4 Evolution of AI in Medical Imaging

Josh Schaefferkoetter

Contents

In the field of medical imaging, the application of computer vision to solve radiologic problems has been proposed since the mid-twentieth century [1]. As computers became more prevalent and imaging became digitized, the infrastructure was in place upon which to build sophisticated analysis pipelines to be used in routine workflow—this workflow has included, and will certainly continue to include, different applications of artificial intelligence. Today, AI is fundamental in many facets of everyday life, from semantic searches on the internet to facial and voice recognition in mobile devices, and it has made remarkable progress in recent years. There are various potential applications of AI in medicine, and AI has already impacted radiology in some regards, introducing quantification into a space which was historically based purely on subjectivity [2, 3]. This however is just the beginning—it is widely recognized that medical imaging is one of the many fields in which advanced AI will cause a complete paradigm shift. Molecular imaging in particular is an especially likely candidate to benefit, and it is in a position which would allow it to readily integrate this technology.

Molecular imaging technologies have continually improved year over year. MRI developments include higher field strength magnets,

J. Schaefferkoetter (✉)
Siemens Medical Solutions USA, Inc., Knoxville, TN, USA
e-mail: joshua.schaefferkoetter@siemens-healthineers.com

P. Veit-Haibach, K. Herrmann (eds.), *Artificial Intelligence/Machine Learning in Nuclear Medicine and Hybrid Imaging*, https://doi.org/10.1007/978-3-031-00119-2_4

improved RF coil arrays increasing acquisition SNR, and a growing catalog of pulse sequences for various applications. Single photon emission computed tomography (SPECT) systems routinely employ advanced correction techniques now producing quantitative images, and modern positron emission tomography (PET) scanners are using smaller crystals leading to better spatial resolution, with detection systems approaching timing resolution close to 200 ps. All of these modalities have realized concurrent progress in data processing as well, including sophisticated reconstruction and motion correction techniques. These advances have yielded extraordinary levels of image quality, but point is approaching where it is becoming less clear how these improvements are practically realized in terms of clinical outcomes. For instance, producing images with superfine resolution for routine examinations might not significantly impact diagnostic reliability, staging, or treatment planning. In fact, the additional time taken for the data acquisition and radiologist interpretation would potentially have adverse effects on the clinical workflow. Furthermore, in recent years, the amount of medical imaging data has grown exponentially, and this has already increased the pressure on radiologists to maintain accuracy at higher throughput. While novel imaging innovations will continue to have impact on patient care and be welcomed by the medical community, it is likely that technological developments in the near future will focus on increasing efficiency, reliably standardizing care, and improving patient safety.

Artificial intelligence, by definition, is the branch of computer science, developing computer algorithms to perform jobs normally requiring human intelligence. Machine learning (ML) is a subgroup of AI connoting any algorithm which improves through experience. There are many different schemes, ranging in complexity from simple regression models and component analyses to more complex methods like random forests and support vector machines. However, most of the remarkable successes and resulting excitement of recent times belong to the class of ML known as deep learning (DL). State-of-the-art results have been achieved in the fields of object detection, classification, image segmentation, speech recognition, and image generation—in fact, DL models have matched and even surpassed human performance in certain tasks [4–6]. It is impossible to ignore that these tasks are ubiquitous components in many aspects of radiology, and novel applications for DL are immediately identified. Indeed, there are many areas of active research in medicine and remarkable successes have been reported. Most reviews or general overviews of DL in medicine cite the growing number of related publications on PubMed, and at the time of this writing, the search phrase “deep learning” returned 5315 results for 2019. This is up from 3004 in the previous year, and for 2020, there are already 3994 results in the first 6 months. This trend is certainly a testament to the applicability and success of DL in medicine.

It is difficult to understand the evolution and future direction of AI without a basic understanding of the recent advances in AI techniques. This section gives an abbreviated overview, detailing a few specific examples. It cannot possibly cover all aspects but will instead focus on DL, since it is, without question, the dominant trend and direction of recent AI research; it has demonstrated promising improvements even over other traditional ML approaches. Almost all DL techniques are based on artificial neural networks (ANNs) comprising layers of numerical weights and “activation” nodes. More specifically, each node within a layer generally consists of a linear operation involving the summed product of its weights and input (the outputs of the previous layer), followed by a nonlinear operation, e.g., sigmoid, hyperbolic tangent, rectified linear—there may be thousands of nodes in a given layer. By stacking many of these layers, through densely interconnected nodes, one can effectively piecewise construct complex functions which are able to be shaped throughout many degrees of freedom. In this sense, a network can be shaped to “learn” mapping functions between different domains. Unlike most other ML approaches, DL does not require inputs which explicitly define the discriminating features of the population; through training, it inherently learns the features which best represent the data for the current task. This data-driven approach allows DL applications to characterize more abstract features and makes these systems more generalizable, but it is predicated on the availability of large amounts of training data to enable accurate characterizations of the sample populations.

Convolutional neural networks (CNNs) are an extension of neural networks, designed to handle data with higher dimensionality, usually in 2D or 3D, and so are well suited for image-based tasks. In conventional ANNs, the weights at each layer have a single, unique value for every combination of nodes of its layer and the nodes of the previous layer, and so the corresponding total number of weights at each layer is the product of these numbers. For CNNs, instead of a single value, there is a matrix of values, which can be thought of as a weighted filter; the size of the matrices is relatively small. The filters are passed over the layer input data like a convolution kernel, resulting in output feature maps of the same dimensionality as the input. This approach exploits the spatial dependencies within the data and makes the network invariant to input translations, while at the same time significantly reducing the total number of network parameters. For example, say we have a single 2D input image with pixel dimensions 100×100, and this feeds a layer with 128 channels. A conventional ANN would handle each of the 10,000 input pixels independently, and so the total number of parameters would be 1,280,000 for that single layer. For a CNN, this corresponding layer would handle the whole image as a single, multidimensional input—with a filter size 3×3, the total number of layer parameters would then only be 1152 ($1 \times 3 \times 3 \times 128$). This scheme is not only more efficient but potentially allows the same network to handle inputs of arbitrary sizes. For these reasons, CNNs are currently the AI technique of choice for image analyses and computer vision tasks.

Various CNN architectures are currently used—a few are explicitly mentioned here, but many of the basic concepts are common with many other networks. The convolution layers typically have filters with sizes between 3 and 5 pixels (for each dimension), and most networks also have multiple resolution downsampling (or encoding) layers. Many of the early uses for CNNs were focused on classification tasks and used a nonconvolutional, densely connected layer at the last layer to sort the output in scalar class probabilities [7]. Fully convolution networks (FCNs), however, do not contain any densely connected layers and preserve the input dimensionality throughout the network—this architecture is better suited to certain analysis tasks, i.e., when requiring a dense prediction map over all pixels [8]. The U-Net architecture has become widely used in image analyses [9] and uses a dedicated encoding and decoding path to produce outputs of the same size as the inputs. A major contribution of U-Net was the introduction of skip connections between the encoding and decoding paths at each resolution level in order to preserve spatial detail throughout the network—this feature makes this architecture popular for medical image segmentation tasks. Another useful architecture is ResNet, which is built on residual blocks containing multiple convolution layers, with the block input directly connected to its output [10]. This direct connection results in an alternate identity path, and so each convolutional block needs only to learn the pixel residuals and is pre-conditioned to learn mappings which are close to identity; the ResNet architecture has facilitated training stability in some of the deepest networks. The last relevant architecture is called Inception [11]. It contains blocks of multiple streams, each with different numbers of convolutions, under the premise that explicit filter sizes need not be defined since the image is now analyzed at multiple scales at the same level, i.e., taking the network wider rather than deeper. There is also a powerful extension of this called Inception-ResNet, which as the name implies, uses Inception blocks, rather than blocks of single-convolution streams, to calculate the block residuals.

Alongside the evolution of network architectures were concurrent advances in network training approaches. In the context of ML, training refers to the minimization of an objective loss metric corresponding to a certain task, i.e., some measure of distance between the network output and target value. In more basic terms, this means the values of the network weight parameters are gradually modified so that the desired outputs are obtained. This is usually accomplished by backpropagating the derivative of the loss through the network. Backpropagation is a computationally efficient method, combining simple mathematical operations, to generate a gradient of partial derivatives comprising the influences on the loss of every parameter in the network. After a complete backpropagation cycle, each network parameter is updated according to a predefined schedule in the direction which minimizes the loss. This process is repeated for many, sometimes millions, of iterations until acceptable performance is achieved.

