The Future of Artificial Intelligence in Clinical Radiology: Savior or False Hope?

Greg Zaharchuk and Janet Mei

AJNR Am J Neuroradiol 2024, 45 (12) 1838-1844
doi: https://doi.org/10.3174/ajnr.A8550
http://www.ajnr.org/content/45/12/1838

The Future of Artificial Intelligence in Clinical Radiology: Savior or False Hope?

Greg Zaharchuk, and Janet Mei

Making predictions is hard, especially about the future.

Yogi Berra

Artificial intelligence (AI), defined here as deep learning and transformer-based computer vision and language models, has been proposed as a revolutionary technology for radiology, possibly more impactful than past revolutions such as new imaging modalities and/or the adoption of PACS. However, some suggest that AI is overhyped and may not meaningfully address the core challenges facing radiology. Medicine is an inherently conservative field, and rapid technology adoption is the exception rather than the norm. In other specialties, changes are often triggered by landmark clinical trials or new guidelines. In radiology, in which much of our work occurs upstream of clinical decision-making, linking practice changes to improved patient outcomes is difficult. Without clear evidence of clinical impact, financial incentives will likely drive technology adoption. This is already visible, where products that have made inroads are largely those have led to improved efficiency or cost-effectiveness. Conversely, adoption has been slower for technologies promising improved diagnostic accuracy. In this editorial, what the next decade may hold for AI in radiology will be explored, drawing on current applications while also exploring the history of similar medical technologies to better understand the potential forms that AI might take and how it will be integrated into practice.

INTRODUCTION

Radiology is primed for disruption. Past models, both private and academic, are struggling to keep pace with the surge in imaging studies, particularly because there has been no commensurate increase in the number of radiologists.[1] Likewise, this increased demand is straining noninterpretive parts of our field, leading to longer service hours, without enough scanners or technologists to run them. Simultaneously, declining reimbursement rates per study and rising operational costs highlight the need for AI solutions that could help radiology departments manage higher volumes cost-effectively. Clearly, the current system is unsustainable.

Will AI be our savior? Perhaps AI will soon become good enough to autonomously read studies without human supervision, or maybe, more realistically, it could provide a “first draft” of the report much like a good trainee, complete with relevant measurements and comparisons with prior studies. Short of this, it could optimize our reading environment, making us more efficient and accurate by standardizing the readout process and reducing distractions. Finally, AI-driven image acceleration and other upstream applications could improve operational efficiency and alleviate financial pressures. The following will explore each of these areas.

AI for Interpretation

Even though only about 20% of the reimbursement for radiology departments comes from the interpretive “professional” fee, it is front and center for practicing radiologists. This 20–80 reimbursement split may understate the value of AI for interpretive work, because there is evidence that the cost (as opposed to the reimbursement) within a radiology practice is actually more equally split between professional and technical areas.[2] Assuming that the number of radiologists is largely capped in the next 10 years (though ideas such as importing foreign medical graduates and using midlevel providers might alter this calculus), how could AI reduce the time spent per study without decreasing quality? Three options are possible (Figure 1):

1) Doctor AI: Read some studies completely autonomously (normal CTA head and neck studies in patients with low-quality indications anyone?).
2) Trainee AI: Predraft reports, including measurement of lesion sizes and comparison with prior studies. A radiologist would then confirm or edit the reports similar to the traditional academic model.
3) House Elf AI: No diagnostic capabilities but AI helps streamline the process of moving from images to reports: better study organization (imagine hanging protocols that work), less need to call/interact with technologists, quality assessment and control, and maybe even executive summaries of cases with links to key images. This improvement could also include “stream of consciousness” reporting. Such AI would reduce hidden costs associated with interpretation.

Note that none of these scenarios are what is currently the rage: one-off AI detection and triage algorithms for individual diseases like intracranial hemorrhage and cervical spine fracture. There are literally hundreds of these one-off implementations, raising questions about whether this is the right approach long term. While the computer vision technology to perform these tasks is remarkable, multiple studies in real-life settings suggest that it does not reduce interpretation time or improve outcomes for neuroradiologic applications.[3-6] Thus, they will remain niche products for early adopters, though they could become building blocks of larger systems described above.

