Prospects of deep learning for medical imaging

Jonghoon Kim[1,2], Jisu Hong[1,2], Hyunjin Park[2,3]

[1]Department of Electronic, Electrical and Computer Engineering, Sungkyunkwan University, Suwon, Korea
[2]Center for Neuroscience Imaging Research, Institute for Basic Science (IBS), Suwon, Korea
[3]School of Electronic and Electrical Engineering, Sungkyunkwan University, Suwon, Korea

Corresponding author:
Hyunjin Park
School of Electronic and
Electrical Engineering,
Sungkyunkwan University, 2066
Seobu-ro, Jangan-gu, Suwon
16419, Korea
Tel: +82-31-299-4956
E-mail: hyunjinp@skku.edu

ABSTRACT

Machine learning techniques are essential components of medical imaging research. Recently, a highly flexible machine learning approach known as deep learning has emerged as a disruptive technology to enhance the performance of existing machine learning techniques and to solve previously intractable problems. Medical imaging has been identified as one of the key research fields where deep learning can contribute significantly. This review article aims to survey deep learning literature in medical imaging and describe its potential for future medical imaging research. First, an overview of how traditional machine learning evolved to deep learning is provided. Second, a survey of the application of deep learning in medical imaging research is given. Third, well-known software tools for deep learning are reviewed. Finally, conclusions with limitations and future directions of deep learning in medical imaging are provided.

Keywords: Deep learning; Diagnostic imaging; Machine learning

INTRODUCTION

Machine learning techniques have been widely used in medical imaging research in the form of many successful classifier and clustering algorithms [1,2]. Many clinicians are well aware of the effectiveness of classifiers, such as support vector machine (SVM), and clustering algorithms, such as k-nearest neighbor (k-NN) [3]. Recently, deep learning (DL) has emerged as the go-to methodology to drastically enhance the performance of existing machine learning techniques and to solve previously intractable problems. In addition, DL is a generic methodology that has a disruptive impact in other scientific fields as well. Thus, it has become imperative for medical imaging researchers to fully embrace DL technology. This review article is borne out of that necessity.

 Medical image processing refers to a set of procedures to obtain clinically meaningful information from various imaging modalities, mostly for diagnosis or prognosis. The modalities are typically *in vivo* types, but *ex vivo* imaging could be used for medical image processing as well. The extracted information could be used to enhance diagnosis and prognosis according to the patient's needs. Distinct medical imaging modalities, such as magnetic resonance imaging (MRI), computed tomography (CT), and positron emission tomography (PET), could provide

distinct information for the patient being imaged. Structural and functional information could be extracted as necessary, and these are used as quantitative features for future diagnosis and prognosis. Research in medical image processing typically aims to extract features that might be difficult to assess with the naked eye. There are two types of features. The first is the well-known semantic feature defined by human experts, and the other is the agonistic feature defined by mathematical equations [4]. The agonistic features suffer from less operator bias than the semantic features. Still, semantic features are well recognized in radiology research, which is an accumulation of years of human expertise. However, many semantic features are time-consuming to compute and sometimes there are inconsistencies among experts. The extracted agonistic features might be used as imaging biomarkers to explain various states of the patient. A recent research approach known as the "radiomics" approach employs hundreds or thousands of agnostic features to obtain clinically relevant information for diagnosis and prognosis [5,6].

Machine learning approaches are applied to associate imaging features obtained from medical image processing with relevant clinical information. Machine learning started as a field in computer science to endow algorithms to solve problems without being explicitly programmed. It typically learns representations from training data, which are generalized in separate test data. The technology has been applied to computer-aided diagnosis (CADx) and computer-aided detection (CADe) in medical imaging [7,8]. The CADx system can identify a disease-related region and quantify the properties of that region, which could be used as a guidance for surgeons and radiologists. Despite their usefulness in analyzing medical imaging, machine learning approaches have several limitations. They show excellent performance when applied to training data but typically suffer losses in performances when applied to independent validation data [9]. This is partly due to the overfitting of the training data. Performance of machine learning techniques must be evaluated with both training and independent validation data. Many machine learning studies have demonstrated great technical potential but only a few have shown actual clinical efficacy including gains in survival. Machine learning techniques also have issues related to feature definition. For example, they rely on a pre-defined set of features. Additionally, sometimes the features are difficult to define for a given problem. Researchers need to choose from different combinations of features, algorithms, and degrees of complexity to sufficiently solve a given problem, and many studies depend on trial-and-error to find the right combination. A major challenge in particular is choosing the right features to correctly model a given problem.

