector dose per projection image equal to 0.36 μ Gy per frame, and angular increments of 1.52° or 0.85° per frame. Iodinated contrast medium was injected into the proximal internal carotid artery or vertebral artery just after the initiation of the fill acquisition. For each subject, “native fill” and subtracted 3D volumes were reconstructed using the vendor’s proprietary software (InSpace Reconstruction, syngo Workplace; Siemens). All reconstructions were performed using the standard filtered back-projection with edge enhancement, normal image characteristic, full FOV ($238 \times 238 \text{ mm}^2$) with a 512×512 image matrix, and 0.46-mm image thickness/increment for a 0.46-mm isotropic voxel size. The effective dose for the acquisition protocols used in this study was 1.1 mSv for the 6-second rotation acquisition and 1.8 mSv for 13-second rotation acquisition, which is similar to the dose level reported by others.^{26,27}

Training Dataset

A training dataset consisting of 13,790 axial images from 35 patients with >150 million labeled voxels was generated using the information from both the conebeam CT image of the fill scan and the subtracted images from the subtracted conebeam projection data. For each patient in the training dataset, vasculature extraction was performed by a manual thresholding of the subtracted images. The selection of the threshold was based on the subjective assessment of complete vasculature segmentation while excluding image artifacts and background noise, with threshold values typically in the range of 500–700 HU. Large vessels, specifically the internal carotid artery, middle cerebral artery, anterior cerebral artery, distal branches of the middle cerebral artery and anterior cerebral artery, vertebral artery, and posterior cerebral artery were isolated through 3D connected component analysis.²⁸ Small regions not connected to a large vessel were assumed to be image artifacts and were excluded from the final vasculature volume. After the previous steps, in the event of remaining intersweep patient motion and streak artifacts, the vasculature volume was subjected to manual artifact removal.

The extraction of bone tissue was performed by subtracting the vasculature volume from the contrast-enhanced images (ie, fill scan) and performing manual thresholding and connectivity analysis (like that of vasculature extraction) in the resulting images. Only connected regions including the skull and mandible were considered bone. The remaining streaking artifacts and metal implants in the bone volume were manually removed. Finally, the soft-tissue class was extracted by thresholding the fill images with thresholds of -400 to 500 HU and applying a morphologic erosion.

The procedure described above generates approximately 0.28 million, 6 million, and 15 million voxels of vasculature, bone, and soft tissue, respectively, for each patient. To mitigate the class imbalance (ie, different number of labeled voxels per tissue class) and reduce redundant training data by a similarity of adjacent voxels, we included only 4.3 million labeled voxels per patient for training, consisting of all vasculature voxels and an equal number

of randomly extracted bone and soft-tissue voxels (ie, random undersampling).²⁹

Validation and Testing Datasets

A validation dataset and a testing dataset were created using the remaining image volumes from 70 subjects divided into 8 examinations for the validation dataset and 62 examinations for the testing dataset. These datasets were created with the same procedure used to generate the training dataset; however, the tissue labels were constrained to a region only containing the following anatomy: ICA, middle cerebral artery, anterior cerebral artery, distal branches of the middle cerebral artery and anterior cerebral artery, vertebral artery, posterior cerebral artery, the base and anterior aspect of the skull, temporal bone, otic capsule, and surrounding soft tissue as opposed to the entire head in the training dataset. Each examination in the validation and testing dataset had approximately the same number of labeled voxels for each tissue class.

Neural Network Architecture and Implementation

A 30-layer CNN³⁰ with a ResNet architecture,^{31,32} as shown in Fig 2, was used. All convolutional layers except the input layer used 3×3 filters with rectified linear units for activation function. The input of the network was a $41 \times 41 \times 5$ volumetric image patch extracted from the contrast-enhanced image volume; the network output consisted of a 3-way fully connected layer with softmax activation. Training and inference were performed on a voxelwise basis, in which the input volumetric image patch was labeled with the tissue class of its central voxel. The DLA model was implemented using Tensorflow (Google, Mountain View, California). Network parameters were initialized using the variance scaling method³³ and trained from scratch using a synchronous stochastic gradient descent method with a batch size of 512 volumetric image patches using 2 GTX 1080 Ti (NVIDIA, Santa Clara, California) graphics processing units (GPUs) (256 image patches per GPU). The time required to process 1 case in this study varied from 1 to 3 minutes, depending on the size of the image volume.