Doctor AI

The most ambitious of these scenarios is autonomous AI, capable of generating complete diagnostic reports from imaging. This possibility is only on the table due to recent explosive improvements in performances of so-called “foundational models,” based on large-vision language models, many of which use self-supervised contrastive learning. Current performance for image to report based AI has mostly focused on the areas where we have the most data, such as 2D imaging like chest X-rays.[7] More

http://dx.doi.org/10.3174/ajnr.A8550

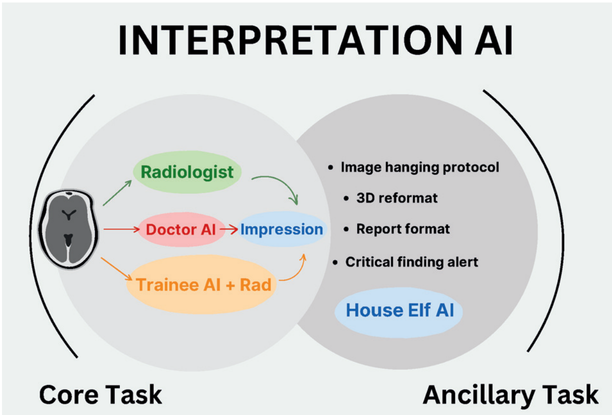

FIG 1. Varying levels of AI integration in radiologic workflows, ranging from full autonomy to efficiency optimization. In the Doctor AI model, AI autonomously generates complete reports with minimal-to-no need for physician input. The Trainee AI model involves AI-assisted prepopulation of reports but with more physician oversight, like an attending overseeing a resident or fellow. The House Elf AI model does not alter the physician's interpretation process but enhances the diagnostic environment, boosting radiologists' productivity through AI-driven efficiency improvements and removing distractions. It helps with better study organization, report structuring and formatting, less need to call/interact with technologists, quality assessment and control, and automatic critical finding alerts. Such AI would reduce hidden costs and wasted time associated with interpretation. Rad indicates radiologist.

recently, several groups have sought to extend these models to 3D imaging modalities, such as abdominal or head CT.[8,9] While impressive progress is clearly being made, the performance remains suboptimal. The examples shown in preprints demonstrate clear errors on relatively simple tasks and overall performance that is well below anything remotely useful.

However, we wouldn't bet against it improving given the rapid pace of improvement, computer vision detection of cats and dogs has quickly given way to conversational AI that can provide image captions or create movies from short text prompts. For this task, there are semi-infinite amounts of training data: all the paired imaging examinations and reports we have created during last 20 or so years. Incredible amounts of money are being spent by early-stage start-ups as well as large incumbents who believe this task is feasible. It is likely that the winners here will be institutions or companies that can combine extremely large, diverse data sets with massive compute budgets. These entities will need to incorporate radiologist domain expertise to navigate all the details that radiologists learn in residency. Another current unsolved problem is how to handle context, from comparison with priors to the integration of information from the electronic health records that is known to be essential for radiologists to provide value.

The biggest challenges here will be determining and then achieving adequate performance (what exactly is the error rate for radiologists anyway?); overcoming the (probably high) regulatory hurdles, given the potential for ultimately reducing the number of radiologists needed (expect your friendly professional association to weigh in on this); the acceptance of such technology by patients;[10] and potential medicolegal concerns (who is on the hook for mistakes?). Perhaps we can find subsets of examinations that might represent initial implementations, such as noncontrast head CTs without priors, or perhaps AI could be trained not to detect individual abnormalities but rather to detect the absence of abnormalities, a so-called "normal detector," and remove them from the worklist. This latter approach is particularly attractive because the rise in the number of imaging studies is being driven largely by low-quality indications, resulting in more studies with normal findings filling our lists.[11-13] I sometimes feel transported back to my days in residency during my screening mammography rotation when I sit down for the evening emergency department shift. Perhaps one day, we will remark on those "old days" when we had to actually read the studies with normal findings, much like some of us talk about using alternators, to be greeted by blank stares.

It may be instructive to look back on a technology that promised such automation. In the 1990s, interpretations of Pap smears by cytologists had all the elements we see in today's radiology: too many studies and too few people to read cases.[14] Computer-based algorithms had reasonable performance. However, the autonomous reading approach was ultimately not successful, given the reluctance of the FDA to approve what they considered a "high-risk" device. Thus, the first implementations were for quality control applications, looking at cases with negative findings with the hope of reducing false-negatives. The most successful paradigm turned out to be one in which cases deemed to have normal findings by the machine were flagged for the cytologist, who then had to review only the 20 "most abnormal" high-power fields flagged by the AI rather than examine the whole slide, thus resulting in time-savings. This improvement leads to our next category of interpretative AI.

Trainee AI

A recent survey showed that only 21% of patients are comfortable with AI providing radiology reports, suggesting the need for alternatives.[10] Thus, if truly autonomous AI, even if possible, is not deemed palatable by regulators, patients, and payors, this model could be scaled back to the idea of providing a draft, which would need to be checked by the radiologist before sign off. This approach is similar to that seen in academic practices, in which residents and fellows review images, interrogate the medical record, pull up relevant priors, and ultimately draft a report that is then edited in response to the attending's feedback. This approach is an attractive example of "human in the loop" AI, in which the radiologist is still liable ultimately for the read, making clear the medicolegal responsibility.[15] It may also be more palatable to regulators and professional societies, though clearly there might be worries about "mills" that just sign off on whatever these models produce. Such models might also be used to address the most common or laborious examinations; however, as for the Doctor AI implementation above, figuring out how to give AI the proper context is still likely to be a challenging undertaking. Of course, for maximal efficiency, both approaches need to be fully integrated into the radiologist's reading environment, interfacing with the electronic medical record and dictation software.

