

Artificial Intelligence Application in Nuclear Medicine

Pooja Sanjay Gore

Department of Computer, Nutan Maharashtra Institute of Engineering & Technology, Samarth Vidya Sankul, Vishnupuri, Talegaon Dabhade, Pune, Maharashtra 410507, India
Email: gorep6[at]gmail.com
Tel- 07984609375

Abstract: *In the field of medication, specifically, clinical imaging, the publicity of late years about man-made consciousness (AI) has had a critical effect. Despite the fact that news in the day to day press and clinical distributions about new capacities and accomplishments of AI is practically overpowering, for some mediators the term and the working of AI stay a "discovery," prompting overstated assumptions from one perspective and unwarranted feelings of dread on the other. Individuals as of now connect with AI in an assortment of routes in day to day existence - for instance, on cell phones, in the vehicle, or while riding the web - however frequently without really acknowledging it. Artificial intelligence likewise can possibly take on an assortment of basic or dreary undertakings in the medical services area sooner rather than later. Be that as it may, AI unquestionably won't make radiologists or atomic medication experts outdated as clinical specialists within a reasonable time-frame. As opposed to the disturbance evoked in certain media, a consistent change can be anticipated; this change doubtlessly will start or has started in the demonstrative disciplines, specifically, clinical imaging. According to the point of view of the radiologist or atomic medication trained professional, this turn of events, rather than being seen as a danger, should be visible as a valuable chance to play a spearheading job inside the medical care area and to effectively shape this change interaction. In this article, we endeavor to give a calculated order of AI, a short rundown of what we view as the main specialized basics, a conversation of potential applications in atomic medication and, at last, a concise thought of the conceivable effect of these innovations on the calling of the doctor.*

Keywords: Artificial intelligence, Machine learning, Nuclear medicine, Deep learning

1. How to Define AI

The term artificial intelligence first appeared in an application for a 6-wk workshop entitled Dartmouth Summer Research Project on Artificial Intelligence at Dartmouth College in Hanover, New Hampshire, and is often defined as “intelligence demonstrated by machines, in contrast to the natural intelligence displayed by humans and animals”. However, since its first appearance, the term has undergone a constant redefinition against the background of what is technically feasible. On the one hand, the definition per se is already vague, because the partial term intelligence is not itself well defined. On the other hand, it depends directly on human perception and evaluation, which change constantly. Only a few decades ago, chess computers were regarded as a classic example of AI, because a kind of “intelligence” was considered a prerequisite for the ability to master this game. With the exponential growth of performance in computer hardware, however, it was soon possible to program chess computers that played masterfully without developing an understanding of the game as human players do. In simple terms, a computer’s memory had stored such a large number of moves from archived chess games between professional human players that the computer could look up an equivalent in a historical game for almost every imaginable game situation and derive the next move from it. This procedure, although simplified here, did produce extremely successful chess computers, but their behavior was predictable in principle and lacked typical human qualities, such as strategy and creativity. This “explainability,” together with a certain wear and tear of the “wow effect,” finally led to the fact that chess computers are no longer regarded as examples of AI by most people today. An attempt to systematize the area of AI

leads to a multitude of different procedures which, only in their entirety, define the field of AI. From the 1950s to the 1980s, AI was strongly dominated by so-called symbolic reasoning, through which AI is implemented by rules engines, expert systems, or so-called knowledge graphs. What these methods have in common is that they model entities of the real world and their logical relationships in the form of symbols with which arithmetic operations can then be performed. The main advantages of these systems are, on the one hand, their often comparatively low demand on the computing capacity of a computer system and, on the other hand, their comprehensible behavior, with which every step of the system (data input, processing, and data output) can be reproduced and understood. The main disadvantage, however, is the necessary step of modeling, in which the part of the real world required for the concrete application domain has to be converted into symbols. This extremely labor-intensive task often has to be performed by people, so that the creation of such systems is mostly reserved for corporations or, recently, well-organized crowdsourcing movements. The larger problem in modeling, however, is that the performance and accuracy of such systems are bound a priori to the human understanding of the real world. Although this situation seems unproblematic in a game such as chess, with a manageable number of game pieces and their well-defined relationships to each other, for other applications (such as medicine), this situation results in considerable difficulties. Thus, many physicians are probably aware that even the most complex medical ontologies and classifications ultimately represent crude simplifications of the underlying biologic systems and do not fully describe the variability of diseases or their dependencies. Moreover, such classification systems can hardly keep pace with the medical knowledge gained in the