Each tissue class had equal probability of being included in a single batch (ie, data resampling), to account for class imbalance.²⁹ The learning rate was initially set to 1×10^{-3} with a momentum of 0.9. The learning rate was reduced to 1×10^{-4} and 1×10^{-5} after 1 and 1.5 epochs, respectively. The validation dataset was used only to monitor the convergence and generalization error during model training. Early stopping was used when the validation error reached a plateau at 2×10^5 iterations.

Statistical Analysis

The trained DLA model was applied for the task of tissue classification in the validation and testing cohorts, consisting of image volumes from 8 and 62 subjects, respectively. The final vasculature tissue class was used to generate the 3D-DLA images. To quantify the generalization error of the trained model, we evaluated the vasculature classification for each labeled voxel in the reference standard for the validation and testing datasets. Two-by-two tables were generated for each patient, and accuracy, sen-

FIG 2. Neural network architecture.

sitivity (also known as recall), positive predictive value (also known as precision), and Dice similarity coefficients were calculated. The 95% CIs for each performance metric were also reported. Finally, the clinical 3DRA and the 3D-DLA images were subjected to a qualitative assessment for the presence of intersweep motion artifacts, and results were expressed as frequencies and percentages.

Summary of performance metrics for vascular classification in the training, validation, and testing datasets

Dataset	Sensitivity (Recall)		PPV (Precision)		DSC (FI Score)		Accuracy	
	TP		TP		2TP		TP + TN	
	TP + FN		TP + FP		2TP + FP + FN		TP + TN + FP + FN	
	Mean	95% CI	Mean	95% CI	Mean	95% CI	Mean	95% CI
Validation ($n = 8$)	97.8%	(96.9%–98.7%)	97.2%	(96.4%–98.1%)	97.5%	(97.0%–98.0%)	98.4%	(98.0%–98.7%)
Testing ($n = 62$)	97.6%	(96.5%–98.6%)	98.5%	(97.6%–99.3%)	98.0%	(97.4%–98.7%)	98.7%	(98.3%–99.1%)

Note:—TP indicates true-positive; FN, false-negative; PPV, positive predictive value; FP, false-positive; TN, true-negative; FI, Dice similarity coefficients.

RESULTS

Contrast-enhanced image volumes from 105 subjects (53.3 ± 13.5 years of age) who underwent clinically indicated rotational angiography examinations for the assessment of cerebrovascular abnormalities were used in the study. Contrast medium was injected via the proximal ICA in 89 patients (85%) and via the vertebral artery in 16 patients (15%). Average and 95% CIs for vasculature classification accuracy, sensitivity, positive predictive value, and Dice similarity coefficient in the testing dataset were 98.7% (95% CI, 98.3%–99.1%), 97.6% (95% CI, 96.5%–98.6%), 98.5% (95% CI, 97.6%–99.3%), and 98.0% (95% CI, 97.4%–98.7%), respectively. The Table summarizes the performance metrics for vascular classification in the training, validation, and testing datasets.