Such an approach will need to prove that its use leads to equivalent outcomes with less cost. This lower cost could be due to faster per-study reading by the radiologist, though little data exist about AI helping speed up interpretation. Concerns exist as to whether radiologists are capable or willing to overrule AI results when they are wrong.[16] Others have suggested combining these draft reports with oversight by nonradiologist midlevel providers. While this approach is attractive because there remains a human in the loop who can step in and rectify obvious mistakes, studies suggest that this may be the worst of both worlds, because individuals with less training are less willing to contradict erroneous AI outputs. Thus, the combination of trainee AI and midlevel providers, while attractive to payors, might not benefit patients.

House Elf AI

Further down on the omniscient scale would be House Elf AI. This recognizes all the inefficiencies in the current system and uses AI to mitigate them, thereby making the radiologist's experience better and allowing faster and more accurate reads. The overall goal is to maximize image-interpretation time and minimize all the other distractors. This could be considered as a radiologist's "personal assistant" who can just fix everything so that when he or she sits down to read cases, everything is perfectly ready for efficiency. The advantages of this approach are many. As stated by the American economist W. Edwards Deming, "Uncontrolled variation is the enemy of quality." As imaging groups get larger, the amount of complexity increases exponentially, exacerbating the problems of variability. Finally, such an AI approach would not have a heavy diagnostic component, mitigating fears of radiologist de-skilling.[17] It should be straightforward to deploy and require minimal regulatory burden or societal resistance.

It is difficult to estimate how much inefficiency there is in the current system, but it is likely high. Radiologists are currently asked to do many noninterpretive tasks as part of their daily job. The first and most obvious is the time spent wrangling with report dictation software. Even under the best of circumstances, time spent in report creation takes away from time interpreting images. How much time do we waste daily looking away from images to find where to insert text and to correct misspellings and faulty capitalization? This is not a good use of a radiologist's time. Some radiologists have begun to use "scribes" who handle this aspect of report creation.[18] This use can lead to noticeable efficiencies, similar to medical scribes that have been used for clinical visits. AI has already been implemented to listen to patient-doctor conversations and transcribe first drafts of clinical notes that doctors can then lightly edit.[19] Implementation of this use in the radiology space would be a game changer for productivity.

Beyond acting as a scribe, AI could finally allow us to use hanging protocols, a dream of radiologists for decades. Hanging protocols enhance efficiency by standardizing the PACS layout of imaging series and prior comparisons and can lead to increased efficiency.[20] As everyone knows, these never seem to work, which is largely due to the variability in the naming conventions that different scanners use or that are flummoxed by simple things like series name changes or repeat studies. While PACS companies continue to promise that they will work "in the next version," we need to fix the problem of variability at the root source. AI can be used to evaluate the DICOM header with natural language processing and computer vision of the pixel data to classify sequences into standardized names, regardless of the names the technologists use at the scanner.[21] Such technology can also "grade" the quality of sequences, evaluating the presence of motion, the correct extent of organ coverage, and the SNR. One time saver here is that AI can determine the best of multiple motion-degraded studies and send only that one to the PACS for interpretation, saving PACS storage space and radiologist's interpretation time, reducing the frustration and risks associated with evaluating suboptimal motion-degraded images.

Finally, as we move to more 3D imaging, we can envision software that automatically performs reformats in standardized imaging planes (Fig 2). Once implemented, this change will result in serial imaging studies that will always be aligned, leading to simple comparisons of lesions and their sizes without the need for time-consuming reformatting by the radiologist. This improvement would also free the technologist from doing the reformats and dealing with radiologists' phone calls when the technologist makes a mistake. This time could be used for additional patient care or even improved throughput for certain imaging studies that are reformat-heavy, such as CT.

The largest challenge to introducing such AI into operations is how to estimate its value. Radiologists and technologists see the value immediately, but administrators may not be familiar with all the hidden costs that are being borne by the organization as variability increases. Studies to demonstrate improved efficiency by technologists (ie, the ability to scan more patients) by freeing them from the need for reformats and attention to potential errors could be performed. For radiologists, even more than accuracy (which should increase), demonstration that they can read more studies per unit time is essential. It is possible that AI-driven software could enable 3D labs at sites that cannot currently staff them or allow them to operate in a timely fashion with fewer full-time employees, including off-hours and holidays. Removing these easily automatable tasks from both groups will reduce the radiologist's and technologist's burnout and help with staff retention. Last, in the age of big data, removing variability should finally enable imaging data to be optimally integrated into downstream processes. These will include internal uses, subsuming the need for smart DICOM routers to send the right sequences to third-party AI applications. Benefits would accrue outside the radiology department, enabling the easier consumption of images into planning software for other departments or into aggregated data sets for research. There is specific synergy here with integrating images into vision language models. Despite imaging being 90% of data by storage volume, the extreme variability makes it difficult to integrate into predictive models due to the infeasibility of manual data organization at scale. Such AI tools might, in fact, accelerate our movement toward the more diagnostically oriented AI in the Doctor and Trainee models.