Volume 11 Issue 6, June 2022

www.ijsr.net

Licensed Under Creative Commons Attribution CC BY

digital age, a fact that inevitably limits symbolic AI systems based on such models. However, with the strongly increasing performance of computer hardware, nonsymbiotic AI systems increasingly came to the fore from the mid-1980s onward. What these systems have in common is that they are data driven and work statistically. These procedures are often summarized under the term machine learning in which computer systems learn to accomplish a task independently-that is, without explicit instructions-and thus perform observational learning from large amounts of data. The obvious advantage of these systems is that the time-consuming and limiting modeling phase is omitted, because the machine largely independently appropriates the internal abstraction of the respective problem and, assuming a sufficient and representative amount of example data, can also record and map its variability. In addition to the high demand for computing capacity during the training phase, these methods primarily have 2 disadvantages. On the one hand, there is a large to very large demand for example datasets during the training phase for almost all methods because, despite all technical advances, the abstraction of a problem is far less efficient than in the human brain. On the other hand, the internal representation of this abstraction in most of these systems is so complex that it can no longer be comprehended and understood by people, so that such systems are often referred to as “black boxes,” and the corresponding output of such systems can no longer be reliably predicted outside the set of tested input parameters. For complex and highly variable input parameters, such as medical image data, these systems thus can produce unexpected results and show a quasi-nondeterministic behavior; for example, an image of an elephant can be placed clearly visible into an image, and a state-of-the-art trained neural network either will most often not see it at all or will mistake it as other objects, such as a chair (2). In principle, machine learning procedures can be divided into supervised and unsupervised learning. In supervised learning, not only the input data but also the desired output data are given during the training phase, and the model learns to generate those outputs from the given inputs. To prevent the model from learning only the example data by memorization (also referred to as overfitting), various techniques are used; the central element is that only part of the data is presented to the model during training, and the performance of the model (i.e., the control of learning success) is measured against the other part of the data. In contrast, in unsupervised learning, the input data are given without any labels. The goal is then to understand the inherent structure in the data. Using clustering methods, for example, the observations to be analyzed are divided into subgroups according to certain features or feature combinations. Generative methods derive a probability distribution from sampled observations that can be used to generate synthetic observations. In the medical domain, in which the cost of labeling the data is high, semi supervised learning could be more useful. Here, only part of the data is labeled, and although the task is similar to supervised learning, the advantage is that the structure of the unlabeled data-which are often more abundant- can be exploited. Another form of classification is the division of the area of machine learning into conventional machine learning and deep learning. Conventional machine learning includes a large number of established methods, such as naive Bayes

classifiers, support vector machines, random forests, or even hidden Markov models, and has been used for years and decades in a wide variety of application areas, such as time series predictions, recommendation engines in e-commerce, spam filters, text translation, and many more. In recent years, however, the field of machine learning has been strongly influenced by deep learning, which is based on artificial neural networks (ANNs). Because of a multitude of layers (so-called hidden layers) between the input and output layers, these neural networks have a much larger space for free parameters and thus allow much more complex abstractions than conventional machine learning methods. An area of medical imaging currently receiving much attention, so-called radiomics, can be reduced to a 2-step process. In the first step, image data are converted by image processing methods into high-dimensional vectors (so-called feature vectors); from these vectors, predictive models-usually a classifier or a regressor-for deriving certain information from the same image data are then generated in the second step using conventional machine learning. Radiomics is currently being evaluated in a multitude of small, often retrospective studies, which often try to predict information such as histologic subtype, mutational status, or a response to a certain therapy from medical images of tumors. Because the first step requires careful feature engineering and strong domain expertise, there are already some attempts to replace the 2-step process in radiomics with deep learning by placing the image data directly into the input layer of an ANN without prior feature extraction. Because an article dedicated to radiomics also appears in this supplement to The Journal of Nuclear Medicine, we will not discuss radiomics further and will focus in particular on other applications of machine learning and deep learning.

