Radiology: Artificial Intelligence

AI IN BRIEF

Deep Learning–based Detection of Intravenous Contrast Enhancement on CT Scans

Zezhong Ye, PhD • Jack M. Qian, MD • Ahmed Hosny, PhD • Roman Zeleznik, PhD • Deborah Plana, PhD • Jirapat Likitlersuang, PhD • Zhongyi Zhang, BS • Raymond H. Mak, MD • Hugo J. W. L. Aerts, PhD • Benjamin H. Kann, MD

From the Artificial Intelligence in Medicine Program, Mass General Brigham, Harvard Medical School, Harvard Institutes of Medicine, Boston, Mass (Z.Y., J.M.Q., A.H., R.Z., D.P., J.L., Z.Z., R.H.M., H.J.W.L.A., B.H.K.); Departments of Radiation Oncology (Z.Y., J.M.Q., A.H., R.Z., J.L., Z.Z., R.H.M., H.J.W.L.A., B.H.K.) and Radiology (H.J.W.L.A.), Dana-Farber Cancer Institute, Brigham and Women's Hospital, Harvard Medical School, 75 Francis St, Boston, MA 02115; Harvard–MIT Division of Health Sciences & Technology, Cambridge, Mass (D.P.); and Department of Radiology and Nuclear Medicine, School for Cardiovascular Diseases (CARIM) & School for Oncology and Reproduction (GROW), Maastricht University, Maastricht, the Netherlands (H.J.W.L.A.). Received November 12, 2021; revision requested January 7, 2022; revision received March 24; accepted April 14. **Address correspondence to** B.H.K. (email: *Benjamin_Kann@dfci.harvard.edu*).

Supported by the National Institutes of Health (H.J.W.L.A.: NIH-USA U24CA194354, NIH-USA U01CA190234, NIH-USA U01CA209414, and NIH-USA R35CA22052; B.H.K.: NIH-K08:DE030216), the European Union–European Research Council (H.J.W.L.A.: 866504), the Radiological Society of North America (B.H.K.: RSCH2017), and the National Institute of General Medical Sciences (D.P.: T32-GM007753).

Conflicts of interest are listed at the end of this article.

Radiology: Artificial Intelligence 2022; 4(3):e210285 • https://doi.org/10.1148/ryai.210285 • Content codes: AI CH CT HN

Identifying the presence of intravenous contrast material on CT scans is an important component of data curation for medical imaging–based artificial intelligence model development and deployment. Use of intravenous contrast material is often poorly documented in imaging metadata, necessitating impractical manual annotation by clinician experts. Authors developed a convolutional neural network (CNN)–based deep learning platform to identify intravenous contrast enhancement on CT scans. For model development and validation, authors used six independent datasets of head and neck (HN) and chest CT scans, totaling 133480 axial two-dimensional sections from 1979 scans, which were manually annotated by clinical experts. Five CNN models were trained first on HN scans for contrast enhancement detection. Model performances were evaluated at the patient level on a holdout set and external test set. Models were then fine-tuned on chest CT data and externally validated. This study found that Digital Imaging and Communications in Medicine metadata tags for intravenous contrast material were missing or erroneous for 1496 scans (75.6%). An EfficientNetB4-based model showed the best performance, with areas under the curve (AUCs) of 0.996 and 1.0 in HN holdout (n = 216) and external (n = 595) sets, respectively, and AUCs of 1.0 and 0.980 in the chest holdout (n = 53) and external (n = 402) sets, respectively. This automated, scan-to-prediction platform is highly accurate at CT contrast enhancement detection and may be helpful for artificial intelligence model development and clinical application.

Supplemental material is available for this article.

The clinical translation of artificial intelligence for medical image analysis faces many challenges (1). Critical among these are data quality and curation, which represent difficult, labor-intensive, and largely manual processes conducted by clinical experts such as radiologists and radiation oncologists (2). The characterization of image-specific parameters relies on metadata tags from the Digital Imaging and Communications in Medicine (DICOM) standard, which was designed as a clinical protocol and not for downstream computational analysis (3). Furthermore, certain types of metadata are input manually by scanner operators and are notoriously poorly documented and error prone (4–6). One such parameter is the administration of intravenous contrast material (5,6).

The presence or absence of intravenous contrast material has large ramifications for computational imaging model performance and is essential knowledge for imaging analyses (7–9). Currently, the only reliable way to detect contrast enhancement on a scan is through manual review by clinical experts, which is time-consuming and often impractical. With the growing interest in using large datasets to develop computational models, there is a need for automated tools that can detect intravenous contrast enhancement with high fidelity. Several conventional computer vision methods, including a hybrid discriminative–generative model (5) and a multiclass LogitBoost classifier (6), have been previously adopted for contrast enhancement detection with adequate performance, although these models were not externally validated and require region localization steps prior to contrast enhancement prediction, which may limit generalizability. Recently, deep learning has demonstrated tremendous promise for medical imaging classification (10). There have been two prior studies investigating deep learning for contrast-phase detection, one for abdominal CT (11) and one for the kidney (12). Both achieved promising performance on internal test sets. To our knowledge, there exist no models for the detection of intravenous contrast enhancement on CT scans that have been externally validated. We hypothesized that a deep learning model implementing convolutional neural networks (CNNs) could be developed and externally validated to reliably and accurately detect intravenous contrast enhancement on CT scans.

Materials and Methods

Study Design and Datasets

This study was conducted in accordance with the Declaration of Helsinki guidelines and received approval

Abbreviations

AUC = area under the curve, CNN = convolutional neural network, DICOM = Digital Imaging and Communications in Medicine, HN = head and neck

Summary

Authors developed and externally validated a deep learning model that accurately detects intravenous contrast enhancement on head and neck CT scans and chest CT scans efficiently and with a nearly perfect performance.