In general, there are two fundamental approaches to training ML systems, supervised and unsupervised. Under supervised approaches, the input data have corresponding labels, and gradient backpropagation begins with a loss calculation over every output element of the network. For example, a CNN designed for classification might predict the correct class for a given input image by finding the maximum of the discrete probabilities calculated over all possible classes—during training, it would compare this prediction to the correct label and backpropagate its error differentials. In a simple classification task, each possible class might be represented as a single node in the output layer. This concept is readily extended to FCNs, in which a classification framework might be used for organ segmentation, for example. In this situation, the loss would be calculated over each pixel, giving the likelihood that it belongs to a given tissue class. Supervised methods provide a direct objective but require manual data labeling or annotating, which is a laborious task and is often the main challenge given the large scale of data typically needed for training. Unsupervised methods, on the other hand, do not require labeled data and instead rely on the algorithm itself to extract the discriminating features within different sample populations to minimize the loss for the task at hand. There are several methods for unsupervised network training, but one approach stands out for its range of applicability and remarkable recent results, and it is designed for image-based tasks performed by CNNs. Generative adversarial networks (GANs), introduced in 2014, comprise a system of two networks [12]. The first is the primary network, the generator, which for simplicity, can be regarded no differently than the networks discussed above—its job is to perform the desired task. However, instead of defining the training loss directly at its top layer with labels, the generator's output is fed into the second network, the discriminator, and the job of this network is to distinguish the generator's outputs from a corresponding set of real samples. During training, the discriminator learns the features that are common to the real and generated populations as a whole and uses this information to discriminate between the two sample sets. However, this same information can also be backpropagated to the generator and used to improve its own output. In this way, the two networks are adversaries in that they are each constantly trying to outperform the other, but at the same time, both the networks can simultaneously improve together. Deep learning systems built on the GAN framework have been tailored for specific applications in a wide range of fields and have demonstrated state-of-the-art performance, especially for image generation, translation, and transformation tasks.

Artificial intelligence has already established applications in the medical field. Novel investigations however, particularly those based on DL, are yielding especially impressive results, and these provide a glimpse of the direction of AI and hint at its potential future role in molecular imaging. The following sections provide an abbreviated outline of its historical and current uses and also highlight some areas of emerging research.

4.1 Disease Characterization

Characterization is a general term implying the segmentation, diagnosis, and staging of disease. These tasks are achieved by identifying and measuring the imaged properties of a pathologic abnormality. A radiologist performing these analyses is therefore required to process large amounts of data for each examination, and he or she must then distill it down into a manageable, and much smaller, number of qualitative features, e.g., size, shape, heterogeneity, to serve as the basis for the final interpretation. Inevitably, some radiological information is lost throughout this process. Furthermore, every physician is different, and there will be unavoidable variability among human observers. Artificial intelligence can help to automate this procedure. It has the capacity to consider large numbers of quantitative features, potentially orders of magnitude greater than a human, and it could perform the task in a fraction of the time in a reproducible way. For example, benign and malignant pulmonary nodules have similar appearances, and hence, the status of malignancy in the lungs is difficult to assess. AI can account for many features simultaneously and automatically determine those which are most relevant to the current case. The relevant features could be treated as imaging biomarkers to be used in the malignancy prediction, along with other clinical endpoints like risk assessment and prognosis [13].

The idea to use AI for disease characterization and diagnosis dates back to the mid-twentieth century [14–17]. Many of these studies focused on the improved interpretation of electrocardiograms by computers [18–21] since these data are particularly suitable for computer analyses. Other related work included the differential diagnosis of hematological diseases [22], automatic biochemical analysis of bodily substances [23], and sclerosis prediction in the coronary arteries [24]. These efforts mostly comprised smaller pilot studies and reported some success. Although larger-scale, definitive experiments were not performed during this time, these efforts led to the general belief that automatic diagnoses by computers were not just feasible, but necessary as part of a comprehensive medical data control system [25–27]. These early studies fostered an optimistic outlook for the potential of machine-assisted diagnosis and led to many advancements in computer-aided diagnostic (CAD) programs.

Dedicated CAD programs have early roots [28], but researchers only started large-scale development toward practical solutions in the 1980s. Significant effort was made in the research arena, but the benefits to the real clinical applications fell short [29], and it was not until 1998 that the FDA approved its use in screening and diagnostic mammography, as well as in plain chest radiography and CT imaging. Today, several systems are in clinical use with screening mammograms [30]. They are typically recommended to serve as a second opinion, complementing the initial radiologist assessment [31], and these led to the development of similar systems for other imaging modalities, including ultrasonography and MRI [32].

These conventional CAD systems generally consist of two components: detection of suspicious lesions and reduction of the false positive findings. The detection system is based on radiologist-defined criteria like tumor volume, shape, texture, etc. which are translated into a pattern-recognition problem where the most robust features are fed into an algorithm to highlight suspicious objects in the image [33]. The false-positive reduction part is also based on traditional ML, but can pose a bigger challenge to these algorithms. Even with sophisticated programs, the general performance of current CAD systems is not good, and this limits their extensive clinical use. Several trials have concluded that these systems, at best, deliver no benefit [34, 35]. It is more concerning though that these systems were actually found to reduce radiological accuracy in some cases [36], leading to higher recall and biopsy rates [37, 38].

Conventional CAD systems are built on rigid ML algorithms, mostly relying on expert knowledge, established a priori, for engineering features to be extracted from regions of interest. In contrast, new programs built on DL algorithms offer potential advantages regarding the degrees of freedom and level of abstraction in which the detection and classification tasks are defined. Furthermore, the performance of conventional CAD systems is notoriously sensitive to image noise and selected scanning protocol, and DL has demonstrated flexibility with regard to these parameters [39].

Largely due to the advances in computer hardware and processing technology, DL applications have emerged only recently for CAD systems—perhaps the earliest use in radiology was first reported in 1990, when a group at the University of Chicago developed an ANN for improving differential diagnosis of interstitial lung diseases using clinical and radiographic information. They claimed that the decision performance of the neural network was comparable to that of the chest radiologists and even superior to that of the senior radiology residents [40]. This led to several subsequent studies at that institution investigating neural network-aided diagnoses of lung disease [41–43]. The first object detection system using CNNs was proposed a few years later in 1995 at Georgetown University Medical Center, using a CNN with four layers to detect nodules in X-ray images [44].

Since then, DL-based CAD systems have been developed for the identification, detection, diagnosis, and risk analysis of various pathologies. Breast cancer, for example, was an obvious target since there was a historical precedent, and recent studies have demonstrated promising results regarding the performance of these next-generation systems in detecting and staging the diseases [45, 46]. In particular, it was reported that the automatic feature exploration and higher noise tolerance of DL-based CAD systems were responsible for the performance gains, which were quantified using different metrics, including sensitivity, specificity, and receiver operating characteristic analyses [47]. Lung cancer detection and screening is another attractive application, and several studies have evaluated the implementation of DL-based CAD systems for this purpose [48, 49]. These have also shown potential to effectively predict lung cancer and classify pulmonary nodules [47, 50]. In dermatology, deep convolutional networks have been used to classify skin lesions according to malignancy [51]. This large study found that AI achieved equivalent performance to all tested experts on two separate classification tasks, and further, it suggested that smartphone cameras could be used in conjunction with this technology to provide low-cost access to vital diagnoses. Other groups have also investigated DL with multi-modal imaging data. One notable study used PET and computed tomography (CT) data together in order to reduce false-positive results in lung lesion detections [52]. Simultaneous PET/CT data have also been used to classify lymph node metastases; a recent work found that this approach yielded higher sensitivities than radiologists [53]. Studies are consistently showing that the detection performance of AI in dedicated tasks is rivaling that of physicians [54], and recent interest in pursuing large-scale CAD solutions suggests the future for developing robust, high-performance systems based on deep learning [55].

Deep learning has also demonstrated success for using radiological information, not just for disease detection and characterization, but for predicting patient diagnosis and prognosis. Early works in this area included survival predictions in patients with lung adenocarcinoma [56] and high-grade gliomas [57]. More recently, DL algorithms have been developed to predict the risk of lung cancer from a patient's current and prior CT volumes [58]. This work achieved a state-of-the-art predictive performance on thousands of national lung cancer screening trial cases and independent clinical validation sets. This work also noted that its AI-based model reduced many risks associated with conventional low-dose CT screening, including false positives, overdiagnoses, and radiation exposure. The computer-aided detection and diagnosis of Alzheimer's disease (AD) is another area of active DL research. SPECT and PET are both used by physicians to image the metabolism, protein aggregation, or amyloid deposition associated with AD, and a few studies have investigated DL-based CAD systems for early AD diagnoses. The flexibility of DL allows brain data from multiple modalities to be assessed together [59–61]. Two notable recent works even used 3D CNNs to classify patients having AD [62, 63]. In other functional neurological studies, Parkinson's disease has been automatically diagnosed in dopamine active transporter SPECT scans, achieving sensitivities around 95% [64, 65]. Other work has been performed with PET/CT and PET/MR data, and the inclusion of multimodal inputs, exploiting functional and structural information, has the potential to further improve the performance of AI-based disease characterization.