Applications in Nuclear Medicine

The rise of AI in medicine is often associated with “superhuman” abilities and precision medicine. At the same time, often overlooked are the facts that large parts of physicians’ everyday work consist of routine tasks and that the delegation of those tasks to AI would give the human workforce more time for higher-value activities (3) that typically require human attributes such as creativity, cognitive insight, meaning, or empathy. The day-to-day work of medical imaging involves a multitude of activities, including the planning of examinations, the detection of pathologies and their quantification, and manual research for additional information in medical records and textbooks-which often tend to bore and demand too little intellectually from the experienced physician but, with continuously rising workloads, tend to overwhelm the beginner. Without diminishing the prospects of “superdiagnostics” and precision medicine, seemingly more easily achievable goals of AI in medicine should not be forgotten because they might relieve people who are highly educated and have specialized skills of repetitive routine tasks. A typical medical imaging work flow can be divided into 4 steps: planning, image acquisition, interpretation, and reporting. Steps such as admission and payment could be included as well. We have deliberately focused on the parts of the work flow in which the physician is directly and primarily involved. Next, we discuss existing or potential AI-based solutions clustered by that structure Planning Before an

examination is performed on a patient at all, whether a planned procedure is medically indicated should be determined. The more unpleasant, risky, or expensive the respective examination is, the more this guideline applies. For example, in the recruitment of amyloid-positive individuals for clinical trials. One of the greatest challenges in the scheduling of medical examinations is “no-shows”; this challenge is particularly problematic in the case of nuclear medicine because of tracer availability, decay, and cost. A highly relevant study from Massachusetts General Hospital demonstrated the feasibility of predicting no-shows in the medical imaging department using relatively simple machine learning algorithms and logistic regression. The authors included 54,652 patient appointments with scheduled radiology examinations in their study. Considering 16 data elements from the electronic medical record grouped by no-show history, appointment-specific factors, and sociodemographic factors, their model had a significant power to predict failure to attend a scheduled radiology examination (area under the curve [AUC], 0.75) (4). Given the recent technical improvements in deep learning, the relatively small number of included predictors in that study, and the recent availability of methods such as continuous (or incremental) learning, it is not far-fetched to hypothesize that the prediction of no-shows at a much higher accuracy could be available soon. Often patient-related information given at the time of referral is sparse, and extensive manual searching through large numbers of unstructured text documents by the physician is necessary to gather all of the information that is needed for optimal preparation and planning of the examination. Although the analysis of text documents may seem easy (compared with, e.g., image analysis) and recent advances in natural language processing and natural language understanding became very visible in gadgets such as Alexa (<https://alexa.amazon.com>), Google Assistant (<https://assistant.google.com>), or Siri (<https://www.apple.com/siri/>), such analysis in fact remains a particularly delicate task for machine learning. Still, the research community is making steady progress, structured reporting that allows straightforward algorithmic information extraction is gaining popularity, and data interoperability standards such as Fast Healthcare Interoperability Resources (FHIR) (<https://www.hl7.org/fhir/>) will gradually become available in clinical systems. Therefore, it can be assumed that, in the future, the time-consuming manual research of patient information will be performed by intelligent artificial assistants and presented to the physician in the form of concise casespecific dashboards. Such dashboards not only will aggregate relevant patient information but also likely will enrich this information by putting it into context. For example, a relatively simple rule-based symbolic AI could automatically check for certain contraindications, such as allergies, or reduce unnecessary duplication of examinations by analyzing prior examinations. Scanning Modern scanner technology already makes increasing use of machine learning, and recent advancements in research suggest considerable technical improvements in the near future (5). In nuclear medicine, attenuation maps and scatter correction remain hot topics for PET and SPECT imaging, so it is not surprising that these are the subjects of intensive research by various AI groups. Hwang et al. used a modified U-Net,