No residual signal from osseous structures was observed for any testing cases generated using 3D-DLA except for small regions in the otic capsule and nasal cavity compared with 37% (23/62) of the 3DRA cases that presented residual bone artifacts. Figure 3 shows a comparison of MIP images derived from 3DRA and the 3D-DLA datasets of a patient evaluated for posterior cerebral circulation. One can see how residual bone artifacts induced by intersweep patient motion are greatly reduced in 3D-DLA, improving the conspicuity of small vessels. Similarly, Fig 4 shows lateral and oblique MIP images derived from 3DRA and the 3D-DLA datasets of a patient evaluated for anterior cerebral circulation. Results show reduced residual bone artifacts for 3D-DLA images, in particular for the anterior aspect of the skull and the temporal bone. Figure 5 shows a comparison of volume-rendering images for both the clinical 3DRA and the 3D-DLA of a patient with a small aneurysm in the anterior communicating artery and a large aneurysm in the middle cerebral artery bifurcation.

DISCUSSION

In this work, a deep CNN was used to learn generic opacified vasculature from contrast-enhanced C-arm conebeam CT datasets to generate a 3D cerebral angiogram, without an explicit definition of cerebrovascular diseases or specific vascular anatomy. The datasets used for model training, validation, and testing were created by applying simple image-processing techniques with minimum manual editing for a total of 82,740 subtracted and contrast-enhanced conebeam CT images from 105 subjects. The proposed DLA method was used to improve image quality by reducing image artifacts caused by misregistration of mask and fill scans in 3DRA, in addition to enabling potential radiation dose reduction.

Many angiographic systems require 2 rotational acquisitions (mask and fill) for reconstruction of a subtracted 3DRA. Others, using vascular segmentation and thresholding algorithms, allow a

FIG 3. Comparison of anterior and lateral views of MIP images derived from 3DRA (A) and 3D-DLA (B) datasets of a patient evaluated for posterior cerebral circulation. Residual bone artifacts induced by intersweep patient motion are greatly reduced in 3D-DLA, improving the conspicuity of small vessels as indicated by the white arrowheads.

3D vessel reconstruction without the availability of a mask. Those that require 2 rotations are susceptible to artifacts caused by potential misregistrations of the mask and fill projections. Those that require the use of segmentation and thresholding algorithms may be subject to errors related to too little contrast intensity and/or improper segmentation. Together, these techniques remain the standard of care for the diagnosis and treatment planning of cerebrovascular diseases. Misregistration artifacts arise in conventional 3DRA imaging primarily due to the following: 1) small variations in the angular range differences occurring from one rotational acquisition to another, and 2) potential patient motion in both mask and fill runs. The mask-free DLA method, by eliminating the need for one of the rotational acquisitions, in theory, would reduce the chance of motion from both mechanical instability and patient motion and effectively reduce the radiation

FIG 4. Comparison of lateral and oblique MIP images derived from 3DRA (A) and 3D-DLA (B) datasets of a patient evaluated for anterior cerebral circulation. Results show reduced residual bone artifacts for 3D-DLA images, in particular for the anterior aspect of the skull and the temporal bone.

FIG 5. Comparison of volume-rendering images for both the clinical 3DRA (A) and the 3D-DLA (B) of a patient with a small aneurysm in the anterior communicating artery and a large aneurysm in the middle cerebral artery bifurcation as indicated by the *arrowheads*.

dose required to obtain a 3DRA by half in those systems that require 2 rotations.

In the context of medical imaging, machine-learning methods have been investigated since early 1990s^{11–14}; however, the recent unprecedented performance of deep learning has made major

advances in solving very difficult problems in science that were thought to be intractable when approached by other means.^{34,35} In addition to clever mathematical techniques and the availability of large annotated datasets, many authors recognize that the massively parallel computing capabilities of GPUs have played a key role in the success of deep learning applications, providing accelerations of 40× to 250× compared with multicore and single-core CPUs.^{10,15} For example, the training procedure of the network used in this study took approximately 23 hours. This training procedure could have taken 4–5 weeks if only a multicore CPU computation architecture was used, making this application impractical. Fortunately, the training procedure is performed off-line, it only needs to be done once, and GPU computing is already widely available within the medical imaging community or accessible via cloud computing services such as the Google Cloud Platform (<https://cloud.google.com/>) or the Amazon Web Service (<https://aws.amazon.com/>). Also, many standard open-source libraries used for deep learning applications are highly optimized to be used in conjunction with GPUs.