The Back Office

While all these applications described above sound great, the big question remains: How will departments afford them? While we have described how return-on-investment (ROI) could be measured for these, most of the reimbursement in our field comes from the "technical component," which is meant to support the

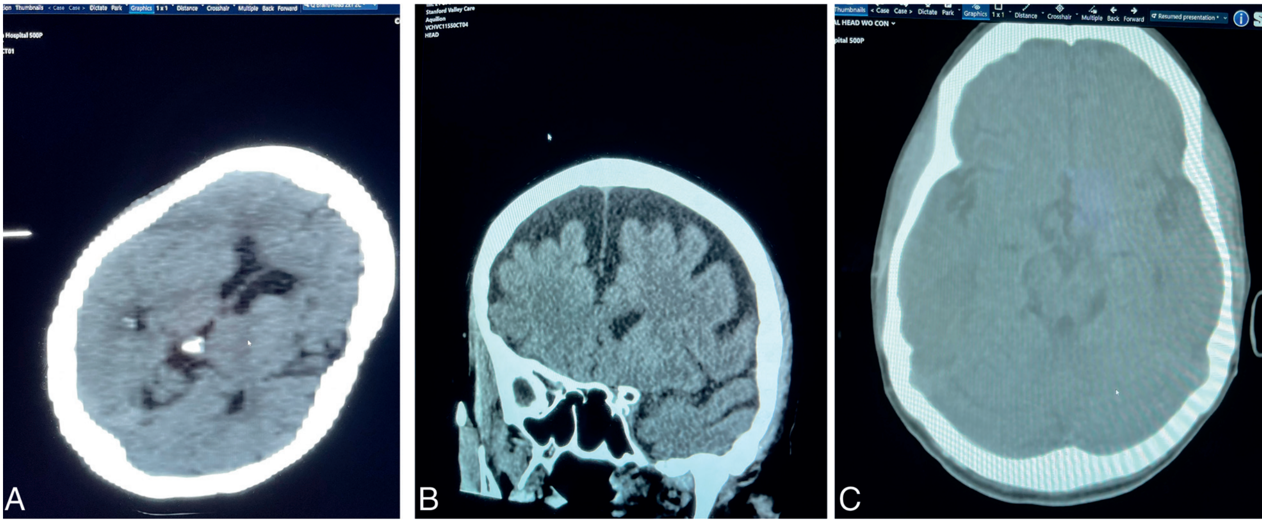

FIG 2. Variability in quality among human-generated 3D reformatted images from noncontrast head CTs, highlighting the potential benefits of automation. *A*, Stairstep artifacts caused by reformats from thicker slab imaging. *B*, Poor alignment of a reformat intended to be coronal. *C*, Inadvertent exclusion of the anterior portion of the head after reformatting.

infrastructure and costs incurred performing an imaging study, typically about 80% of total reimbursement. Thus, AI that enhances scanning efficiency without adding resources has become one of the early AI success stories in radiology—a trend expected to continue and accelerate.

Several companies have successfully introduced AI tools for accelerating image acquisition and reducing the radiation dose. These innovations were initially pioneered by independent companies and later adopted by larger scanner vendors, yielding 20%–50% reductions in scan times without compromising diagnostic quality.[22,23] To date, millions of scans have been expedited using these methods, particularly in MRI and PET, two of the most time-consuming and costly imaging modalities. However, we believe we are at the beginning of a very long road on what is possible here. The combination of low-field MRI with AI image-quality improvement is largely unexplored. A mid- or low-field MRI scanner, with improved performance around metallic implants due to reduced susceptibility artifacts, could potentially achieve image quality comparable with that of a 3T scanner with the help of AI. Who would not want such a system? In addition to reduced cost and sitting requirements, lower electricity costs also make it more environmentally friendly.[24]

AI also has the potential to significantly cut down repeat scans due to patient motion (Fig 3*B*), which currently accounts for about 15% of all MRI scans.[25] Synthetic MRI can help complete examinations for patients who cannot tolerate the full scan (Fig 3*C*), avoiding the need to reschedule or resort to sedation.[26] Most important, we will need the "House Elf AI" described above to decide when and how to apply such capabilities, because the radiologist does not have time to oversee them.