4.2 Segmentation

Segmentation is an important component of medical image analyses—indeed, many of the aforementioned applications regarding the characterization of disease may be predicated on accurate delineations of organs, tissue or pathologic region of interest. It can often be a tedious and arduous task, and techniques to reliably speed the process would be welcomed by medical practitioners. Automatic segmentation methods using computer vision date back to the 1980s [66], with continual improvement over the following decades. Early approaches were based on clustering to isolate areas of similar intensities or region growing algorithms which spatially expanded regions around a user-selected seed point until homogeneity dropped below a certain criterion [67]. The next-generation algorithms used statistical learning and optimization to improve accuracy. One such approach is the watershed algorithm, in which image values are used to construct topology-like maps [68]. More advanced systems were able to use previous knowledge to construct a probability map to inform the segmentations. This approach is analogous to Bayesian inference, and the use of prior information lends itself, for example, to situations where objects are ill-defined in terms pixel intensities. The use of probability maps has proven especially helpful for oncologic segmentation within patient populations, since they contain information regarding the expected location of tumors [69]. Other segmentation systems based on prior knowledge-based probability maps have also been applied to radiotherapy planning in head and neck CT images [70] and segmenting gliomas in brain MRIs [71].

These past techniques have realized some success in the clinical workflow, but the algorithms are somewhat inflexible and were designed for specific tasks. Segmentation programs built on DL technology will significantly outperform their predecessors, and for these applications, fully convolutional networks are well suited. A major step toward semantic segmentation by FCNs was reported by UC Berkeley in 2015 [8]. This group first constructed FCNs by "decapitating" the fully connected layers from conventional CNNs, and replacing them with new layers to expand the resolution. This resulted in a network which produced an output having the same dimensions as its input, and by fine-tuning only the new layers, the parameters of the original layers which had already been trained on millions of images for classification tasks were not affected. The result was a network which was able to exploit the feature extraction mechanisms of the original network and apply this information to a dense prediction matrix. These researchers achieved impressive results, effectively using an FCN to segment detailed regions based on multiclass probabilities predicted for every discrete pixel [72]. Although this work focused only on natural images, the concept is readily extended to medical images.

Substantial attention has been paid to CNNs to resolve the challenges associated with medical imaging segmentation. Many techniques have been evaluated for various applications—a few specific examples include the automatic segmentation of lungs [73], biological cells and membranes [74, 75], tibial cartilage [76], bone tissue [77], brain structures [78], prostate [79], and tumors [80–83]. An important contribution came in 2015 with the introduction of the U-Net architecture and skip connections [9]. U-Net has been the de facto choice for many applications, including segmenting multiple organs on thoracic CT images with 3D data [84] or as incorporated into a GAN framework [85]. This network architecture also led to other derivatives like V-Net, which introduced a novel loss function directly based on the Dice coefficient [86].

Segmentation platforms built on DL offer other general advantages over older AI techniques as well. One study describes that DL methods for brain MRI segmentation completely eliminate the need for image registration required by other approaches like atlas-based methods [87]. It has also been reported that a single DL system is able to perform diverse segmentation tasks, without task-specific training, across multiple modalities and tissue types, including brain MRI, breast MRI, and cardiac CT angiography [88]. Considering this with the fact that current

DL technologies are already equivalent in many regards to radiologists' performance for segmentation [89], it is expected that the presence of DL-based segmentation algorithms in routine clinical tools will increase dramatically in the near future.

4.3 Image Generation/Reconstruction

Images are fundamental in radiology and diagnostic medicine. It was Wilhelm Roentgen who first discovered X-rays could be used to image bone just prior to the turn of the twentieth century. These early images were created directly, simply by exposing photographic film with the high-energy radiation. Over the next few decades, several other scanners were developed and some became digitized. This included the first positron-annihilation coincident detection system in the 1950s. A simple rectilinear scanner with sodium iodide detectors was designed and built by Gordon Brownell at Massachusetts General Hospital to image tumors in the brain. As imaging technology advanced throughout the century, so did the methods used to process the acquired data and produce the images. Certainly, one of the most groundbreaking inventions was the CT scanner in the 1970s by Sir Godfrey Hounsfield. This achievement ushered in the era of volumetric tomography, i.e., cross-sectional imaging of a 3D body, in the medical setting. The CT scanner acquired X-ray projection data at various angles for sequential axial positions. The projection data were used to reconstruct image slices by filtered back-projection (FBP), a direct reconstruction technique which is still used even today. FBP was used to reconstruct projection data for emission modalities as well like PET and SPECT as they made their way into nuclear medicine departments in the 1980s and 1990s. During this time, MRI systems also became a mainstream diagnostic tool. MR is unique from the others in that its images are generated directly through inverse Fourier transforms of the acquired frequency and phase data. For all imaging modalities, processing methods have made great strides over recent years, and through many recent advances, the images which are routinely produced in the clinic are of unprecedented quality. Artificial intelligence has the potential to push this even higher.

Until recent times, AI had not realized an overwhelming presence in image reconstruction. Conventional approaches relied on physics and closed-form mathematics to define the acquisition process and translate the data into images. However, recent decades have seen processing schemes which have become less rigid and more adaptive. Although these may not be considered AI, per se, they incorporate some of the same components. For example, direct reconstruction methods like FBP have been replaced by iterative algorithms. The objective of these algorithms is to find the image which is the most likely source of the projections—this framework can account for data which may be incomplete which results in far less image noise. The optimal image may be found by maximizing some likelihood or minimizing some cost measure, a technique which is often used in clustering machine learning algorithms. Also, many MR systems are moving toward compressed sensing to perform routine examinations in fractions of the time. Combining these under sampled data with prior information, images of high fidelity can still be produced.

Deep learning algorithms based on CNNs have incredible potential for applications in image reconstruction and generation. Research in this field is rapidly increasing, with the large majority of work focusing on MRI—only a relatively small subset of studies is mentioned here. A popular area is looking to AI for acceleration of MR imaging through improving compressed sensing techniques [90, 91]. Neural networks have demonstrated the ability to learn spatio-temporal dependencies which enable them to improve the accuracy of reconstructed MR images from highly undersampled complex-valued k-space data. This concept can be applied to dynamic MR imaging and may be especially interesting for cardiac cine protocols [92]. Furthermore, this idea has been extended to various MRI acquisition strategies. Recent algorithms have proved to be flexible for treating the MR reconstruction process as a supervised learning task, mapping the scanner sensors to resultant images [93]. Deep learning has also been used to reduce the gadolinium dose in contrast-enhanced brain MRI by an order of magnitude while preserving the quality of the images [94] and for inferring advanced MRI diffusion parameters from limited data [95]. Quantitative susceptibility mapping, which aims to estimate the magnetic susceptibility of biological tissue, is currently a growing field in MRI research [96, 97]. The estimation of magnetic susceptibility from local magnetic fields is an ill-posed problem, and recent AI methods are being used here as well. One work developed a CNN based on the U-Net architecture which was able to generate high-quality susceptibility maps from single orientation data [98]. MR-fingerprinting (MRF) is another recent technique [99]. The idea is to use a pseudo-randomized acquisition that captures a unique signal from different tissues. These tissue "fingerprints" are then mapped back to standard parameters, T1, T2, proton density, etc. by matching them to a predefined dictionary of predicted signal evolutions. This mapping is a difficult problem and has usually employed a pattern recognition approach—deep learning methodology is now being investigated for this purpose. A four-layer neural network was trained to map the recorded signal magnitudes to their corresponding tissue T1 and T2 values [100]. This group found reconstruction times using this approach were 300–5000 times faster than conventional dictionary-matching techniques in both phantom and human brain studies. Other similar approaches have been used to predict quantitative tissue parameter values from undersampled MRF data [101, 102].

Although MRI has so far realized the largest number of deep learning research efforts, these have potential applications extending to many areas in medical imaging on a more general scale. The last few years have seen impressive results for synthesizing photo-realistic images, especially using GANs [12, 103–105], and these techniques have also been used for biological image synthesis [106, 107]. One recent study designed a system to generate synthetic tumors in otherwise normal brain images [108]. This approach highlights a tremendously powerful use for generative networks, namely creating or augmenting training data. This is highly interesting for medical imaging as datasets are often sparse or imbalanced, with few examples of pathological findings. Overcoming this challenge would help alleviate a huge limitation commonly encountered in training deep learning models. This approach has been used for brain tumor segmentation [109], synthesizing realistic prostate lesions [110], augmenting data for improved liver lesion classification [111], and generating synthetic retinal fundus images [112]. GANs have also been used for unsupervised generation of T1-weighted brains [113] and image synthesis for tissue recognition and computer-assisted intervention [114, 115]. Inter-modality translation has even been performed by GANs, transforming MR to CT images [116, 117] and to PET images [118]. This work even showed that the generated images can be used in CAD systems for improving the diagnosis of Alzheimer's disease when the patient data are incomplete.