Once the parameters of the model have been learned, the process of analyzing new data that were not used for training the model (ie, inference) can be further optimized for production. The method proposed in this study uses a voxelwise training and inference in which the input of the network is a small image region of $41 \times 41 \times 5$ voxels around the voxel of interest. This approach has multiple benefits: 1) Inference can be parallelized; in other words, the classification of multiple voxels can be performed at the same time. Therefore, the time required to analyze a new case is directly proportional to the number of voxels to be classified (eg, the entire head or a targeted ROI) and the number and generation of available GPUs. The throughput of the particular research implementation of the CNN model used in this study is approximately 2500 voxels/s/GPU. 2) This approach results in a large training dataset consisting of 150 million labeled voxels derived from 13,790 axial images and 35 examinations. In addition to a large training dataset, it is important to have a large testing cohort to assure a good model generalization that better reflects how this technique could be used in practice. Having a large testing cohort also helps to determine whether the training dataset is large enough to achieve a desired level of performance.

Although DLA images were successfully created for all validation and testing cases and were subjected to quantitative and qualitative image analysis, this study still has some limitations: First, the use of a very specific image-acquisition protocol and reconstruction with selective intra-arterial contrast media injection into the proximal internal carotid artery or vertebral artery may limit its clinical application. Fine-tuning of the model and clinical validation with prospective reader studies are required to further generalize these results to the vasculature of other organ systems, to complex or uncommon vascular abnormalities, as well as to angiography studies acquired using different image-acquisition protocols and modalities (eg, injection of IV contrast media, time-resolved 3DRA, multidetector CT, and so forth). This kind of prospective reader study would also overcome the limitation of the current qualitative evaluation in our study by a single reader (C.S.).

Second, in this study, a specific type of deep CNN with 30

layers, implemented with a ResNet^{31,32} architecture, was used to demonstrate the DLA application. The selection of the well-known 30-layer ResNet is only because this architecture won an image contest among computer scientists (ie, ILSVRC 2015 classification task; <http://www.image-net.org/challenges/LSVRC/2015/results>); namely, it outperformed other types of networks such as AlexNet, VGGNet, and GoogLeNet, for the task of natural image classification using the ImageNet (<http://image-net.org/>) dataset. This type of network architecture has also outperformed other types of networks in medical imaging classification tasks with deeper models (ie, increased number of layers) having improved classification accuracy.^{20,22,32,36} However, it remains unknown whether other architectures can be used for DLA and what would be the advantage or disadvantage among all these networks. Furthermore, additional optimization and fine-tuning of the DLA model hyperparameters (eg, number of layers, number of hidden units per layer, learning rate, regularization schemes, and so forth) are required for optimal on-line implementation and compatibility with clinical workflow.

Third, even though metallic objects are automatically subtracted in the 3DRA images that were used to create the training dataset, small movements of metallic implants (eg, an aneurysm clip or a coil mass) that occur during a cardiac cycle are, in the case of subtracted images, usually sufficient to create enough misregistration artifacts to allow detection of an implant presence. This situation, in addition to the high x-ray attenuation and proximity to vasculature of metallic implants, could result in their imitation in the final DLA images. The presence of a high-attenuating object (eg, metal or Onyx [Covidien, Irvine, California]) is also known to be an intrinsic limitation of mask-free angiography (vendors who provide a method to obtain 3DRAs without a mask also offer the ability to perform a mask and fill acquisition in situations in which metal objects are known to be present), and its clinical implications need to be addressed with an expert reader study.

CONCLUSIONS

A DLA method based on CNNs that generates 3D cerebral angiograms from a contrast-enhanced C-arm conebeam CT without mask data acquisition was developed. Results indicate that the proposed method can successfully reduce misregistration artifacts induced by intersweep patient motion and, by eliminating the need for a mask acquisition, can reduce the radiation dose in future clinical 3D angiography.