Key Points

- We used 1979 head and neck (HN) and chest CT scans from multiple institutions to develop and validate a deep learning model to detect intravenous contrast enhancement.
- An EfficientNetB4-based model yielded areas under the curve (AUCs) of 0.996 in the internal validation set (n = 216) and 1.0 in the external test set (n = 595) for HN scans; by using transfer learning, the model was retrained to detect contrast enhancement on chest scans, yielding an AUC of 1.0 for the internal validation set (n = 53) and an AUC of 0.980 for the external test set (n = 402).

Keywords

CT, Head and Neck, Supervised Learning, Transfer Learning, Convolutional Neural Network (CNN), Machine Learning Algorithms, Contrast Material

from the local institutional review board. A waiver of consent was obtained from the institutional review board prior to research initiation as a result of using public datasets or conducting a retrospective study. Data from five institutions and one national clinical trial from 2001 through 2015 were included (Fig E1 [supplement]). The head and neck (HN) cancer dataset consists of four publicly available, de-identified patient cohorts, each downloaded and curated from The Cancer Imaging Archive, as follows: cohort 1 (n = 558) (13); cohort 2 (n = 101) and cohort 3 (n = 61) (14); and cohort 4 (n = 603) (15). The lung cancer dataset includes two patient cohorts, as follows: cohort 5 (n = 262) and cohort 6 (n = 402), which were derived from a national clinical trial (16). These subsets represent all scans that passed the initial quality control of DICOM metadata. Scans that excluded the HN portions (n = 3) and whole-body scans (n = 6) were excluded from analyses. Data from all patients (n = 1979) in this study have been used in previous publications, yet none of these studies have focused on intravenous contrast enhancement detection. CT scanning parameters are found in Appendix E1 (supplement).

Image Review and Annotations

All CT images were manually reviewed and annotated at the image axial-section level and the scan level for intravenous contrast material presence by a radiation oncologist (J.M.Q.), with 4 years of clinical experience (Fig E5 [supplement]), and then were further reviewed by a board-certified radiation oncologist with 7 years of clinical experience (B.H.K.) to confirm. CT image preprocessing steps are found in Appendix E1 (supplement).

Model Development, Training, and Validation

Five CNN models were investigated: one simple CNN model (Fig E2 [supplement]); three representative, published deep CNN models that have been top performers in classifications of large imaging datasets (ResNet101V2 [17], InceptionV3 [18], and EfficientNetB4 [19]); and a transfer learning approach based on ResNet101V2 with pretraining weights on ImageNet (20) (architecture details are found in Appendix E1 [supplement]).

After data preprocessing, HN scans from three patient cohorts (cohorts 1–3) were shuffled and randomly split into 70:30 for model training (n = 504 patients and 33 264 images) and internal validation (n = 216 patients and 14 256 images; Fig E1 [supplement]). The data partition was stratified by intravenous contrast enhancement. Scans from cohort 4 (n = 595 patients and 39 270 images) were used for the independent external test. Models were first trained and validated on the image level. Each CNN model was trained for up to 100 epochs on the training dataset and validated on the validation set. Models were constructed and trained by using TensorFlow 2.0 frameworks in Python version 3.8 on a Titan RTX graphics processing unit (NVIDIA) (Appendix E1 [supplement]).

To determine whether a model largely based on HN CT scans could generalize to chest CT, 80% of the cohort 5 chest CT dataset (n = 209 patients and 14 840 images) was used to fine-tune the HN model, and the remaining data (n = 53 patients and 3710 images) were used for internal validation. A separate cohort 6 dataset (n = 402 patients and 28 140 images) was used for the external test.

Model Performance and Statistical Analysis

The Pearson χ^2 test and the Kruskal-Wallis H test were performed to test the statistically significant differences among training, validation, and test datasets. Model performance at the patient level was primarily evaluated by using the patient probability score, calculated by averaging the probability scores of all the images of each scan (Fig 1). A value of .5 was used as the probability threshold to determine the model prediction class (contrast vs noncontrast) at both the image level and the patient level. Receiver operating characteristic analysis and area under the curve (AUC) analysis were adopted to assess model discrimination of intravenous contrast enhancement. Sensitivity and specificity values were calculated by using the optimal cutoff point with the Youden index. Precision-recall curves and F1 scores were calculated to provide information complementary to the receiver operating characteristic curve. The 95% CIs were calculated on the basis of results from more than 10 000 bootstrapped iterations. Statistical metrics and curves were calculated by using Scikit-learn packages in Python. The overall study workflow is found in Figure 1.

All source code and the model can be found at *https://github.com/AIM-Harvard/DeepContrast*. National Lung Screening Trial data including raw CT images may be requested from The Cancer Image Archive *(https://www.*