Artificial intelligence has provided a new paradigm for solving inverse problems in medical imaging [119–123]. Furthermore, studies have demonstrated the ability of DL to not only improve existing image reconstructions [124, 125] but also replace the reconstruction altogether, generating images directly from acquisition data [126]. This work found that a deep convolutional encoder–decoder network could be successfully used to generate quantitatively accurate PET images in a fraction of the time taken by conventional reconstruction methods. These works, and others like them, are incredibly encouraging. As a result, they have provoked a new, and necessary, avenue for research focusing solely on the potential pitfalls of DL-based reconstruction, and it has been found that deep learning can often cause unstable reconstruction methods. One recent work reported that these instabilities occur in several forms including: severe reconstruction artifacts caused by small perturbations in both the image or sampling domain; incomplete or incorrect representation of small structural changes, e.g., tumors; and more training samples yielded poorer reconstruction performance for several of the models investigated [127]. Numerical accuracy and stability are essential components of medical image reconstruction, and so the limitations of new technology are important to understand before it can be reliably used in the clinic. It is likely that, in the future, the image reconstruction process will be omitted altogether for certain applications, since a computer can theoretically extract any information contained in an image directly from the acquired data. For now, however, since humans perform the clinical interpretation, medical images need to be generated, and AI will continue to impact this process in unprecedented ways.

4.4 Data Corrections

As alluded to in the previous section, the methods to create medical images must be accurate and stable in order to be reliable—these requirements become even more critical when medical decisions depend on measurements of precisely quantified image values. Hence, the entire reconstruction process may comprise multiple steps to address different aspects. The backprojection algorithm, the cornerstone of tomographic reconstruction, can help to illustrate this. Data that are acquired as projections are mathematically regarded as a set of 1D line integrals, and backprojection seeks to invert this process and transform the sets of projections back to their original 2D form. However, due to the nature of the acquisition, low frequencies have a stronger latent prevalence within the projections than do the higher frequencies. So, to avoid a blurry reconstructed image dominated by low frequencies, the projection data must first be convolved with a ramp filter to boost the high frequencies. Additionally, the cylindrical geometry of the detection system results in nonuniform radial sampling, and this nonuniformity must also be accounted for in the reconstruction. This example demonstrates some of the steps necessary for a correct reconstruction approach, but backprojection is considered a direct method—newer, more sophisticated techniques usually require many additional considerations.

In addition to the corrections needed to compensate for the limitations of the acquisition method, the acquired data themselves may not be of high inherent quality. For PET, the true data come from pure annihilation photons, detected within a small coincidence window. However, the scanner also captures coincident events arising from scattered and random photons which must be corrected. These are not generally abled to be measured directly, so they must be estimated—this is currently accomplished by modeling the underlying physics. Photon scattering and absorption also leads to signal attenuation, and this requires an additional correction, usually based on an accompanying anatomical map. For MRI, the quality of the acquired data depends on the homogeneity of the static magnetic field, linearity of the gradients and stability of the receiver coils. These properties are bound by engineering limitations, and many techniques are routinely used to correct anomalies; for example, shimming is used to adjust the field homogeneity and spherical harmonic polynomial models can be used to characterize high-order gradient nonlinearities. However, sometimes these attempts are insufficient. Additionally, the MR scanner is very sensitive to environmental perturbations, and these can also lead to image noise. Artificial intelligence has proven adept at finding solutions to inference problems and should be able to help with issues related to incomplete or corrupted imaging data—indeed, it has already attained some notable successes.

Deep learning has recently been introduced to image denoising for many applications. In one study, neural networks were specifically developed to learn the implicit brain manifolds in MR images [113]. This group tested their approach by adding various levels of noise to several hundred T1-weighted brain images and reported improved performance over current denoising methods in terms of peak signal-to-noise ratios. Denoising has also been applied to dynamic contrast enhanced MR data, using multiple networks to improve the signal quality, both spatially and temporally [128]. Emission modalities have also been a focus of AI denoising research since they are inherently noisy. For instance, each projection bin of a routine PET acquisition may contain only a few coincident events, introducing uncertainty into the reconstruction. Several works within the last few years have reported success for PET image denoising using both supervised and unsupervised training approaches [129–131]. One notable study incorporated a 2D network pretrained on millions of natural images as a perceptual loss network [132]. This group reported that image resolution and noise properties were improved by optimizing the perceptual loss in this way, rather than simply using a per-pixel supervised loss like L1- or L2-norm. This approach has also been successfully applied for denoising CT images at various noise levels [133]. These reported successes have driven other research to investigate the potential clinical impacts of these methods. One such work reported improvements in physician lesion detectability performance when low-count PET images where denoised by a CNN [134].

Artifacts are another common nuisance in medical images—physiological or random patient motion, metal implants and temporal or spatial aliasing all cause distortions in the reconstructions. Deep learning methods have been used for correcting these. Techniques have been applied to automatically detect and correct patient motion for both MRI [135] and PET [136]. Motion does not only compromise imaging data. It can also affect techniques like MR spectroscopy, and approaches based on DL have been developed to remove ghosting artifacts in these studies [137, 138]. Regardless of their source, artifacts degrade the reconstructed spatial resolution. This of course limits the value of medical images for diagnoses, since good resolution properties are required to extract fine details from small pathological foci.

Improving medical image resolution has been the sole focus of many research efforts. Super-resolution in MRI has been around for over a decade [139]. These approaches enabled the reconstruction of a 3D volume with high isotropic resolution by acquiring the data typically through regular angular sampling about a common frequency encoding axis [140] or through modulation of the longitudinal magnetization to acquire independent k-space data [141]. Studies have reported success for estimating quantitative high-resolution T1 maps from a corresponding set of low-resolution maps [142] and even using conventional machine learning techniques to generate 7T-like MR images from 3T data [143]. Within the last few years, image super-resolution has become an interesting application for DL methods. Novel methods have produced state-of-the-art results for resolution up-sampling in natural images [144], and applications specific to MRI followed closely. Deep convolutional networks have constructed super-resolution brain [145] and musculoskeletal [146] images. These networks have also been adapted to generate super-resolution images from another modality [147].

The transformational mapping between multiple image domains is yet another exciting application for DL [148]. Due in part to recent advances in unsupervised training methods [149], this concept has found applications in medical research. Deep convolutional networks have been developed for transforming Flair to T1 MRI [150], CT to PET [151], and T1 MRI to CT [117]. Clinical interpretations and therapy planning based on images synthesized from another, unrelated modality could have far-reaching effects in the future of diagnostic and therapeutic medicine; this should be approached cautiously though, as synthesized images may contain incorrect pathological information and could lead to critical errors [150]. Notwithstanding this, image transformation based on DL may have the immediate potential to be a valuable tool for some technical problems. One popular current focus is related to PET/MR systems, transforming MR data to CT for PET attenuation correction. In order to produce quantitative images, photon attenuation must be corrected in all PET scans. This can be accurately estimated when an anatomical correlate of quantified attenuation values is available for directly generating a correction map, as it is with PET/CT. For PET/MR, however, this problem is more complicated since MR data do not contain information regarding photon scattering and absorption. Transforming MR images into quantified CT data has been implemented by several groups with promising results [152–154].

Furthermore, the PET/MR attenuation correction problem has also been addressed by omitting the CT transformation step altogether, using a CNN to estimate the correction map directly from the attenuated PET data themselves [155].

4.5 Image Registration

Once accurate medical images are produced, the image data must be translated into information which can be used for clinical patient management by a physician. In certain situations, the information obtained from multiple images read concurrently may be of much greater value than that obtained from reading them independently. The frequency of these situations dramatically increased at the turn of the twenty-first century for multimodal imaging with the invention of the PET/CT [156]. Multimodality imaging brought a new perspective into the field of clinical imaging. In this case, the combination of functional information with anatomical and morphological information provided an advanced medical tool, and countless studies over the past two decades have unequivocally established its diagnostic value. Other situations in which multiple images may be analyzed simultaneously include dynamic acquisitions, longitudinal comparisons or multiparametric MRI. In each of these cases, it is helpful, or even necessary, for the images to be spatially matched. For this reason, image registration is a constant focus of research, and techniques continue to evolve.

There are many potential sources of misregistration between two images of the same object, but assuming the differences are only spatially variant, one space can be mapped to the other through linear and nonlinear transformations. It is then the job of the registration algorithm to find the optimal transformation. For rigid structures, e.g., the head, linear transformations comprising global translations and rotations may be sufficient for coregistration. However, most other natural movement contains local, elastic deformations, and more complex methods are additionally needed to characterize and compensate for it. This is conventionally handled by projecting one image onto a grid, which is then deformed in a way which increases some joint similarity measure. Many different similarity metrics have been proposed and investigated, but common ones include correlation (for single-modality data) or mutual information (for multimodal data). The optimization algorithm typically combines these approaches within some convergence framework to try and maximize the relative similarity.