ACKNOWLEDGMENTS

The authors thank Dr John W. Garrett for technical and editorial assistance.

Disclosures: Guang-Hong Chen—RELATED: Grant: National Institutes of Health*; UNRELATED: Grants/Grants Pending: GE Healthcare, Siemens, Hologic, Comments: Principal Investigator of research contracts for unrelated research topics*; *Money paid to the institution.

REFERENCES

- Go AS, Mozaffarian D, Roger VL, et al; American Heart Association Statistics Committee and Stroke Statistics Subcommittee. **Executive summary: heart disease and stroke statistics—2013 update: a report from the American Heart Association.** *Circulation* 2013;127:143–52 CrossRef Medline
- Wiebers DO. **Unruptured intracranial aneurysms: natural history, clinical outcome, and risks of surgical and endovascular treatment.** *Lancet* 2003;362:103–10 CrossRef Medline
- Mohr JP, Parides MK, Stapf C, et al; international ARUBA investigators. **Medical management with or without interventional therapy for unruptured brain arteriovenous malformations (ARUBA): a multicentre, non-blinded, randomised trial.** *Lancet* 2014;383:614–21 CrossRef Medline
- Ogilvy CS, Stieg PE, Awad I, et al; Special Writing Group of the Stroke Council, American Stroke Association. **AHA scientific statement: recommendations for the management of intracranial arteriovenous malformations: a statement for healthcare professionals from a special writing group of the Stroke Council, American Stroke Association.** *Stroke* 2001;32:1458–71 CrossRef Medline
- Hacein-Bey L, Provenzale JM. **Current imaging assessment and treatment of intracranial aneurysms.** *AJR Am J Roentgenol* 2011;196:32–44 CrossRef Medline
- Anxionnat R, Bracard S, Ducrocq X, et al. **Intracranial aneurysms: clinical value of 3D digital subtraction angiography in the therapeutic decision and endovascular treatment.** *Radiology* 2001;218:799–808 CrossRef Medline
- Gandhi D, Chen J, Pearl M, et al. **Intracranial dural arteriovenous fistulas: classification, imaging findings, and treatment.** *AJNR Am J Neuroradiol* 2012;33:1007–13 CrossRef Medline
- Fahrig R, Fox AJ, Lownie S, et al. **Use of a C-arm system to generate true three-dimensional computed rotational angiograms: preliminary in vitro and in vivo results.** *AJNR Am J Neuroradiol* 1997;18:1507–14 Medline
- Strobel N, Meissner O, Boese J, et al. **3D imaging with flat-detector C-arm systems.** In: Reiser MF, Becker CR, Nikolaou K, et al, eds. *Multislice CT*. Berlin: Springer-Verlag; 2009:33–51
- Erickson BJ, Korfiatis P, Akkus Z, et al. **Machine learning for medical imaging.** *Radiographics* 2017;37:505–15 CrossRef Medline
- Chan HP, Lo SC, Sahiner B, et al. **Computer-aided detection of mammographic microcalcifications: pattern recognition with an artificial neural network.** *Med Phys* 1995;22:1555–67 CrossRef Medline
- Wu Y, Doi K, Giger ML, et al. **Computerized detection of clustered microcalcifications in digital mammograms: applications of artificial neural networks.** *Med Phys* 1992;19:555–60 CrossRef Medline
- Zhang W, Doi K, Giger ML, et al. **Computerized detection of clustered microcalcifications in digital mammograms using a shift-invariant artificial neural network.** *Med Phys* 1994;21:517–24 CrossRef Medline
- Wang S, Summers RM. **Machine learning and radiology.** *Med Image Anal* 2012;16:933–51 CrossRef Medline
- LeCun Y, Bengio Y, Hinton G. **Deep learning.** *Nature* 2015;521:436–44 CrossRef Medline
- Gulshan V, Peng L, Coram M, et al. **Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs.** *JAMA* 2016;316:2402–10 CrossRef Medline
- Huynh BQ, Li H, Giger ML. **Digital mammographic tumor classification using transfer learning from deep convolutional neural networks.** *J Med Imaging (Bellingham)* 2016;3:034501 CrossRef Medline
- Antropova N, Huynh BQ, Giger ML. **A deep feature fusion methodology for breast cancer diagnosis demonstrated on three imaging modality datasets.** *Med Phys* 2017;44:5162–71 CrossRef Medline
- Akkus Z, Galimzianova A, Hoogi A, et al. **Deep learning for brain MRI segmentation: state of the art and future directions.** *J Dig Imaging* 2017;30:449–59 CrossRef Medline
- Korfiatis P, Kline TL, Lachance DH, et al. **Residual deep convolutional neural network predicts MGMT methylation status.** *J Dig Imaging* 2017;30:622–28 CrossRef Medline
- Nakao T, Hanaoka S, Nomura Y, et al. **Deep neural network-based computer-assisted detection of cerebral aneurysms in MR angiography.** *J Magn Reson Imaging* 2017 Aug 24. [Epub ahead of print] CrossRef Medline