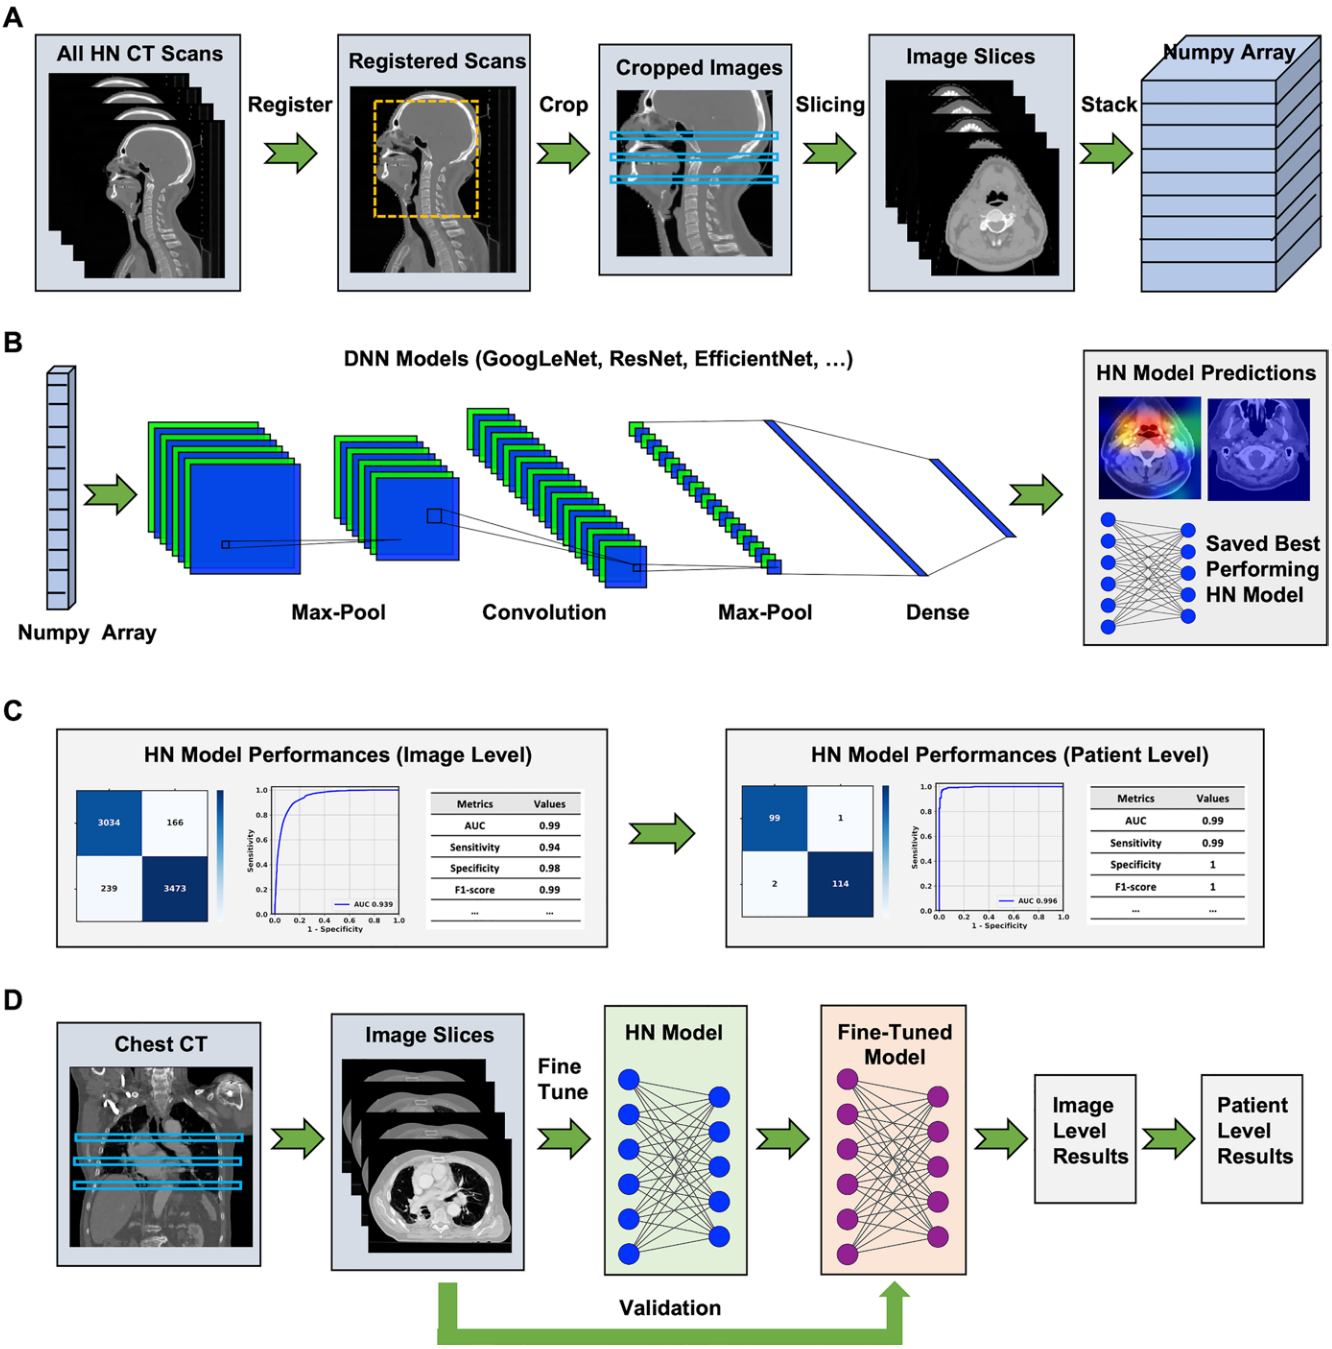

Figure 1: Workflow of deep neural networks (DNNs) for contrast enhancement detections. **(A)** All of the head and neck (HN) cancer CT scans were first co-registered to each other. The scans were then cropped to include only HN portions and exclude most background areas. Two-dimensional image sections were extracted from each scan and stacked together before being converted to NumPy arrays. **(B)** NumPy arrays with corresponding labels of each image section were fed into DNNs for model development and validation. We tested multiple published two-dimensional DNNs, including ResNet101V2, EfficientNetB4, InceptionV3, and a simple convolutional neural network (CNN). The models and prediction results were saved. **(C)** Image-level model performances were evaluated directly from model predictions. Patient-level model performances were then calculated by averaging the probability scores of each image section from each patient. **(D)** Chest CT scans went through the same imaging preprocessing before being input for training. We used a portion of lung images to fine-tune the saved models from HN datasets and applied other portions of lung images to validate the model performances at both the image level and the patient level. AUC = area under the curve.

cancerimagingarchive.net). Although raw CT imaging data cannot be shared, all measured results to replicate the statistical analysis are shared at the GitHub webpage: *https://github.com/AIM-Harvard/DeepContrast*. Furthermore, we include test samples from a publicly available dataset with deep learning and expert reader annotations.