which is a specialized convolutional network architecture for biomedical image segmentation, to generate the attenuation maps for whole-body PET/ MRI (6). They used activity and attenuation maps estimated from the maximum-likelihood reconstruction of activity and attenuation algorithm as inputs to create a CT-derived attenuation map and compared this method with the Dixon-based 4-segment method. Compared with the CT-derived attenuation map, the U-Net-based approach achieved significantly higher agreement (Dice coefficient, 0.77 vs. 0.36). Instead of an analytic approach based on image segmentation, it is also possible to use generative adversarial networks (GANs) to directly translate 1 imaging modality into another. The feasibility of direct MR-to-CT image translation using context-aware GANs was demonstrated by Nie et al. in a small study involving 15 brain and 22 pelvic examinations. Another topic of research is the improvement of image quality. Hong et al. used a deep residual convolutional neural network (CNN) to enhance the image resolution and noise property of PET scanners with large pixelated crystals. Kim et al. showed that iterative PET reconstruction using a denoising CNN with local linear fitting improved image quality and was robust against noise-level disparities. Improvements in reconstructed image quality could also be translated to dose savings, as shown by multiple groups that estimated full-dose PET images from lowdose scans (i.e., reduction in applied radioactivity) using CNNs or GANs with favorable results. Obviously, this approach could also be translated to shorter acquisition times and result in higher patient throughput. In addition, improved image quality could also be translated to higher temporal resolution, as shown by Cui et al., who used stacked sparse autoencoders (unsupervised ANNs that learn a representation by training the network to ignore noise) to improve the quality of dynamic PET images. Berg and Cherry used CNNs to estimate time-offlight directly from the pair of digitized detector waveforms for a coincident event; this method improved timing resolution by 20% compared with leading-edge discrimination and 23% compared with constant fraction discrimination. In principle, this method could make the additional MRI scan obsolete. In nuclear medicine, scanning depends directly on the application of radiotracers, the development of which is a time-consuming and costly process. As in the pharmaceutical industry, the prediction of drug-target interactions (DTI) is an important part of this process in the radiopharmaceutical industry and has been performed with computer assistance for quite some time; AI-based methods are increasingly being used (6). For example, Wen et al. were able to predict the interactions between ziprasidone or clozapine and the 5-hydroxytryptamine receptor 1C (or 2C) or alprazolam and g-aminobutyric acid receptor subunit r-2 with a deep-belief network. Interpretation many interpreters maintain a list of examinations that they have to interpret and that they process chronologically in a first-in, firstout order. In reality, however, some studies have findings that require prompt action and therefore should be prioritized. Recently, a deep learning-based triage system that detects free gas, free fluid, or fat stranding in abdominal CTs was published (7), and multiple studies have already demonstrated the feasibility of detecting critical findings in head CT scans. In the future, such systems could work directly on raw data, such as sinograms, and raise alerts during the scan time,

even before reconstruction. In such a scenario, the technician could modify or extend the planned scan protocol to accommodate the unexpected finding; for example, an intracranial hemorrhage detected during a lowdose PET/CT scan could trigger an immediate full-dose CT scan of the head. However, the automatic detection of pathologies also offers other interesting possibilities beyond the prioritization of studies. For example, the processing of certain examinations, such as bone or thyroid scans, could be automated or at least accelerated with preliminary assessments, or an AI assistant working in the background could alert the human interpreter to possibly overlooked findings. Another, often disregarded possibility is that recurring secondary findings could be automatically detected and included in the report, freeing the human interpreter from an often annoying task. Many studies have already addressed the early detection of Alzheimer disease and mild cognitive impairment using deep learning. Ding et al. were able to show that a CNN with InceptionV3 architecture could make an Alzheimer disease diagnosis with 82% specificity at 100% sensitivity (AUC, 0.98) on average 75.8 mo before the final diagnosis based on 18F-FDG PET/CT scans and outperformed human interpreters (majority diagnosis of 5 interpreters). Similar network architecture was used by Kim et al. in the diagnosis of Parkinson disease from 123I-ioflupane SPECT scans; the test sensitivity was 96.3% at 66.7% specificity (AUC, 0.87). Li et al. used a 3-step process of automatic segmentation, feature extraction, and classification using support vector machines and random forests to automatically detect pancreas carcinomas on 18F-FDG PET/CT scans (4). On their test dataset of 80 scans, they found a sensitivity of 95.23% at a specificity of 97.51%. Perk et al. combined threshold-based detection with machine learning-based classification to automatically evaluate 18F-NaF PET/CT scans for bone metastases in patients with prostate cancer. A combination of statistically optimized regional thresholding and random forests resulted in a sensitivity of 88% at a specificity of 89% (AUC, 0.95). However, the ground truth in learning data originated from only 1 human interpreter, so that the performance of the machine learning approach must be evaluated with care. Interestingly, in a subset of patients who were evaluated by 3 additional nuclear medicine specialists, the machine learning classification performance was high when the ground truth originated from any of the 4 physicians (AUC range, 0.91–0.93), whereas the agreement between the physicians was only moderate (k , 0.53). That study (8) underlined the importance of reliable ground truth not only during validation but also during training when supervised learning is used. Nevertheless, it should not be forgotten that although existing systems sometimes provide excellent results with regard to the detection of 1 or more classes of pathologies, they still cannot generalize results as well as a human diagnostician. For this reason, human supervision remains absolutely mandatory in most scenarios.