22. Wang X, Peng Y, Lu L, et al. **ChestX-ray8: hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases.** In: *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, Hawaii. July 21–26, 2017
23. Lakhani P, Sundaram B. **Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks.** *Radiology* 2017;284:574–82 [CrossRef Medline](#)
24. Cybenko G. **Approximation by superpositions of a sigmoidal function.** *Mathematics of Control, Signals and Systems* 1989;2:303–14 [CrossRef](#)
25. Hornik K. **Approximation capabilities of multilayer feedforward networks.** *Neural Networks* 1991;4:251–57 [CrossRef](#)
26. Struffert T, Hauer M, Banckwitz R, et al. **Effective dose to patient measurements in flat-detector and multislice computed tomography: a comparison of applications in neuroradiology.** *Eur Radiol* 2014;24:1257–65 [CrossRef Medline](#)
27. Lang S, Göllitz P, Struffert T, et al. **4D DSA for dynamic visualization of cerebral vasculature: a single-center experience in 26 cases.** *AJNR Am J Neuroradiol* 2017;38:1169–76 [CrossRef Medline](#)
28. Shapiro LG. **Connected component labeling and adjacency graph construction.** *Machine Intelligence and Pattern Recognition* 1996;19:1–30
29. He H, Garcia EA. **Learning from imbalanced data.** *IEEE Transactions on Knowledge and Data Engineering* 2009;21:1263–84
30. LeCun Y, Boser BE, Denker JS, et al. **Handwritten digit recognition with a back-propagation network.** In: David ST, ed. *Advances in Neural Information Processing Systems* 2. San Francisco: Morgan Kaufmann Publishers;1990:396–404
31. He K, Zhang X, Ren S, et al. **Identity mappings in deep residual networks.** In: Leibe B, Matas J, Sebe N, et al, eds. *Proceedings of Computer Vision—ECCV 2016: 14th European Conference*, Amsterdam, the Netherlands. October 1–14, 2016
32. He K, Zhang X, Ren S, et al. **Deep residual learning for image recognition.** In: *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, Washington. June 27–30, 2016:770–78
33. He K, Zhang X, Ren S, et al. **Delving deep into rectifiers: surpassing human-level performance on ImageNet classification.** In: *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, Los Alamitos, California. December 7–13, 2015:1026–34
34. Silver D, Huang A, Maddison CJ, et al. **Mastering the game of Go with deep neural networks and tree search.** *Nature* 2016;529:484–89 [CrossRef Medline](#)
35. Silver D, Schrittwieser J, Simonyan K, et al. **Mastering the game of Go without human knowledge.** *Nature* 2017;550:354–59 [CrossRef Medline](#)
36. Shin HC, Roth HR, Gao M, et al. **Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning.** *IEEE Trans Med Imaging* 2016;35:1285–98 [CrossRef Medline](#)