Results

Patient and CT Scan Characteristics

The HN patient cohort consisted of 1315 patients (Table E7 [supplement]). Manual contrast annotation took 7.6 clinician hours for the HN scans (n = 1315), and 798 (60.7%)

Classification Performance of EfficientNetB4 Model on HN CT Scans and Lung CT Scans

Scan Type	Validation Type	Evaluation Level	AUC	Sensitivity (%)	Specificity (%)	F1 Score
HN CT	Internal validation	Image level (n = 33 264)	0.988 (0.988, 0.988)	95.9 (95.9, 96.0)	96.6 (96.6, 96.7)	0.964
		Patient level (n = 216)	0.996 (0.996, 0.996)	98.9 (98.8, 98.9)	99.9 (99.8, 99.9)	0.991
	External test	Image level (n = 39 270)	0.976 (0.976, 0.976)	95.1 (95.1, 95.1)	97.5 (97.5, 97.6)	0.970
		Patient level (n = 595)	1 (1,1)	100 (100, 100)	100 (100, 100)	0.999
Lung CT	Internal validation	Image level (n = 3710)	0.998 (0.998, 0.998)	96.0 (95.9, 96.0)	99.6 (99.5, 99.6)	0.973
		Patient level (n = 53)	1 (1,1)	100 (100, 100)	100 (100, 100)	1
	External test	Image level (n = 28 140)	0.948 (0.948, 0.949)	85.8 (85.6, 85.9)	91.4 (91.3, 91.5)	0.821
		Patient level (n = 402)	0.980 (0.980, 0.981)	96.9 (96.8, 97.0)	95.9 (95.8, 96.1)	0.923

Note.—Data in parentheses are 95% CIs. AUC = area under the curve, HN = head and neck.

scans were labeled as contrast enhanced. There were 491 scans (67.8%) documented with dental artifact among 724 patients coming from cohorts 1, 2, and 3 (Table E8 [supplement]). DICOM metadata for contrast material information (tags entered by technologists) was missing or erroneous in 808 scans (61.4%). The lung cancer patient cohort consisted of 664 chest scans (Table E9 [supplement]). Manual review required 6.1 hours of clinician time, and 197 scans (29.7%) were labeled as contrast enhanced. DICOM metadata for intravenous contrast material was missing for all the chest CT scans. Representative scans with contrast and without contrast enhancement can be found in Figure E5 (supplement).

Model Performance on HN and Chest CT Scans

For HN scans, all five models yielded excellent results, with patient-level AUCs greater than 0.98 and F1 scores greater than 0.96 on the internal holdout validation sets and external test sets at the patient level (Tables E10, E12–E15 [supplement]). EfficientNetB4 (Table, Fig 2) was selected as the most favorable model because it had the highest combined patient-level AUC and F1 score (AUC, 0.996; F1 score, 0.991) and because it has fewer parameters than ResNet101V2. On evaluation of performance metrics on the external test set, EfficientNetB4 yielded perfect patient-level classification performance, with an AUC of 1.0 (95% CI: 1.0, 1.0), a sensitivity of 100% (95% CI: 100%, 100%), a specificity of 100% (95% CI: 100%, 100%), and an F1 score of 1.0 for the patient-level prediction. EfficientNetB4 confusion matrices showed excellent agreement (Fig E3 [supplement]). Compared with artificial intelligence, the available metadata yielded an AUC of 0.185 and an AUC of 0.572 for intravenous contrast enhancement detection in the internal validation set and the external test set, respectively, with significant discordance indicated by confusion matrices (Fig E4 [supplement]). With model fine-tuning, the EfficientNetB4 model demonstrated an AUC of 1 (95% CI: 1.0, 1.0) and an AUC of 0.980 (95% CI: 0.980, 0.981) for the internal validation set and the external test set, respectively, at the patient level (Table, Fig 2). EfficientNetB4 still showed the best overall model performance among the five CNN models (Tables E11–E15 [supplement]). Including the image preprocessing, data loading, and prediction, the pipeline employing EfficientNetB4 took 2.1 hours to analyze all HN scans (n = 1315 scans, n = 86 790 axial sections) and 1.1 hours to analyze all chest scans (n = 664 scans, n = 46 690 axial sections). Error analysis indicated that faint contrast, artifacts, and dense vessels could cause false-positive predictions or false-negative predictions with the model (Fig E6 [supplement]).

Gradient-weighted Class Activation Maps for Model Interpretability

Qualitative analysis of gradient-weighted class activation heatmaps demonstrated that regions of importance were centered around the central blood vessels of the neck and chest (Fig 3).