Overall, however, the detection of pathologies during interpretation often accounts for only a small part of the total effort for the experienced interpreter. The increasing demand for quantification and segmentation usually involves much more effort, although these tasks are intellectually not very challenging and often are rather

ting. Therefore, the reasons for the wish to delegate these tasks to intelligent systems seem obvious. Rocca et al. used machine learning to estimate the arterial input function for the noninvasive full quantification of the regional cerebral metabolic rate for glucose in 18F-FDG PET (9). Instead of measuring the arterial input function during the scan with an invasive arterial blood sampling procedure, it was predicted with data from medical health records and dynamic PET imaging data. Before planned radiotherapy, it is necessary to precisely quantify the target structures by segmentation which, in the case of nasopharyngeal carcinomas, is often a particularly difficult and time-consuming activity because of the anatomic location. Zhao et al. showed, for a small group of 30 patients, that the automatic segmentation of such tumors on 18F-FDG PET/CT data was, in principle, possible using the U-Net architecture (mean Dice score of 87.47%). Other groups applied similar approaches to head and neck cancer and lung cancer (10). Still, fully automated tumor segmentation remains a challenge, probably because of the extremely diverse appearance of these diseases. Such an approach requires correspondingly large amounts of training data, for which the necessary ground truth in the form of segmentation masks usually has to be generated in a labor-intensive manual or semiautomatic task. Intelligent systems can also support the interpreter with classification and differential diagnosis. Many studies have shown possible applications for radiology, such as the differentiation of liver masses in MRI, bone tumor diagnosis in radiography, classification of interstitial lung diseases in CT, or diagnosis of acute infarct like myocarditis in MRI.

The training dataset was considerably large, including 780 samples of Graves' disease, 438 samples of Hashimoto disease, 810 samples of subacute thyroiditis, and 860 samples of normal cases. However, their validation strategy remains unclear, so the reported numbers must be evaluated with care (11). Reporting Medical imaging professionals are often confronted with referrer questions that, according to current knowledge and the state of the art, cannot be answered reliably or at all with the possibilities of imaging. In health care, AI is often intuitively associated with superhuman performance, so it is not surprising that there is such a high level of research activity in the area of prediction of unknown outcomes. Despite the high sensitivity and specificity of procedures such as PET/CT in tumor detection, it is still not possible to detect so-called micro metastases or early metastatic disease, although the detection of tumor spread has significant effects on the treatment concept. In an animal study of 28 rats injected with breast cancer cells, Ellmann et al. were able to predict later skeletal metastasis with an ANN based on 18F-FDG PET/CT and dynamic contrast-enhanced MRI data on day 10 after injection with an accuracy of 85.7% (AUC, 0.90). Future prospective studies will show whether these results can also be achieved in people, but the approach seems promising. Another group achieved promising results in the detection of micro metastases in lymph nodes in head and neck cancers by combining radiomics analysis of CT data and 3-dimensional CNN analysis of 18F-FDG PET data through evidential reasoning. Another important question in oncology—one that often cannot be answered with imaging—is the prediction of the response to therapy and overall survival. A small study by Xiong et al. of 30 patients with