The registration problem comprises a challenging combination of many factors; decisions regarding the spatial transformations, similarity metrics, optimization strategies and numerical framework all play important roles in the performance. Machine learning techniques have been applied successfully for some specific applications in the past. However, as with other traditional ML techniques, these algorithms require explicitly handcrafting the features and have limited flexibility. In many cases, they are unable to meet the accuracy requirements of high-resolution medical imaging [157–160]. Recently, DL methods have been applied to image registration in order to improve accuracy and speed [161]. Image registration depends fundamentally on the identification of relevant information in the images, and this is a strength of deep neural networks. Convolution stacked auto-encoder networks, for example, have demonstrated the ability to identify intrinsic features in image patches [162], and CNNs have been developed for regressing the transformation parameters of the registration for multimodal data [163]. The flexibility of DL makes it well suited to address applications involving deformable registrations [162, 164]. Many groups have reported recent successes for specific tasks including elastic registration between 3D MRI and transrectal ultrasound for guiding prostate biopsy [165], deformable brain MRI registration [166], unsupervised CNN-based deformable registration for CT and MRI [167–169], and DL-based 2D/3D registration for registration of preoperative 3D data and intraoperative 2D X-ray images in image-guided therapy [170].

As diagnostic medicine continues to evolve, more complementary and multiparametric tissue information will be acquired in space and time—accurate image registration will become increasingly critical. Methods based on AI have shown impressive results and will undoubtedly play important roles in the automated clinical workflow, enabling quantitative comparisons at multiple timepoints and across different imaging modalities.

4.6 Radiology Reporting

The underlying goal of any medical imaging examination is a noninvasive survey of pathological information. Regardless of the imaging modality, the radiologic data must be read and translated into reports which are able to be used toward guiding patient management—these reports lie at the intersection of radiology and multiple downstream clinical subspecialties. These reports are sensitive to errors in the previous steps of the imaging pipeline, and so great care must be taken to clearly and accurately outline the relevant findings. This makes it an arduous and time-consuming task. Furthermore, subjectivity and inter-reader variability may introduce communication inconsistencies between radiology and other physicians. AI presents an attractive option for increasing speed and improving standardization of radiology reports.

Artificial intelligence algorithms for voice recognition and text generation were first proposed nearly two decades ago [171], and today, they are used routinely for radiologic reporting. Since then, machine learning techniques have made great strides in natural language processing, and now several vendors have developed powerful tools capable of speech-to-text translation, along with compatible hardware, e.g., dictation microphones [172]. These solutions have proven themselves invaluable for automatic transcription without the need for typing dictation content from radiologists, substantially reducing report generation times and improving clinical workflow.

Radiologic tools driven by deep learning algorithms have the potential to further streamline this process. Recently, DL has been used to automatically produce captions for natural photographic images [173], and this has led to many studies investigating potential applications for generating textual descriptions for medical images [174–181] and also for identifying findings in radiology reports [182–184]. Such AI tools could also replace the conventional qualitative nature of radiologic reporting with a more interactive quantitative one, and this approach has been shown to improve collaboration between radiology and oncology [185]. For example, it is plausible to expect that in the future, an AI-powered platform would be able to identify and diagnose pathological abnormalities and annotate them in a textual format that included quantified information about size, location, and probability of malignancy with associated confidence levels. These data would reduce subjective bias in decisions regarding patient management. Additionally, these well-structured reports would prove very beneficial to population sciences and big data mining initiatives. Another related avenue of DL research is using the generated radiologic reports themselves to annotate and label the imaging data. Medical PACS systems typically store thousands of free-text reports containing valuable information describing the images. Parsing this text and turning it into accurate annotations or labels requires sophisticated text-mining method—this is a field in which DL is currently being applied. Reports with higher degrees of structure more readily lend themselves to this purpose, and there are already some emerging applications. For example, there has been work reporting success leveraging radiologists' BI-RADS categorizations for training deep neural networks for characterizing breast lesions [174]. Considering the point that labeled data can be used to improve classification accuracy, one study was motivated by the fact that large amounts of annotated data might be unobtainable. This work proposed to create semantic description labels for the data, using both images and textual reports [186]. This group reported that semantic information can increase classification accuracy for different pathologies in medical images. Advanced AI algorithms are also being applied in other ways to improve efficiency in radiology practice. Convolutional neural networks can be used to determine scanning protocols from short text classification [187] and to improve time-sensitive decisions by prioritizing urgent cases [188]. One of the most interesting recent endeavors, however, addressed the challenges summarizing and representing patient data from electronic health records [189]. This work presented a novel unsupervised DL method for constructing general-purpose patient representations. This value of such data would be huge, since it could then potentially facilitate clinical predictive modeling on a large scale.

The applications mentioned above involved, to some degree, image interpretations based on human perception. Years of collecting data in routine clinical practice have produced an incredibly rich resource of quantified radiological data along with the associated clinical outcomes. These data are being leveraged to refine the field of radiomics. Radiomics in medicine refers to the high-throughput extraction of large amounts of features from medical images [190]. Radiomic analyses, sometimes involving high order statistics, can be used to identify patterns related to disease characteristics—patterns which may be undetectable by a traditional observer. Radiomics emerged from the field of oncology with the hypothesis that imaged tumors may reveal distinctive features pertaining to the disease which can be useful for predicting prognoses and planning personalized therapy [191, 192]. Early work in radiomics involved analyzing large sets of images and building correlations among various predefined features characterizing, for example, tumor morphology, intensity, and texture. Following this, many efforts have successfully applied radiomic evaluations for assisting clinical decision-making in oncology. For example, radiomics has been used to predict metastatic patterns in lung adenocarcinoma [193] as well as disease recurrence [194] and prognoses [195]. Recently, deep learning has been applied in this space [161]. As with many other examples presented in this chapter, DL poses advantages over traditional methods for automatically extracting the relevant features, while simultaneously providing information regarding their clinical relevance. Deep learning and radiomics are two rapidly evolving technologies, and their symbiosis will likely lead to a single unified framework to support clinical decisions—this has the potential to completely transform the field of precision medicine [13].

4.7 Conclusion

Fundamentally, medical images are generated in order to be presented to physicians for evaluation—optimizing the appearance of images for human viewers almost always includes simplification and down-sampling of the raw data. Quantitative approaches like radiomics represent a step toward automatic image interpretation using the latent information embedded in the images, and following this evolutionary track, it is expected in the future that the presence of automated, AI-driven analyses in routine clinical workflow will continue to increase. In this paradigm, processed medical images may become altogether unnecessary for certain indications. This would avoid the loss of information inherent in the creation of images, leading to reproducible analyses which were faster and more accurate.

In conclusion, AI has made great advances, especially recently, but it is not expected that it will outperform humans for general clinical planning and patient management in the near future. Instead, both will improve together. Although AI is currently able to provide advantages for specific quantitative tasks, medical decisions cannot be strictly regarded as such. They are based on knowledge obtained through life experience and philosophy. To incorporate these characteristics into an AI program, one would be faced with many challenges including data collection and algorithm development [29]. Considering this, it is likely that the trend in AI will move toward advanced unsupervised learning approaches, allowing the immense amounts of readily-available, unlabeled data to be utilized. In any case, the synergy between AI and physicians will certainly grow and continue to be mutually beneficial within the field of medical imaging, leading to unprecedented levels of precision and quality in patient care.