Discussion

We developed a CNN-based deep learning platform for automated intravenous contrast enhancement detection on CT scans that demonstrated nearly perfect classification performance on several large datasets from a variety of institutions, clinical settings, and scanner types. With fine-tuning on small datasets, an intravenous contrast enhancement detection model developed for one anatomic site (HN) could be successfully applied to another region (chest). Data curation and quality assessment, including intravenous contrast enhancement confirmation, are extremely time- and resource-intensive manual processes. Similar to prior studies (7–9), this study found that intravenous contrast enhancement annotation from DICOM clinical metadata was often poorly documented and unreliable, with more than 70% of scans in our study missing or containing an erroneous contrast material status. Our model was more efficient than an expert clinician on contrast enhancement detection. It provides a usable tool that can be incorporated into research and clinical settings, obviating time-intensive manual annotation and review. Researchers conducting automated imaging classification and segmentation studies will find this platform useful in curating and performing quality assurance on their studies, saving a substantial amount of time and manual effort on annotation. Integrating this platform into the radiology workflow could help stratify contrast-enhanced and unenhanced CT scans and aid in the accurate reporting of study techniques and protocols. The platform could also be applied to clinical use cases, such as the identification of scans with

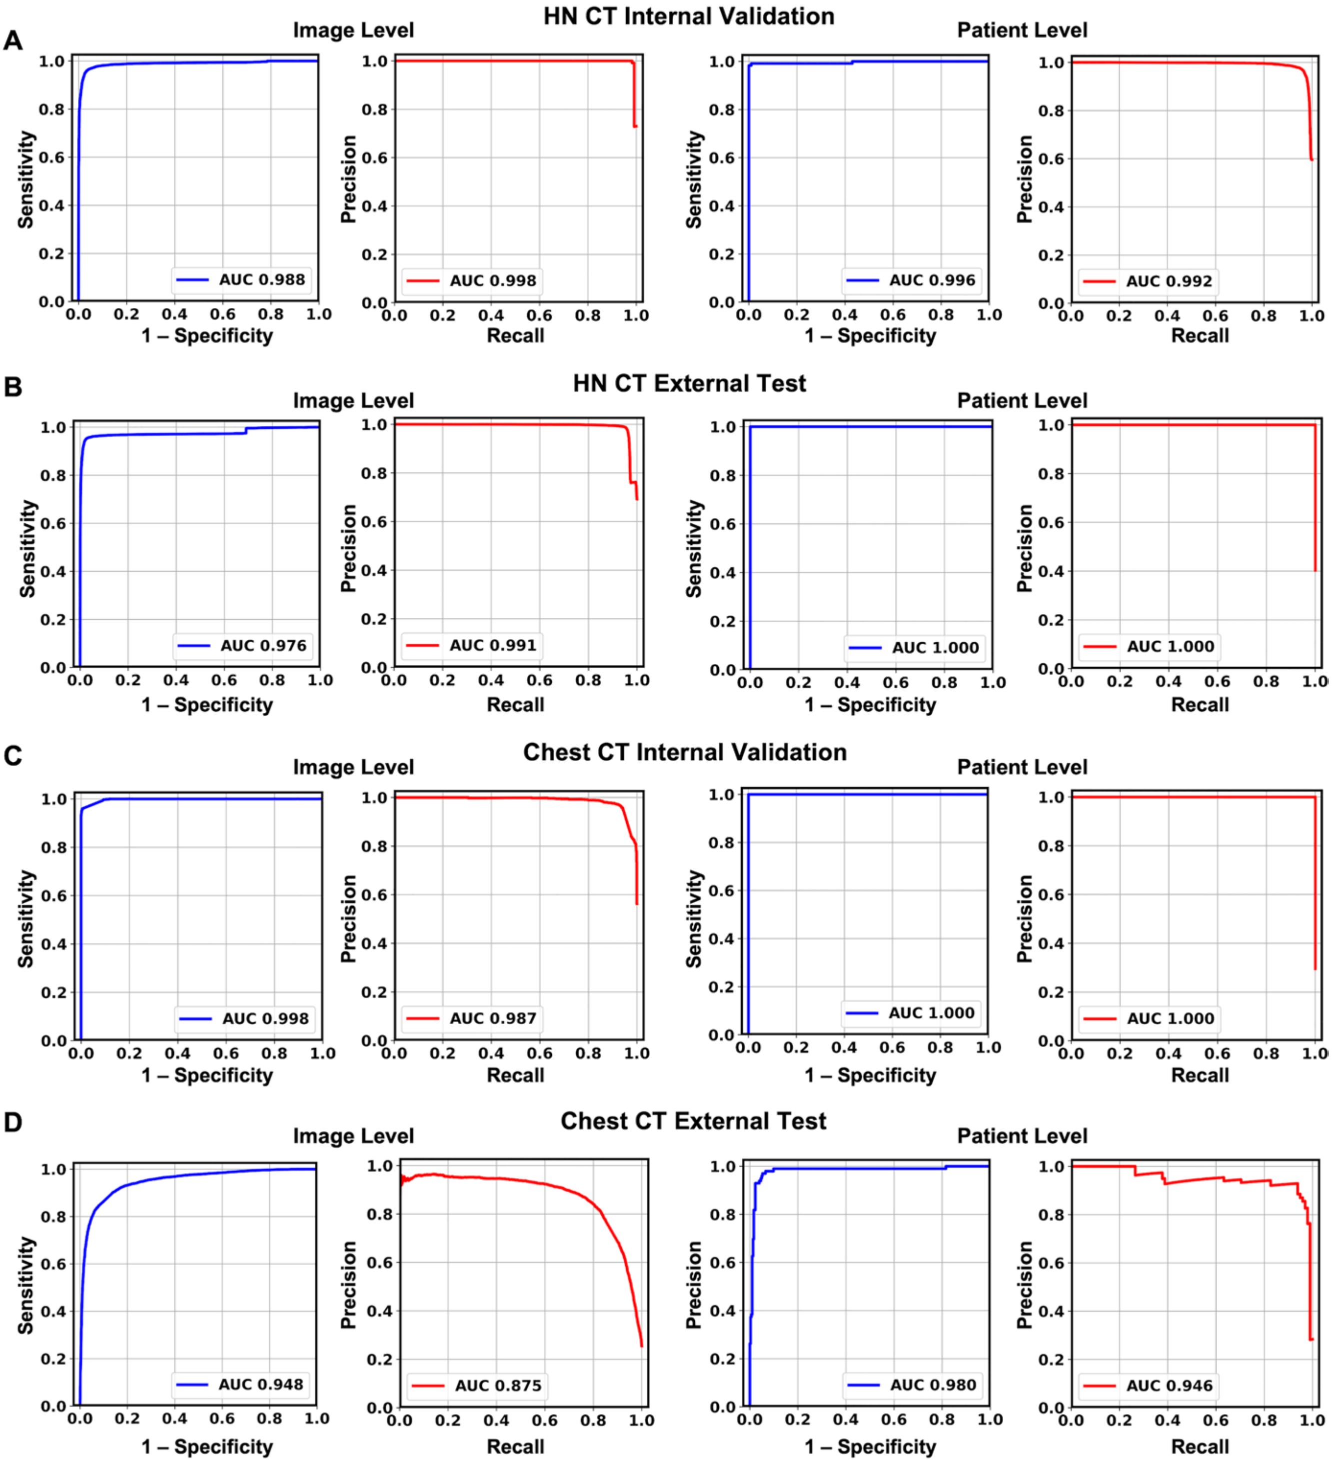

Figure 2: Receiver operating characteristic (ROC) curves and precision-recall (PR) curves calculated with the EfficientNetB4 model at both the image level and the patient level for **(A)** the head and neck (HN) cancer internal validation set (*n* = 216 patients and 33 264 images), **(B)** the HN cancer external test set (*n* = 595 patients and 39 270 images), **(C)** the lung cancer internal validation set (*n* = 53 patients and 3710 images), and **(D)** the lung cancer external test set (*n* = 402 patients and 28 140 images). All six ROC curves showed high areas under the curve (AUCs), indicating strong sensitivity and specificity in detecting these contrast enhancements at both the image level and the patient level. The PR curve of the lung CT external test set at the image level showed a slightly lower AUC than those of the other PR curves.

retained intravenous contrast material from prior outside hospital imaging during stroke workup, which would have lower sensitivity for acute stroke evaluation.

Limitations of this study include the possibility of uncaptured confounders within our datasets that vary from other institutions and our datasets being limited to HN scans and chest scans with a single phase of contrast enhancement and a known cancer diagnosis. We recommend that future users of the pipeline conduct small, local tests on their institutional scans prior to implementation at scale.