All these solutions clearly address current problems, and the ROI for potential buyers is easy to calculate. It is also attractive to patients who prefer shorter, lower-radiation-dose scans and could yield a competitive advantage in a saturated radiology market. Nevertheless, smaller groups with fewer information technology (IT) resources still struggle to implement such solutions. The other headwind here is the fear by some radiologists that these methods might "hallucinate" (add new structures or remove existing ones) otherwise not be able to provide quality care; so far, these are theoretic risks that have not been seen under typical operating conditions.[27]

AI Applications No One Really Wants but That We Will Likely Get Anyway

The overlap of what is doable and what increases revenue means that we are going to see some applications that may not be what patients or radiologists really need. One example is software that scans radiologists' reports to determine whether downstream studies recommended by radiologists are being performed, and if not, it contacts the patients and providers to follow up. This software sounds good,[28] but there have not been randomized trials to determine whether such approaches improve patient care. Instead, much of the discussion is around avoiding potential litigation. Many recommendations made by radiologists are not evidence-based and are likely to drive up costs without providing benefit. This issue has been evaluated with respect to the recent practice of reporting small thyroid nodules and recommending follow-up ultrasound, an approach no longer recommended because it could not be shown to have benefit.[29] There might be niche indications in which such recommendations are of value, such as those that incorporate well-vetted guidelines for screening. Opportunistic AI screening might fall into this category, which, if used sparingly and wisely, could allow radiology to help our referring colleagues in truly unique ways, by starting conversations about occult diseases and risk factors,[30] though evidence of this value should precede implementation.

Will hospitals view these applications as "money makers" and justify them to avoid patients "falling through the cracks" of the system? Realistically though, we question the value and ethics of alerting already alert-fatigued physicians about a 1-mm questionable aneurysm to be followed up yearly on an outpatient basis (especially if the patient has really good insurance).