References

1. Ledley RS, Lusted LB. Reasoning foundations of medical diagnosis. Science. 1959;130(3366):9–21.
2. Haug PJ. Uses of diagnostic expert systems in clinical care. In: Proceedings of the annual symposium on computer application in medical care. Bethesda: American Medical Informatics Association; 1993.
3. Ambinder EP. A history of the shift toward full computerization of medicine. J Oncol Pract. 2005;1(2):54.
4. Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. Adv Neural Inf Proces Syst. 2012;25:1097–105.
5. He K, et al. Delving deep into rectifiers: surpassing human-level performance on imagenet classification. In: Proceedings of the IEEE international conference on computer vision. New York: IEEE; 2015.
6. Silver D, et al. Mastering the game of Go with deep neural networks and tree search. Nature. 2016;529(7587):484–9.
7. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
8. Long J, Shelhamer E, Darrell T. Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
9. Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2015.
10. He K, et al. Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.
11. Szegedy C, et al. Going deeper with convolutions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
12. Goodfellow I, et al. Generative adversarial nets. In: Advances in neural information processing systems. 2014.
13. Hosny A, et al. Artificial intelligence in radiology. Nat Rev Cancer. 2018;18(8):500–10.
14. Ledley RS. Using electronic computers in medical diagnosis. IRE Trans Med Electron. 1960;4:274–80.
15. Amosov N, Shkabara E. Experience in determining diagnosis with the aid of diagnostic machines. Eksp Khirurgiia. 1961;6:15–22.
16. Rikli AE, et al. Computer analysis of electrocardiographic measurements. Circulation. 1961;24(3):643–9.
17. Paycha F. Diagnosis with the aid of artificial intelligence: demonstration of the 1st diagnostic machine. Presse Therm Clim. 1968;105(1):22.
18. DeCote R, Horvath WJ. An electronic computer for vector electrocardiography. IRE Trans Med Electron. 1957;1957:31–7.
19. Caceres CA. How can the waveforms of a clinical electrocardiogram be measured automatically by a computer? IRE Trans Biomed Electron. 1962;9(1):21–2.
20. Pipberger HV, Stallmann F. Use of computers in ECG interpretation. Am Heart J. 1962;64:285.
21. Steinberg C, Abraham S, Caceres C. Pattern recognition in the clinical electrocardiogram. IRE Trans Biomed Electron. 1962;9(1):23–30.
22. Lipkin M, Hardy JD. Mechanical correlation of data in differential diagnosis of hematological diseases. J Am Med Assoc. 1958;166(2):113–25.
23. Jonnard R. Random selection system for automatic biochemical analysis-partial functional analysis. IRE Trans Biomed Electron. 1961;8(2):83–98.
24. Moyer D, Talbott G. Instrumentation for the diagnosis of coronary-artery disease. Trans Am Inst Electr Eng. 1962;80(6):717–21.
25. Ledley RS, Lusted LB. Computers in medical data processing. Oper Res. 1960;8(3):299–310.
26. Gillon J. Is automatic diagnosis in the future? Concours Med. 1962;84:3829–33.
27. Schweisheimer W. Can electronic machines facilitate and improve medical diagnosis? Hippokrates. 1962;33:162.
28. Lodwick GS, Keats TE, Dorst JP. The coding of roentgen images for computer analysis as applied to lung cancer. Radiology. 1963;81(2):185–200.
29. Tang X. The role of artificial intelligence in medical imaging research. BJR Open. 2019;2(1):20190031.
30. Champaign JL, Cederbom GJ. Advances in breast cancer detection with screening mammography. Ochsner J. 2000;2(1):33–5.
31. Shiraishi J, et al. Computer-aided diagnosis and artificial intelligence in clinical imaging. In: Seminars in nuclear medicine. London: Elsevier; 2011.
32. Ayer T, et al. Computer-aided diagnostic models in breast cancer screening. Imaging Med. 2010;2(3):313.
33. Nagaraj S, Rao G, Koteswararao K. The role of pattern recognition in computer-aided diagnosis and computer-aided detection in medical imaging: a clinical validation. Int J Comput Appl. 2010;8(5):18–22.
34. Cole EB, et al. Impact of computer-aided detection systems on radiologist accuracy with digital mammography. Am J Roentgenol. 2014;203(4):909–16.
35. Lehman CD, et al. Diagnostic accuracy of digital screening mammography with and without computer-aided detection. JAMA Intern Med. 2015;175(11):1828–37.
36. Fenton JJ, et al. Influence of computer-aided detection on performance of screening mammography. N Engl J Med. 2007;356(14):1399–409.
37. Gilbert FJ, et al. Single reading with computer-aided detection for screening mammography. N Engl J Med. 2008;359(16):1675–84.
38. Oakden-Rayner L. The rebirth of CAD: how is modern AI different from the CAD we know? Oak Brook: Radiological Society of North America; 2019.

39. Lee J-G, et al. Deep learning in medical imaging: general overview. Korean J Radiol. 2017;18(4):570–84.
40. Asada N, et al. Potential usefulness of an artificial neural network for differential diagnosis of interstitial lung diseases: pilot study. Radiology. 1990;177(3):857–60.
41. Lin J-S, et al. Reduction of false positives in lung nodule detection using a two-level neural classification. IEEE Trans Med Imaging. 1996;15(2):206–17.
42. Ashizawa K, et al. Artificial neural networks in chest radiography: application to the differential diagnosis of interstitial lung disease. Acad Radiol. 1999;6(1):2–9.
43. Ashizawa K, et al. Effect of an artificial neural network on radiologists' performance in the differential diagnosis of interstitial lung disease using chest radiographs. AJR Am J Roentgenol. 1999;172(5):1311–5.
44. Lo S-C, et al. Artificial convolution neural network techniques and applications for lung nodule detection. IEEE Trans Med Imaging. 1995;14(4):711–8.
45. Wang D, et al. Deep learning for identifying metastatic breast cancer. arXiv preprint arXiv:1606.05718, 2016.
46. Kallenberg M, et al. Unsupervised deep learning applied to breast density segmentation and mammographic risk scoring. IEEE Trans Med Imaging. 2016;35(5):1322–31.
47. Cheng J-Z, et al. Computer-aided diagnosis with deep learning architecture: applications to breast lesions in US images and pulmonary nodules in CT scans. Sci Rep. 2016;6(1):1–13.
48. Hua K-L, et al. Computer-aided classification of lung nodules on computed tomography images via deep learning technique. Onco Targets Ther. 2015;8:2015–22.
49. Kumar D, Wong A, Clausi DA. Lung nodule classification using deep features in CT images. In: 2015 12th conference on computer and robot vision. New York: IEEE; 2015.
50. Chen J, et al. Use of an artificial neural network to construct a model of predicting deep fungal infection in lung cancer patients. Asian Pac J Cancer Prev. 2015;16(12):5095–9.
51. Esteva A, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature. 2017;542(7639):115–8.
52. Teramoto A, et al. Automated detection of pulmonary nodules in PET/CT images: ensemble false-positive reduction using a convolutional neural network technique. Med Phys. 2016;43(6): 2821–7.
53. Wang H, et al. Comparison of machine learning methods for classifying mediastinal lymph node metastasis of non-small cell lung cancer from 18 F-FDG PET/CT images. EJNMMI Res. 2017;7(1):11.
54. Chen T, Metaxas D. Medical image computing and computer-assisted intervention—Miccai 2000. Vol. 1935 of lecture notes in computer science. New York: Springer; 2000.
55. Kooi T, et al. Large scale deep learning for computer aided detection of mammographic lesions. Med Image Anal. 2017;35:303–12.
56. Paul R, et al. Deep feature transfer learning in combination with traditional features predicts survival among patients with lung adenocarcinoma. Tomography. 2016;2(4):388.
57. Nie D, et al. 3D deep learning for multi-modal imaging-guided survival time prediction of brain tumor patients. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2016.
58. Ardila D, et al. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. Nat Med. 2019;25(6):954–61.
59. Suk H-I, et al. Hierarchical feature representation and multimodal fusion with deep learning for AD/MCI diagnosis. NeuroImage. 2014;101:569–82.
60. Suk H-I, Shen D. Deep learning-based feature representation for AD/MCI classification. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2013.
61. Liu S, et al. Early diagnosis of Alzheimer's disease with deep learning. In: 2014 IEEE 11th international symposium on biomedical imaging (ISBI). New York: IEEE; 2014.
62. Hosseini-Asl E, Gimel'farb G, El-Baz A. Alzheimer's disease diagnostics by a deeply supervised adaptable 3D convolutional network. arXiv preprint arXiv:1607.00556, 2016.
63. Payan A, Montana G. Predicting Alzheimer's disease: a neuroimaging study with 3D convolutional neural networks. arXiv preprint arXiv:1502.02506, 2015.
64. Choi H, et al. Refining diagnosis of Parkinson's disease with deep learning-based interpretation of dopamine transporter imaging. NeuroImage. 2017;16:586–94.
65. Kim DH, Wit H, Thurston M. Artificial intelligence in the diagnosis of Parkinson's disease from ioflupane-123 single-photon emission computed tomography dopamine transporter scans using transfer learning. Nucl Med Commun. 2018;39(10):887–93.
66. Haralick RM, Shapiro LG. Image segmentation techniques. Comput Vis Graph Image Process. 1985;29(1):100–32.
67. Pham DL, Xu C, Prince JL. Current methods in medical image segmentation. Annu Rev Biomed Eng. 2000;2(1):315–37.
68. Grau V, et al. Improved watershed transform for medical image segmentation using prior information. IEEE Trans Med Imaging. 2004;23(4):447–58.
69. Sharma N, Aggarwal LM. Automated medical image segmentation techniques. J Med Phys. 2010;35(1):3.
70. Han X, et al. Atlas-based auto-segmentation of head and neck CT images. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2008.