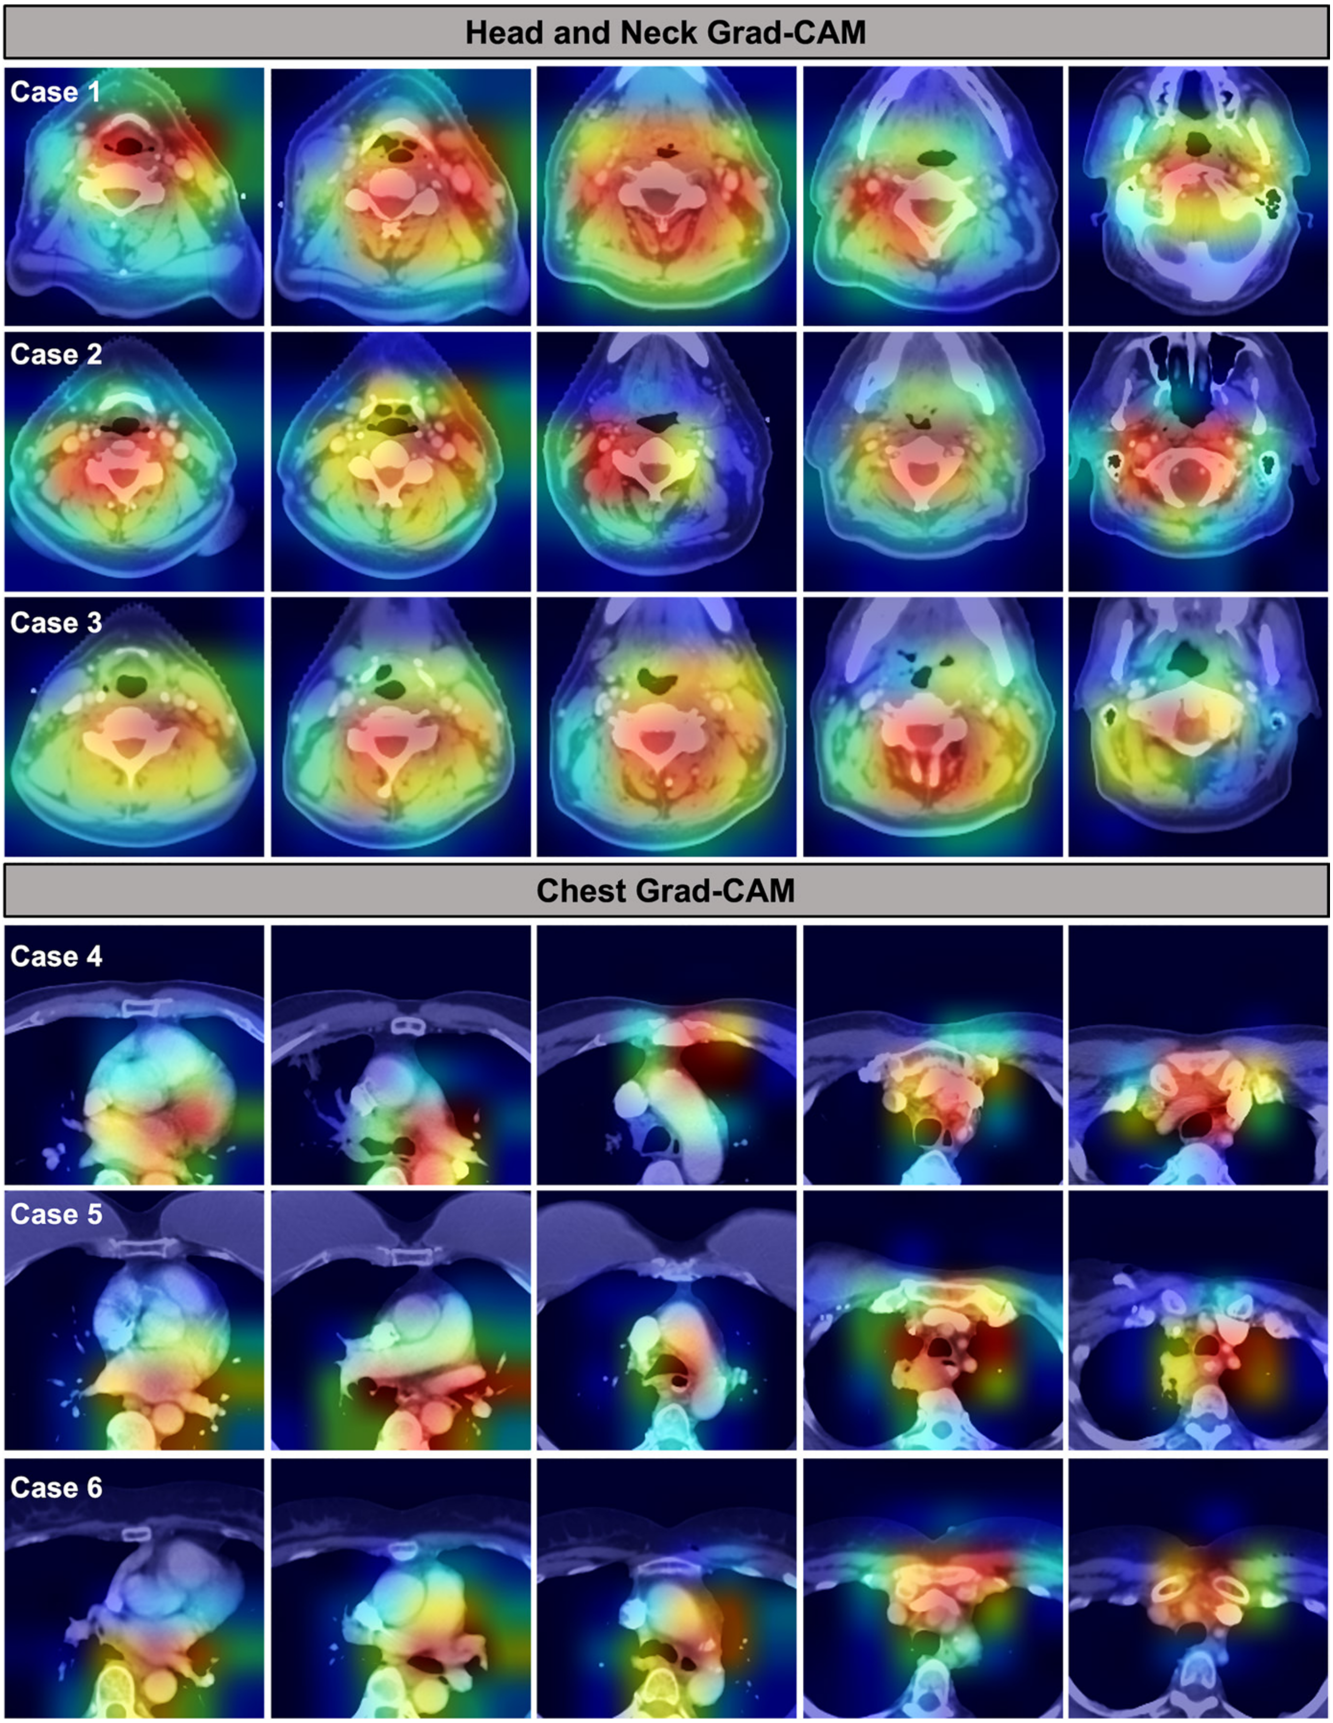

Figure 3: Gradient-based class activation maps (Grad-CAM) from the EfficientNetB4 model. Six representative scans from patients with head and neck cancer (cases 1–3) and patients with lung cancer (cases 4–6) with five different image sections shown. The last convolutional layer in the model was used for the generation of class activation maps. Test input images are shown with overlaid activation maps, in which red colors highlight regions with a higher contribution and blue colors represent areas with a lower weight value.

In conclusion, we developed—and made publicly available—a CNN-based deep learning model that accurately detects intravenous contrast enhancement on HN and chest CT scans across multiple institutions with nearly perfect performance, enabling scan-to-prediction automated contrast enhancement detection.

Acknowledgments: This article was prepared by using data from datasets (RTOG-0617) from the Data Archive of the National Cancer Institute's National Clinical Trials Network. Data were originally collected from a clinical trial (identifier NCT00533949; "High-Dose or Standard-Dose Radiation Therapy and Chemotherapy with or without Cetuximab in Treating Patients with Newly Diagnosed Stage III Non–Small Cell Lung Cancer That Cannot Be Removed by Surgery"). All analyses and conclusions in this article are the sole responsibility of the authors and do not necessarily reflect the opinions or views of the clinical trial investigators, the NCTN, or the NCI.