esophageal cancer demonstrated the feasibility of predicting local disease control with chemoradiotherapy using radiomics features from 18F-FDG PET/CT and machine learning models. Milgrom et al. analyzed 18F-FDG PET scans of 251 patients with stage I or II Hodgkin lymphoma. They found that 5 features extracted from mediastinal sites were highly predictive of primary refractory disease when incorporated into a machine learning model. In a study conducted to predict overall survival in glioblastoma multiforme by integrating clinical, pathologic, semantic MRI-based, and O-(2-18F-fluoroethyl)-L-tyrosine PET/CT-derived information as well as treatment features into a machine learning model, PET/CT was not found to provide additional predictive power; however, the fraction of patients with available PET data was relatively low (68/189), and 2 different PET reconstruction methods were used (8). A study by Papp et al. included L-S-methyl-11C-methionine PET features, histopathologic features, and patient characteristics in a machine learning model to predict 36-mo survival in 70 patients with treatment-naïve gliomas; an AUC of up to 0.9 was achieved. Ingrisich et al. tried to predict the outcome of 90Y radioembolization in patients with intrahepatic tumors from pretherapeutic baseline parameters (10). They trained a random survival forest with baseline levels of cholinesterase and bilirubin, type of primary tumor, age at radioembolization, hepatic tumor burden, presence of extrahepatic disease, and sex. Their model achieved a moderate predictive power, with a concordance index of 0.657, and identified baseline cholinesterase and bilirubin as the most important variables (10). Reporting in nuclear cardiology often involves the prediction of coronary artery disease and the associated risk of major adverse cardiac events. The same group also evaluated the added predictive value of combining clinical information and myocardial perfusion imaging using the LogitBoost algorithm to predict major adverse cardiac events. Finally, when complex cases or rare diseases are being reported, it is often helpful to compare them with similar cases from databases and case collections. Although a textual search for example, in archived reports is uncomplicated, an image-based search is often not possible. Through AI-based automatic image annotations (11) and content-based image retrieval, conducting large, direct image-based and ad hoc database searches and thereby finding potentially similar cases that might be helpful in a real diagnostic situation are increasingly possible.

2. Limitations of AI

Although the use of AI in health care certainly holds great potential, its limitations also need to be acknowledged. A well-known problem is the interpretability of the models. Although symbolic AI or simple machine learning models, such as decision trees or linear regression, are still fully understood by people, understanding becomes increasingly difficult with more advanced techniques and is now impossible with many deep learning models; this situation can lead to unexpected results and nondeterministic behavior (1). Although this issue also applies to other procedures in medicine in which the exact mechanisms of action are often poorly understood (e.g., pharmacotherapy), whether predictive AI can and may be used for far-reaching decisions if the exact mode of action is unclear remains

unresolved. However, in cases in which AI acts as an assistant that provides hints or produces results that can be replicated by people or visually verified (e.g., by volumetry), the lack of interpretability of the underlying models may not be an obstacle to clinical application. For other cases, especially in image recognition and interpretation, certain techniques (such as activation maps) can provide high-level visual insights into the inner workings of ANNs. The problem of interpretability is the subject of intensive research and various initiatives, although whether these will be able to keep pace with the rapid progress in the development of increasingly complex ANN architectures is unclear. Another problem is that many machine learning applications will always deliver a result on an input but cannot provide a measure of the certainty of their prediction. Thus, a human operator often cannot decide whether to trust the result of AI-based software or not. Possible solutions for this problem are the integration of probabilistic reasoning and statistical analysis in machine learning as well as quality control. Bias and prejudice are well-known problems in medicine. However, training AI systems with biased data will make the resulting models generate biased predictions as well (12) this issue is especially problematic because many users perceive such systems as analytically correct and unprejudiced and therefore tend not to question their predictions in terms of bias. One of the largest hurdles for AI in health care is the need for large amounts of structured and annotated data for supervised learning. Many studies therefore work with small datasets, which are accompanied by overfitting and poor generalizability and reproducibility. Therefore, increased collaboration and standardization are needed to generate large machine-readable datasets that reflect variability in real populations and that have as little bias as possible