71. Parisot S, et al. A probabilistic atlas of diffuse WHO grade II glioma locations in the brain. PLoS One. 2016;11(1):e0144200.
72. Shelhamer E, Long J, Darrell T. Fully convolutional networks for semantic segmentation. IEEE Trans Pattern Anal Mach Intell. 2017;39(4):640–51.
73. Middleton I, Damper RI. Segmentation of magnetic resonance images using a combination of neural networks and active contour models. Med Eng Phys. 2004;26(1):71–86.
74. Ning F, et al. Toward automatic phenotyping of developing embryos from videos. IEEE Trans Image Process. 2005;14(9):1360–71.
75. Ciresan D, et al. Deep neural networks segment neuronal membranes in electron microscopy images. In: Advances in neural information processing systems. 2012.
76. Prasoon A, et al. Deep feature learning for knee cartilage segmentation using a triplanar convolutional neural network. In: International conference on medical image computing and computer-assisted intervention. 2013. Springer.
77. Cernazanu-Glavan C, Holban S. Segmentation of bone structure in X-ray images using convolutional neural network. Adv Electron Comput Eng. 2013;13(1):87–94.
78. Moeskops P, et al. Automatic segmentation of MR brain images with a convolutional neural network. IEEE Trans Med Imaging. 2016;35(5):1252–61.
79. Zhu Q, et al. Deeply-supervised CNN for prostate segmentation. In: 2017 international joint conference on neural networks (IJCNN). New York: IEEE; 2017.
80. Rastgarpour M, Shanbehzadeh J. Application of ai techniques in medical image segmentation and novel categorization of available methods and in tools. In: Proceedings of the international multiconference of engineers and computer scientists. Princeton: Citeseer; 2011.
81. Pereira S, et al. Brain tumor segmentation using convolutional neural networks in MRI images. IEEE Trans Med Imaging. 2016;35(5):1240–51.
82. Roth HR, et al. Deep learning and its application to medical image segmentation. Med Imaging Technol. 2018;36(2):63–71.
83. Tang X, Wang B, Rong Y. Artificial intelligence will reduce the need for clinical medical physicists. J Appl Clin Med Phys. 2018;19(1):6.
84. Feng X, et al. Deep convolutional neural network for segmentation of thoracic organs-at-risk using cropped 3D images. Med Phys. 2019;46(5):2169–80.
85. Dong X, et al. Automatic multiorgan segmentation in thorax CT images using U-net-GAN. Med Phys. 2019;46(5):2157–68.
86. Milletari F, Navab N, Ahmadi S-A. V-net: fully convolutional neural networks for volumetric medical image segmentation. In: Fourth international conference on 3D Vision (3DV). New York: IEEE; 2016. p. 2016.
87. de Brebisson A, Montana G. Deep neural networks for anatomical brain segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. 2015.
88. Moeskops P, et al. Deep learning for multi-task medical image segmentation in multiple modalities. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2016.
89. Ghafoorian M, et al. Location sensitive deep convolutional neural networks for segmentation of white matter hyperintensities. Sci Rep. 2017;7:5110.
90. Sun J, Li H, Xu Z. Deep ADMM-Net for compressive sensing MRI. In: Advances in neural information processing systems. 2016.
91. Wang S, et al. Accelerating magnetic resonance imaging via deep learning. In: 2016 IEEE 13th international symposium on biomedical imaging (ISBI). New York: IEEE; 2016.
92. Qin C, et al. Convolutional recurrent neural networks for dynamic MR image reconstruction. IEEE Trans Med Imaging. 2018;38(1):280–90.
93. Zhu B, et al. Image reconstruction by domain-transform manifold learning. Nature. 2018;555(7697):487–92.
94. Gong E, et al. Deep learning enables reduced gadolinium dose for contrast-enhanced brain MRI. J Magn Reson Imaging. 2018;48(2):330–40.
95. Golkov V, et al. Q-space deep learning: twelve-fold shorter and model-free diffusion MRI scans. IEEE Trans Med Imaging. 2016;35(5):1344–51.
96. Deistung A, et al. Toward in vivo histology: a comparison of quantitative susceptibility mapping (QSM) with magnitude-, phase-, and $R2_*$-imaging at ultra-high magnetic field strength. NeuroImage. 2013;65:299–314.
97. Deistung A, Schweser F, Reichenbach JR. Overview of quantitative susceptibility mapping. NMR Biomed. 2017;30(4):e3569.
98. Yoon J, et al. Quantitative susceptibility mapping using deep neural network: QSMnet. NeuroImage. 2018;179:199–206.
99. Ma D, et al. Magnetic resonance fingerprinting. Nature. 2013;495(7440):187–92.
100. Cohen O, Zhu B, Rosen MS. MR fingerprinting deep reconstruction network (DRONE). Magn Reson Med. 2018;80(3):885–94.
101. Hoppe E, et al. Deep learning for magnetic resonance fingerprinting: a new approach for predicting quantitative parameter values from time series. In: GMDS. 2017.
102. Fang Z, et al. Quantification of relaxation times in MR fingerprinting using deep learning. In: Proceedings of the International Society for Magnetic Resonance in Medicine... Scientific Meeting and Exhibition. International Society for Magnetic Resonance in Medicine. Scientific Meeting and Exhibition. Bethesda: NIH Public Access; 2017.

103. Creswell A, et al. Generative adversarial networks: an overview. IEEE Signal Process Mag. 2018;35(1):53–65.
104. Hong Y, et al. How generative adversarial networks and their variants work: an overview. ACM Comput Surv. 2019;52(1):1–43.
105. Huang H, Yu PS, Wang C. An introduction to image synthesis with generative adversarial nets. arXiv preprint arXiv:1803.04469, 2018.
106. Osokin A, et al. Gans for biological image synthesis. In: Proceedings of the IEEE international conference on computer vision. New York: IEEE; 2017.
107. Antipov G, Baccouche M, Dugelay J-L. Face aging with conditional generative adversarial networks. In: 2017 IEEE international conference on image processing (ICIP). New York: IEEE; 2017.
108. Shin H-C, et al. Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In: International workshop on simulation and synthesis in medical imaging. New York: Springer; 2018.
109. Mok TC, Chung AC. Learning data augmentation for brain tumor segmentation with coarse-to-fine generative adversarial networks. In: International MICCAI brainlesion workshop. New York: Springer; 2018.
110. Kitchen A, Seah J. Deep generative adversarial neural networks for realistic prostate lesion MRI synthesis. arXiv preprint arXiv:1708.00129, 2017.
111. Frid-Adar M, et al. Synthetic data augmentation using GAN for improved liver lesion classification. In: 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018). New York: IEEE; 2018.
112. Guibas JT, Virdi TS, Li PS. Synthetic medical images from dual generative adversarial networks. arXiv preprint arXiv:1709.01872, 2017.
113. Bermudez C, et al. Learning implicit brain MRI manifolds with deep learning. In: Medical imaging 2018: image processing. Bellingham: International Society for Optics and Photonics. p. 2018.
114. Zhang Q, et al. Medical image synthesis with generative adversarial networks for tissue recognition. In: 2018 IEEE international conference on healthcare informatics (ICHI). New York: IEEE; 2018.
115. Nie D, et al. Medical image synthesis with context-aware generative adversarial networks. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2017.
116. Nie D, et al. Estimating CT image from MRI data using 3D fully convolutional networks. In: Deep learning and data labeling for medical applications. New York: Springer; 2016. p. 170–8.
117. Wolterink JM, et al. Deep MR to CT synthesis using unpaired data. In: International workshop on simulation and synthesis in medical imaging. New York: Springer; 2017.
118. Li R, et al. Deep learning based imaging data completion for improved brain disease diagnosis. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2014.
119. Strack R. Imaging: AI transforms image reconstruction. Nat Methods. 2018;15(5):309–10.
120. McCann MT, Jin KH, Unser M. Convolutional neural networks for inverse problems in imaging: a review. IEEE Signal Process Mag. 2017;34(6):85–95.
121. Jin KH, et al. Deep convolutional neural network for inverse problems in imaging. IEEE Trans Image Process. 2017;26(9):4509–22.
122. Schlemper J, et al. A deep cascade of convolutional neural networks for MR image reconstruction. In: International conference on information processing in medical imaging. New York: Springer; 2017.
123. Lucas A, et al. Using deep neural networks for inverse problems in imaging: beyond analytical methods. IEEE Signal Process Mag. 2018;35(1):20–36.
124. Kim K, et al. Penalized PET reconstruction using deep learning prior and local linear fitting. IEEE Trans Med Imaging. 2018;37(6):1478–87.
125. Gong K, et al. Iterative PET image reconstruction using convolutional neural network representation. IEEE Trans Med Imaging. 2018;38(3):675–85.
126. Häggström I, et al. DeepPET: a deep encoder–decoder network for directly solving the PET image reconstruction inverse problem. Med Image Anal. 2019;54:253–62.
127. Antun V, et al. On instabilities of deep learning in image reconstruction and the potential costs of AI. Proc Natl Acad Sci. 2020;117(48):30088–95.
128. Benou A, et al. Ensemble of expert deep neural networks for spatio-temporal denoising of contrast-enhanced MRI sequences. Med Image Anal. 2017;42:145–59.
129. Liu C-C, Qi J. Higher SNR PET image prediction using a deep learning model and MRI image. Phys Med Biol. 2019;64(11):115004.
130. Lu W, et al. An investigation of quantitative accuracy for deep learning based denoising in oncological PET. Phys Med Biol. 2019;64(16):165019.
131. Cui J, et al. PET image denoising using unsupervised deep learning. Eur J Nucl Med Mol Imaging. 2019;46(13):2780–9.
132. Gong K, et al. Pet image denoising using a deep neural network through fine tuning. IEEE Trans Radiat Plasma Med Sci. 2018;3(2):153–61.
133. Yang Q, et al. CT image denoising with perceptive deep neural networks. arXiv preprint arXiv:1702.07019, 2017.
134. Schaefferkoetter J, Yan J, Ortega C, et al. Convolutional neural networks for improving image quality with noisy PET data. EJNMMI Res. 2020;10:105. https://doi.org/10.1186/s13550-020-00695-1.