Author contributions: Guarantors of integrity of entire study, **Z.Y., J.L., Z.Z., B.H.K.**; study concepts/study design or data acquisition or data analysis/interpretation, all authors; manuscript drafting or manuscript revision for important intellectual content, all authors; approval of final version of submitted manuscript, all authors; agrees to ensure any questions related to the work are appropriately resolved, all authors; literature research, **Z.Y., B.H.K.**; clinical studies, **R.H.M., H.J.W.L.A., B.H.K.**; experimental studies, **Z.Y., J.M.Q., A.H., R.Z., H.J.W.L.A.;** statistical analysis, **Z.Y., J.L., Z.Z., H.J.W.L.A.**; and manuscript editing, **Z.Y., J.M.Q., A.H., R.Z., D.P., J.L., R.H.M., H.J.W.L.A., B.H.K.**

Disclosures of conflicts of interest: **Z.Y.** No relevant relationships. **J.M.Q.** No relevant relationships. **A.H.** Consultant for Altis Labs; shareholder in Altis Labs. **R.Z.** No relevant relationships. **D.P.** No relevant relationships. **J.L.** No relevant relationships. **Z.Z.** No relevant relationships. **R.H.M.** Contract/grant from ViewRay; consulting for ViewRay and AstraZeneca; payment for expert testimony from U.S. District Attorney's Office of New York. **H.J.W.L.A.** No relevant relationships. **B.H.K.** RSNA Research Scholar Award NIH K08DE030216.

References

1. Bi WL, Hosny A, Schabath MB, et al. Artificial intelligence in cancer imaging: clinical challenges and applications. CA Cancer J Clin 2019;69(2):127–157.
2. Willemink MJ, Koszek WA, Hardell C, et al. Preparing medical imaging data for machine learning. Radiology 2020;295(1):4–15.
3. Källman HE, Halsius E, Olsson M, Stenström M. DICOM metadata repository for technical information in digital medical images. Acta Oncol 2009;48(2):285–288.
4. Gueld MO, Kohnen M, Keysers D, et al. Quality of DICOM header information for image categorization. In: Siegel EL, Huang HK, eds. Proceedings of SPIE: medical imaging 2002—PACS and integrated medical information systems: design and evaluation. Vol 4685. Bellingham, Wash: International Society for Optics and Photonics, 2002.
5. Criminisi A, Juluru K, Pathak S. A discriminative-generative model for detecting intravenous contrast in CT images. In: Fichtinger G, Martel A, Peters T, eds. Medical image computing and computer-assisted intervention – MICCAI 2011. Lecture notes in computer science. Vol 6893. Berlin, Germany: Springer, 2011;49–57.
6. Sofka M, Wu D, Sühling M, et al. Automatic contrast phase estimation in CT volumes. In: Fichtinger G, Martel A, Peters T, eds. Medical image computing and computer-assisted intervention – MICCAI 2011. Lecture notes in computer science. Vol 6893. Berlin, Germany: Springer, 2011; 166–174.
7. Lartaud PJ, Rouchaud A, Rouet JM, Nempont O, Boussel L. Spectral CT based training dataset generation and augmentation for conventional CT vascular segmentation. In: Shen D, Liu T, Peters TM, et al, eds. Medical image computing and computer assisted intervention – MICCAI 2019. Lecture notes in computer science. Vol 11765. Cham, Switzerland: Springer, 2019; 768–775.
8. Kakino R, Nakamura M, Mitsuyoshi T, et al. Comparison of radiomic features in diagnostic CT images with and without contrast enhancement in the delayed phase for NSCLC patients. Phys Med 2020;69:176–182.
9. He L, Huang Y, Ma Z, Liang C, Liang C, Liu Z. Effects of contrast-enhancement, reconstruction slice thickness and convolution kernel on the diagnostic performance of radiomics signature in solitary pulmonary nodule. Sci Rep 2016;6:34921.
10. Hosny A, Parmar C, Quackenbush J, Schwartz LH, Aerts HJWL. Artificial intelligence in radiology. Nat Rev Cancer 2018;18(8):500–510.
11. Tang Y, Lee HH, Xu Y, et al. Contrast phase classification with a generative adversarial network. In: Išgum I, Landman BA, eds. Proceedings of SPIE: medical imaging 2020—image processing. Vol 11313. Bellingham, Wash: International Society for Optics and Photonics, 2020; 1131310.
12. Philbrick KA, Yoshida K, Inoue D, et al. What does deep learning see? insights from a classifier trained to predict contrast enhancement phase from CT images. AJR Am J Roentgenol 2018;211(6):1184–1193.
13. Kwan JYY, Su J, Huang SH, et al. Data from radiomic biomarkers to refine risk models for distant metastasis in oropharyngeal carcinoma. The Cancer Imaging Archive. . Published 2019. Accessed June 30, 2021.
14. Vallières M, Kay-Rivest E, Perrin LJ, et al. Data from head-neck-PET-CT. The Cancer Imaging Archive. Published 2017. Accessed June 30, 2021.
15. Grossberg A, Elhalawani H, Mohamed A, et al. 2020 HNSCC [dataset]. The Cancer Imaging Archive. Published 2019. Accessed June 30, 2021.
16. Bradley J, Forster K. Data from NSCLC-Cetuximab. The Cancer Imaging Archive. Published 2018. Accessed June 30, 2021.
17. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Piscataway, NJ: IEEE, 2016; 770–778.
18. Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z. Rethinking the inception architecture for computer vision. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Piscataway, NJ: IEEE, 2016; 2818–2826.
19. Tan M, Le Q. EfficientNet: rethinking model scaling for convolutional neural networks. 36th International Conference on Machine Learning, ICML 2019, Long Beach, Calif, June 9–15, 2019.
20. Deng J, Dong W, Socher R, Li LJ, Li K, Li F. ImageNet: a large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. Piscataway, NJ: IEEE, 2009; 248–255.