3. Outlook and Future Perspective

Many publications on the topic of AI in medicine deal with some degree of automation. Whether it is the measurement (quantification and segmentation) of pathologies, the detection of pathologies, or even automated diagnosis, AI does not necessarily have to be superhuman to have a benefit for medicine. However, it is obvious that AI is already better than people in some areas, and this development is a consequence of technologic progress. Therefore, many physicians are concerned that they will be replaced by AI in the future—a concern that is partly exacerbated by insufficient knowledge of how AI works. On the other hand, Geoffrey Hinton, undoubtedly one of the most renowned AI experts, made the statement, “People should stop training radiologists now!” at a conference in 2016. This statement triggered a lot of contradiction and is perhaps best explained by a lack of understanding of medicine in general and medical imaging in particular on his part. Although most experts and surveys reject the fear of AI replacing physicians, this fact does not mean that AI will have no impact on the medical profession. In fact, it is highly likely that AI will transform the medical profession and medical imaging in particular. In the near future, the automation of labor-intensive but cognitively undemanding tasks, such as image segmentation or finding prior examinations across different PACS repositories, will be available for clinical application. This change should be

perceived not as a threat but as an opportunity to relieve oneself of this work and as a stimulus for the further development of the profession. In fact, it is imperative for the profession to grow into the role it will be given in the future by AI. The increasing use of large amounts of digital data in medicine will create the need for new skills, such as clinical data science, computer science, and machine learning, especially in diagnostic disciplines. It can even be assumed that the boundaries between the diagnostic disciplines will become blurred, as the focus will increasingly be less on the detection and classification of individual findings and more on the comprehensive analysis and interpretation of all available data on a patient (13). Although prospective physicians can be confident that medical imaging offers them a bright future, it is important for them to understand that this future is open only to those who are willing to acquire competencies like those mentioned earlier. Without the training of and necessary expertise among physicians, precision health care, personalized medicine, and super diagnostics are unlikely to become clinical realities. As Chan and Siegel and others have stated, physicians will not be replaced by AI, but physicians who opt out from AI will be replaced by others who embrace it.

Acknowledgment

The Author would like to thanks Shruti Kulkarni for valuable help for collecting data.

References

- [1] McCarthy J, Minsky ML, Rochester N, Shannon CE. A proposal for the Dart-mouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AIMag*.2006; 27:12–14.
- [2] McCulloch WS, Pitts W. A logical calculus of the ideas immanent in nervous activity.1943.*BullMathBiol*.1990;52:99–115
- [3] Rosenblatt F. The perceptron: a probabilistic model for information storage and organization in the brain. *PsycholRev*.1958; 65: 386–408.
- [4] Hainc N, Federau C, Stieltjes B, Blatow M, Bink A, Stippich C. The bright, artificial intelligence-augmented future of neuroimaging reading. *Front Neurol*.2017; 8: 489.
- [5] Harvey HB, Liu C, Ai J, et al. Predicting no-shows in radiology using regression modeling of data available in the electronic medical record. *J Am Coll Radiol*.2017; 14:1303–1309.
- [6] Hong X, Zan Y, Weng F, Tao W, Peng Q, Huang Q. Enhancing the image quality via transferred deep residual learning of coarse PET sinograms. *IEEE Trans Med Imaging*.2018;37:2322–2332
- [7] Kim K, Wu D, Gong K, et al. Penalized PET Reconstruction using deep learning prior and local linear fitting. *IEEE Trans Med Imaging*.2018;37:1478–1487.
- [8] Xiang L, Qiao Y, Nie D, An L, Wang Q, Shen D. Deep auto-context convolutional neural networks for standard-dose PET image estimation from low-dose PET/MRI. *Neurocomputing*.2017; 267:406–416.
- [9] Ding Y, Sohn JH, Kawczynski MG, et al. A deep learning model to predict diagnosis of Alzheimer disease by using ¹⁸F-FDG PET of the brain. *Radiology*.2019;290:456–464
- [10] Huang B, Chen Z, Wu P-M, et al. Fully automated delineation of gross tumor volume for head and neck cancer on PET-CT using deep learning: a dual-center study. *Contrast Media Mol Imaging*.2018;2018:8923028
- [11] Milgrom SA, Elhalawani H, Lee J, et al. A PET radiomics model to predict refractory mediastinal Hodgkin lymphoma. *Sci Rep*.2019; 9:1322.
- [12] Robinson R, Valindria VV, Bai W, et al. Automated quality control in image segmentation: application to the UK Biobank cardiovascular magnetic resonance imaging study. *J Cardiovasc Magn Reson*.2019;21:18
- [13] Jha S, Topol EJ. Adapting to artificial intelligence: radiologists and pathologists as information specialists. *JAMA*.2016; 316:2353–2354.