135. Küstner T, et al. Automated reference-free detection of motion artifacts in magnetic resonance images. Medicine. 2018;31(2):243–56.
136. Li T, et al. Motion correction of respiratory-gated PET images using deep learning based image registration framework. Phys Med Biol. 2020;65(15):155003.
137. Gurbani SS, et al. A convolutional neural network to filter artifacts in spectroscopic MRI. Magn Reson Med. 2018;80(5):1765–75.
138. Kyathanahally SP, Döring A, Kreis R. Deep learning approaches for detection and removal of ghosting artifacts in MR spectroscopy. Magn Reson Med. 2018;80(3):851–63.
139. Robinson MD, et al. New applications of super-resolution in medical imaging. Super-Resolut Imaging. 2010;2010:384–412.
140. Shilling RZ, et al. A super-resolution framework for 3-D high-resolution and high-contrast imaging using 2-D multislice MRI. IEEE Trans Med Imaging. 2008;28(5):633–44.
141. Ropele S, et al. Super-resolution MRI using microscopic spatial modulation of magnetization. Magn Reson Med. 2010;64(6):1671–5.
142. Van Steenkiste G, et al. Super-resolution T1 estimation: quantitative high resolution T1 mapping from a set of low resolution T1-weighted images with different slice orientations. Magn Reson Med. 2017;77(5):1818–30.
143. Bahrami K, et al. 7T-guided super-resolution of 3T MRI. Med Phys. 2017;44(5):1661–77.
144. Ledig C, et al. Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.
145. Liu C, et al. Fusing multi-scale information in convolution network for MR image super-resolution reconstruction. Biomed Eng Online. 2018;17(1):114.
146. Chaudhari AS, et al. Super-resolution musculoskeletal MRI using deep learning. Magn Reson Med. 2018;80(5):2139–54.
147. Zeng K, et al. Simultaneous single-and multi-contrast super-resolution for brain MRI images based on a convolutional neural network. Comput Biol Med. 2018;99:133–41.
148. Isola P, et al. Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.
149. Zhu J-Y, et al. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. 2017.
150. Cohen JP, Luck M, Honari S. Distribution matching losses can hallucinate features in medical image translation. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2018.
151. Armanious K, et al. Unsupervised medical image translation using cycle-MedGAN. In: 2019 27th European signal processing conference (EUSIPCO). New York: IEEE; 2019.
152. Han X. MR-based synthetic CT generation using a deep convolutional neural network method. Med Phys. 2017;44(4):1408–19.
153. Torrado-Carvajal A, et al. Dixon-VIBE deep learning (DIVIDE) pseudo-CT synthesis for pelvis PET/MR attenuation correction. J Nucl Med. 2019;60(3):429–35.
154. Leynes AP, et al. Zero-echo-time and Dixon deep pseudo-CT (ZeDD CT): direct generation of pseudo-CT images for pelvic PET/MRI attenuation correction using deep convolutional neural networks with multiparametric MRI. J Nucl Med. 2018;59(5):852–8.
155. Spuhler KD, et al. Synthesis of patient-specific transmission data for PET attenuation correction for PET/MRI neuroimaging using a convolutional neural network. J Nucl Med. 2019;60(4):555–60.
156. Beyer T, et al. A combined PET/CT scanner for clinical oncology. J Nucl Med. 2000;41(8):1369.
157. Wohlhart P, Lepetit V. Learning descriptors for object recognition and 3d pose estimation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
158. Dollár P, Welinder P, Perona P. Cascaded pose regression. In: 2010 IEEE computer society conference on computer vision and pattern recognition. New York: IEEE; 2010.
159. Zach C, Penate-Sanchez A, Pham M-T. A dynamic programming approach for fast and robust object pose recognition from range images. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
160. Mottaghi R, Xiang Y, Savarese S. A coarse-to-fine model for 3d pose estimation and sub-category recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
161. Litjens G, et al. A survey on deep learning in medical image analysis. Med Image Anal. 2017;42:60–88.
162. Wu G, et al. Scalable high-performance image registration framework by unsupervised deep feature representations learning. IEEE Trans Biomed Eng. 2015;63(7):1505–16.
163. Sloan JM, Goatman KA, Siebert JP. Learning rigid image registration-utilizing convolutional neural networks for medical image registration. 2018.
164. Yang X, et al. Quicksilver: fast predictive image registration–a deep learning approach. NeuroImage. 2017;158:378–96.
165. Haskins G, et al. Learning deep similarity metric for 3D MR–TRUS image registration. Int J Comput Assist Radiol Surg. 2019;14(3):417–25.
166. Cao X, et al. Deformable image registration using a cue-aware deep regression network. IEEE Trans Biomed Eng. 2018;65(9):1900–11.
167. Shan S, et al. Unsupervised end-to-end learning for deformable medical image registration. arXiv preprint arXiv:1711.08608, 2017.

168. Kearney V, et al. An unsupervised convolutional neural network-based algorithm for deformable image registration. Phys Med Biol. 2018;63(18):185017.
169. de Vos BD, et al. A deep learning framework for unsupervised affine and deformable image registration. Med Image Anal. 2019;52:128–43.
170. Zheng J, et al. Pairwise domain adaptation module for CNN-based 2-D/3-D registration. J Med Imaging. 2018;5(2):021204.
171. Antoniol G, et al. Radiological reporting based on voice recognition. In: International conference on human-computer interaction. New York: Springer; 1993.
172. Liu Y, Wang J. PACS and digital medicine: essential principles and modern practice. Boca Raton: CRC Press; 2010.
173. Karpathy A, Fei-Fei L. Deep visual-semantic alignments for generating image descriptions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
174. Kisilev P, et al. Medical image description using multi-task-loss CNN. In: Deep learning and data labeling for medical applications. New York: Springer; 2016. p. 121–9.
175. Shin H-C, et al. Interleaved text/image deep mining on a very large-scale radiology database. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
176. Shin H-C, et al. Learning to read chest X-rays: recurrent neural cascade model for automated image annotation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.
177. Wang X., et al. Unsupervised category discovery via looped deep pseudo-task optimization using a large scale radiology image database. arXiv preprint arXiv:1603.07965, 2016.
178. Jing B, Xie P, Xing E. On the automatic generation of medical imaging reports. arXiv preprint arXiv:1711.08195, 2017.
179. Li Y, et al. Hybrid retrieval-generation reinforced agent for medical image report generation. In: Advances in neural information processing systems. 2018.
180. Moradi M, et al. Bimodal network architectures for automatic generation of image annotation from text. In: International conference on medical image computing and computer-assisted intervention. New York: Springer; 2018.
181. Zhang Y, et al. Learning to summarize radiology findings. arXiv preprint arXiv:1809.04698, 2018.
182. Pons E, et al. Natural language processing in radiology: a systematic review. Radiology. 2016;279(2):329–43.
183. Zech J, et al. Natural language–based machine learning models for the annotation of clinical radiology reports. Radiology. 2018;287(2):570–80.
184. Goff DJ, Loehfelm TW. Automated radiology report summarization using an open-source natural language processing pipeline. J Digit Imaging. 2018;31(2):185–92.
185. Folio LR, et al. Quantitative radiology reporting in oncology: survey of oncologists and radiologists. Am J Roentgenol. 2015;205(3):233–43.
186. Schlegl T, et al. Predicting semantic descriptions from medical images with convolutional neural networks. In: International conference on information processing in medical imaging. New York: Springer; 2015.
187. Lee YH. Efficiency improvement in a busy radiology practice: determination of musculoskeletal magnetic resonance imaging protocol using deep-learning convolutional neural networks. J Digit Imaging. 2018;31(5):604–10.
188. Annarumma M, et al. Automated triaging of adult chest radiographs with deep artificial neural networks. Radiology. 2019;291(1):196–202.
189. Miotto R, et al. Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. Sci Rep. 2016;6(1):1–10.
190. Lambin P, et al. Radiomics: extracting more information from medical images using advanced feature analysis. Eur J Cancer. 2012;48(4):441–6.
191. Aerts HJ. The potential of radiomic-based phenotyping in precision medicine: a review. JAMA Oncol. 2016;2(12):1636–42.
192. Kumar V, et al. Radiomics: the process and the challenges. Magn Reson Imaging. 2012;30(9):1234–48.
193. Coroller TP, et al. CT-based radiomic signature predicts distant metastasis in lung adenocarcinoma. Radiother Oncol. 2015;114(3):345–50.
194. Huynh E, et al. Associations of radiomic data extracted from static and respiratory-gated CT scans with disease recurrence in lung cancer patients treated with SBRT. PLoS One. 2017;12(1):e0169172.
195. Parmar C, et al. Machine learning methods for quantitative radiomic biomarkers. Sci Rep. 2015;5:13087.