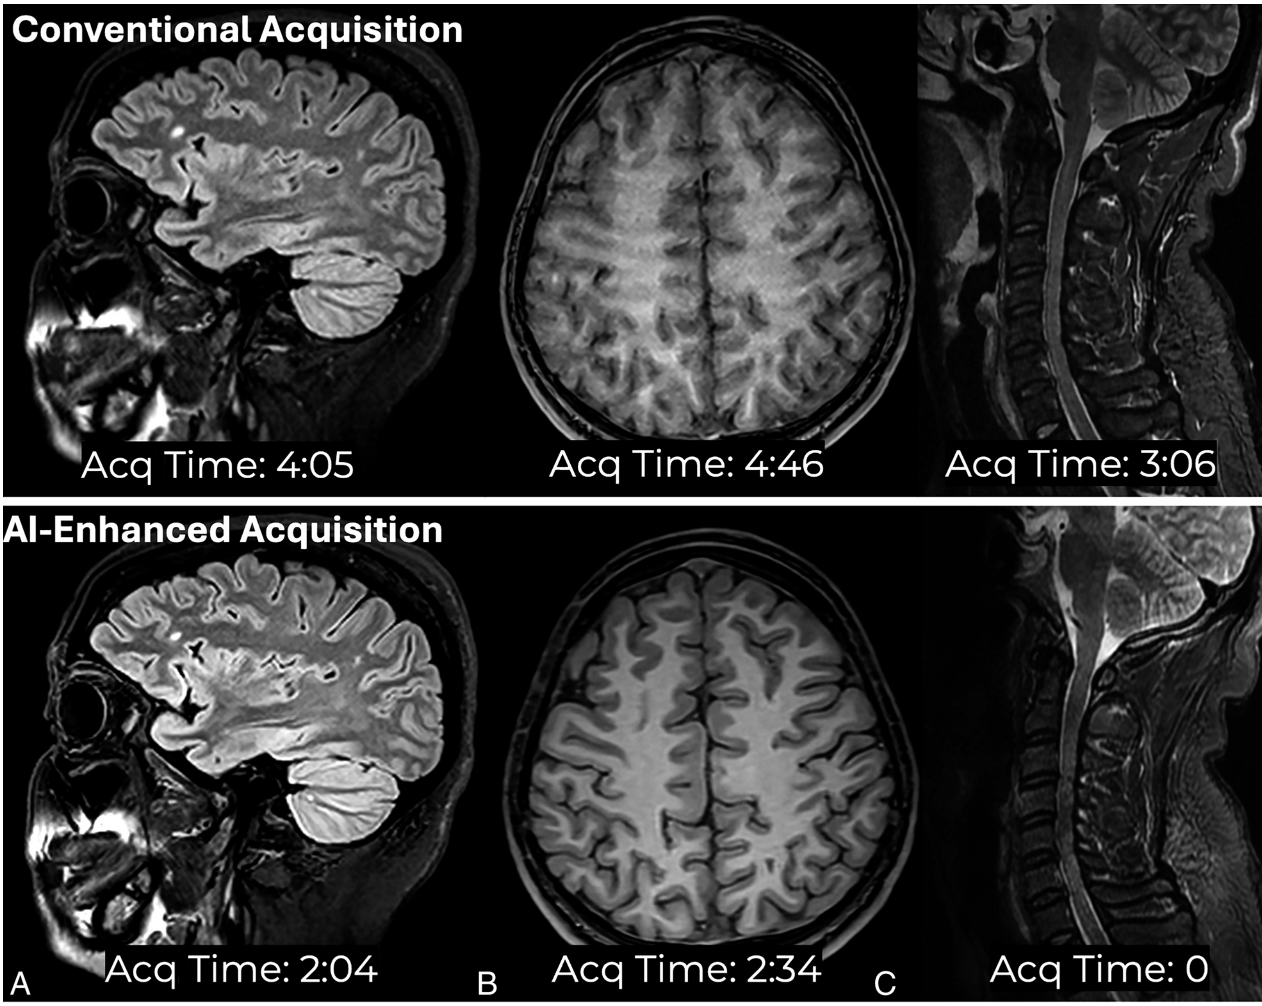

FIG 3. Examples of AI to accelerate MRI acquisition: noise reduction (*A*), super-resolution (*B*), and image synthesis (*C*). Note in *B* that the shorter scan also avoided the patient motion evident in the longer scan. Acq indicates acquisition.

While we do not totally exclude the possibility that there could be some value here, we suspect most will just increase medical costs and inconvenience the patient and their providers. But knowing how our system works, we wouldn't bet against such applications being part of our future.

Another application that has gotten some traction is the idea of creating "patient-friendly" reports with AI. This leveling of language and removal of jargon is a strong suit of current large language models, such as GPT-4 (https://openai.com/index/gpt-4/).[31] Currently, radiology reports are meant to convey information between physicians. However, some hospitals are driven by financial models that reward patient satisfaction and engagement, and providing patients with such targeted reports could improve these metrics. There is nothing inherently problematic about this practice, but we have larger problems to fix in the field, and if patients wish to understand and interact with their reports more deeply, there are many publicly available resources available. Studies demonstrating that such methods improve patient care would change my opinion, but until then, such an AI application is a distraction from more pressing problems even if it is an "easy win" for departments.

Implementation Woes: Is Radiology IT Up to the Job?

Predicting the future of AI in radiology must take into account the bandwidth and the number of clinical informatics personnel to implement and maintain these technologies. Given the need to support legacy systems, implement cybersecurity, and get all the legacy systems to work together, there is not much time for new projects in most IT departments. Adding AI software trials and long-term monitoring to their responsibilities poses a significant challenge.

If we are to make any significant progress during the next 10 years, departments and practices must invest strategically in their IT capabilities to make it happen. This problem is a chicken-and-egg one, in which the financial benefits will only accrue after significant financial resources are provided to implement AI algorithms. Otherwise, we will end up with a few choices from legacy vendors whose products may or may not satisfy radiologists' needs.

Closing Thoughts

There will be new AI applications in the next 10 years (Table 1). The sweet spot will be applications that improve patient care and radiologists' productivity with either net neutral or financially

Benefits and challenges for clinical AI integration in the next 10 years

Benefits

1) Improved radiologist efficiency: AI can automate repetitive tasks like image segmentation and preliminary reporting, allowing radiologists to focus on more complex cases. With an aging population and growing diagnostic needs, AI can help manage the rising volume of cases, ensuring timely diagnosis and care.
2) Shorter scan times and improved standardization: AI can optimize scan protocols, leading to quicker examinations with fewer artifacts and errors, reducing the need for repeat imaging. By standardizing image naming, orientation, and organization, display protocols can be used to reduce radiologists' interpretation times.
3) Cost reduction for health care payors: By automating image analysis and improving workflow efficiency, AI can reduce the overall cost of radiologic examinations and decrease the burden on health care systems.
4) Environmental impact: The use of AI to reduce scan time or create better image quality can reduce the per-person carbon footprint of imaging. AI-based augmentation of low-dose CT and MR contrast examinations can reduce downstream marine environment pollution.

Challenges

1) Demonstrating the value of interpretative AI: Clinical studies must focus on the reporting times of radiologists and demonstrate that AI applications really reduce the time per study and can increase the overall number of studies read by radiologists in a shift, without loss of quality.
2) Technical infrastructure: Skilled IT personnel are needed to deploy, test, and maintain AI systems, possibly posing a challenge, especially for smaller or underfunded institutions. These professionals are in high demand in most health care systems.
3) Regulatory issues: Navigating the evolving regulatory landscape for AI in health care (such as FDA approval) is complex and could slow down the adoption of AI tools. In-house-developed applications might not be FDA-approved but could increase medicolegal risks.
4) Integration into clinical workflow: All radiologists work differently. Incorporating AI into established workflows without disrupting the clinical routine or overburdening radiologists with additional tasks is challenging. Implementing AI so that it fits into the ecosystem of our legacy systems (scanners, routers, PACS) is not easy.
5) Clinician's trust: Radiologists may be hesitant to adopt AI, especially if they are unsure about its accuracy or reliability, making widespread adoption slower.

positive characteristics. The tsunami of imaging studies coupled with the limited number of radiologists will require some solution and where we end up on the Doctor AI/Trainee AI/House Elf AI spectrum will be determined by a range of factors including the capability of vision language models and the burden of regulations that will be imposed by external forces. if we had to bet, we suspect will land somewhere between the Doctor and Trainee AI paradigms, meanwhile taking the "easy wins" from the House Elf model. There will be much handwringing regarding the decline in radiologists' independent clinical skill levels, though this decline has been a constant part of our history as a field. There will be just too much work for us to expect an unchanging high level of individual acumen, though we will continue to value expert diagnosticians as we do now.

The ability of AI to impact and replace technologic cost with software should usher in new, cheaper devices that are easier to maintain and better for the environment, reducing technical component costs while improving the patient experience. There is no question that the task will be difficult, with plenty of issues surrounding these technologies, from data distribution drift to enshrinement of old biases. Pivotal clinical trials of new technologies will be challenging, especially given the pace of innovation. However overall, radiology has always embraced new technologies and found ways to innovate and improve patient care even with incomplete knowledge. This evolution, however, must be a 2-way street. Radiologists need to actively engage with AI developers, contributing their domain expertise and ethical judgment. Meanwhile, regulatory bodies need to trust radiologists to oversee and implement AI applications responsibly. Only then can we fully realize the potential of AI in radiology during the next decade.

ACKNOWLEDGMENT

The authors thank Laura Wilhelm for her illustrations, which enhanced this editorial.

REFERENCES

1. **Why scan-reading artificial intelligence is bad news for radiologists.** November 29, 2017. *The Economist*. https://www.economist.com/free-exchange/2017/11/29/why-scan-reading-artificial-intelligence-is-bad-news-for-radiologists.
2. Anzai Y, Heilbrun ME, Haas D, et al. **Dissecting costs of CT study: application of TDABC (time-driven activity-based costing) in a tertiary academic center.** *Acad Radiol* 2017;24:200–08 CrossRef Medline
3. Voter AF, Larson ME, Garrett JW, et al. **Diagnostic accuracy and failure mode analysis of a deep learning algorithm for the detection of cervical spine fractures.** *AJNR Am J Neuroradiol* 2021;42:1550–56 CrossRef Medline
4. Savage CH, Tanwar M, Elkassem AA, et al. **Prospective evaluation of artificial intelligence triage of intracranial hemorrhage on noncontrast head CT examinations.** *AJR Am J Roentgenol* 2024 Sep 4. [Epub ahead of print] CrossRef Medline
5. Arbabshirani MR, Fornwalt BK, Mongelluzzo GJ, et al. **Advanced machine learning in action: identification of intracranial hemorrhage on computed tomography scans of the head with clinical workflow integration.** *NPJ Digit Med* 2018;1:9 CrossRef Medline
6. Del Gaizo AJ, Osborne TF, Shahoumian T, et al. **Deep learning to detect intracranial hemorrhage in a national teleradiology program and the impact on interpretation time.** *Radiology Artif Intell* 2024;6:e240067 CrossRef Medline
7. Thawkar O, Shaker A, Mullappilly SS, et al. **XrayGPT: chest radiographs summarization using medical vision-language models.** *arXiv.org* 2023. https://arxiv.org/abs/2306.07971v1.
8. Yang L, Xu S, Sellergren A, et al. **Advancing multimodal medical capabilities of Gemini.** *arXiv.org* 2024. https://arxiv.org/abs/2405.03162v1.
9. Blankemeier L, Cohen JP, Kumar A, et al. **Merlin: a vision language foundation model for 3D computed tomography.** *arXiv.org* 2024. https://arxiv.org/abs/2406.06512v1.

10. Ney E, Berger E. **Patients are more comfortable with AI making diagnoses than answering the phone at their doctor's office, but overwhelmingly concerns persist**. *Survery Bain & Company* 2024. https://www.bain.com/about/media-center/press-releases/2024/patients-are-more-comfortable-with-ai-making-diagnoses-than-answering-the-phone-at-their-doctors-office-but-overwhelmingly-concerns-persist/.
11. Tu LH, Malhotra A, Venkatesh AK, et al. **Head and neck CTA utilization: analysis of ordering frequency and nonroutine results communication, with focus on the 50 most common emergency department clinical presentations.** *AJR Am J Roentgenol* 2022;218:544–51 CrossRef Medline
12. Tung M, Sharma R, Hinson JS, et al. **Factors associated with imaging overuse in the emergency department: a systematic review.** *Am J Emerg Med* 2018;36:301–09 CrossRef Medline
13. Mehan WA, Shin D, Buch K. **Effect of provider type on overutilization of CT angiograms of the head and neck for patients presenting to the emergency department with nonfocal neurologic symptoms.** *J Am Coll Radiol* 2024;21:890–95 CrossRef Medline
14. Lew M, Wilbur DC, Pantanowitz L. **Computational cytology: lessons learned from pap test computer-assisted screening.** *Acta Cytol* 2021;65:286–300 CrossRef Medline
15. American College of Radiology. **Thinking about AI.** 2024. https://www.acr.org/Practice-Management-Quality-Informatics/ACR-Bulletin/Articles/June-2021/Thinking-About-AI.
16. Bernstein MH, Atalay MK, Dibble EH, et al. **Can incorrect artificial intelligence (AI) results impact radiologists, and if so, what can we do about it? A multi-reader pilot study of lung cancer detection with chest radiography.** *Eur Radiol* 2023;33:8263–69 CrossRef Medline
17. Brady AP, Allen B, Chong J, et al. **Developing, Purchasing, Implementing and Monitoring AI Tools in Radiology: Practical Considerations—A Multi-Society Statement from the ACR, CAR, ESR, RANZCR and RSNA**. *Radiol Artif Intell* 2024;6:e230513 CrossRef Medline
18. Sami F. **TeleRadiology Scribing in Action.** YouTube 2020 Nov 6. https://www.youtube.com/watch?v=Uvc9opoZs7A.
19. Tierney AA, Gayre G, Hoberman B, et al. **Ambient artificial intelligence scribes to alleviate the burden of clinical documentation.** *N Engl J Med Catal* 2024;5 CrossRef
20. Geis JR. **Medical imaging informatics: how it improves radiology practice today.** *J Digit Imaging* 2007;20:99–104 CrossRef Medline
21. Arnold TC, Duffy B, Zaharchuk G, et al. **Harmonization of Longitudinal Multicenter Imaging Data Through Natural Language Processing and Brain Alignment Algorithms.** In: *Proceedings of the Annual Meeting of the American Society of Neuroradiology*. May 18–22, 2024. Las Vegas, Nevada; 466
22. Kim M, Kim HS, Kim HJ, et al. **Thin-slice pituitary MRI with deep learning-based reconstruction: diagnostic performance in a postoperative setting.** *Radiology* 2021;298:114–122 CrossRef Medline
23. Bash S, Wang L, Airriess C, et al. **Deep learning enables 60% accelerated volumetric brain MRI while preserving quantitative performance: a prospective, multicenter, multireader trial**. *AJNR Am J Neuroradiol* 2021;42:2130–37 CrossRef Medline
24. Heye T, Knoerl R, Wehrle T, et al. **The energy consumption of radiology: energy- and cost-saving opportunities for CT and MRI operation.** *Radiology* 2020;295:593–605 CrossRef Medline
25. Duffy BA, Zhao L, Sepehrband F, et al; Alzheimer's Disease Neuroimaging Initiative. **Retrospective motion artifact correction of structural MRI images using deep learning improves the quality of cortical surface reconstructions**. *Neuroimage* 2021;230:117756 CrossRef Medline
26. Tanenbaum LN, Bash SC, Zaharchuk G, et al. **Deep learning-generated synthetic MR imaging STIR spine images are superior in image quality and diagnostically equivalent to conventional STIR: a multicenter, multireader trial.** *AJNR Am J Neuroradiol* 2023;44:987–93 CrossRef Medline
27. Radmanesh A, Muckley MJ, Murrell T, et al. **Exploring the acceleration limits of deep learning variational network-based two-dimensional brain MRI.** *Radiol Artif Intell* 2022;4:e210313 CrossRef Medline
28. White T, Aronson MD, Sternberg SB, et al. **Analysis of radiology report recommendation characteristics and rate of recommended action performance.** *JAMA Netw Open* 2022;5:e2222549 CrossRef Medline
29. Hoang JK, Langer JE, Middleton WD, et al. **Managing incidental thyroid nodules detected on imaging: white paper of the ACR Incidental Thyroid Findings Committee.** *J Am Coll Radiol* 2015;12:143–50 CrossRef Medline
30. Pickhardt P, Summers R, Garrett J, et al. **Opportunistic screening: Radiology Scientific Expert Panel.** *Radiology* 2023;307:e222044 CrossRef Medline
31. Lyu Q, Tan J, Zapadka ME, et al. **Translating radiology reports into plain language using ChatGPT and GPT-4 with prompt learning: results, limitations, and potential.** *Vis Comput Ind Biomed Art* 2023;6:9 CrossRef Medline