The development of an artificial neural network (ANN) largely circumvents this problem by learning feature representation directly from the raw input data, skipping the feature extraction procedure [10]. ANN attempts to mimic human brain processes using a network composed of interconnected nodes. The initial ANN model was a simple feed-forward network known as perceptron that could perform linear classification tasks [11]. The early perceptron models required complex computation power beyond what was typically available at that time. Multilayer perceptron (MLP) was proposed to improve the simple perceptron model by adding hidden layers and developing learning techniques, such as back-propagation [12]. MLP formed the basis for the modern DL approaches. The "deep" portion of DL refers to having many layers whose structures are suitable to model big data. On a practical side, the DL approach requires a high computational load. With the recent developments of computational infrastructure, such as graphical processing units (GPUs) and cloud computing systems, DL has become practical and has achieved groundbreaking results in many fields. In particular, the visual recognition challenge based on large-scale data (ImageNet) performs the task of classifying among 1,000 objects leveraging 1,200,000 training and 100,000 test images. Since the initiation of the challenge in 2010, the performance of DL algorithms gradually increased, and it began to exceed human accuracy beginning in 2015 (Fig. 1).

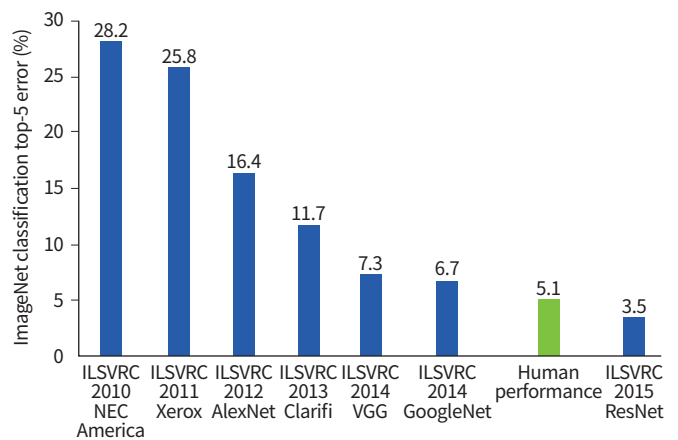DL has been successfully applied to many research fields

Fig. 1. ImageNet classification performance in chronological order. ILSVRC, ImageNet Large Scale Visual Recognition Challenge.

and has become ubiquitous. As such, many studies have already adopted DL to improve medical imaging research, and increasingly more studies will adopt DL for medical imaging research in the future. This review article aims to survey DL literature in medical imaging and describe its potential for future medical imaging research. First, we provide an overview of how traditional machine learning evolved to DL. Second, we survey the application of DL in medical imaging research. Third, software tools for DL are reviewed. Finally, we conclude with limitations and future directions of DL in medical imaging.

FROM TRADITIONAL MACHINE LEARNING TO DEEP LEARNING

Machine learning involves building data-driven models to solve research problems [13]. There are two categories: supervised learning and un-supervised learning. In supervised learning, we train models using input data with matched labels [14]. The model is a mathematical model that can associate input data with the matched labels and a predictive model, which is validated using unseen test data. For example, we start with MRI images labeled either as a normal control or diseased, and machine learning would lead to a mathematical model that could associate MRI images with diagnosis in both training and test data. Supervised learning is commonly used in the following two tasks. In classification,

the model associates input data with pre-defined categorical results (i.e., normal vs. diseased) [15-17]. The output is a discrete categorical variable in classification. In regression, the model associates input data with often continuous results (i.e., the degree of symptoms) [18]. The output is typically a continuous variable in the regression. In un-supervised learning, we use unlabeled input data to learn intrinsic patterns in the input data. Un-supervised learning is commonly used in clustering. With clustering, we might identify subgroups within a given group of patients diagnosed with the same disease. Supervised learning is more costly to prepare, as it involves annotating input data with labels, which often requires human intervention.

ANN is a statistical machine learning method inspired by brain mechanism from neuroscience (Fig. 2) [19]. Researchers designed a learning algorithm that resembles how the brain handles information. A neuron is the basic unit of the brain mechanism. The neuron is an electrically excitable cell that receives signals from other neurons, processes the received information, and transmits electrical and chemical signals to other neurons. The input signal to a given neuron needs to exceed a certain threshold for it be activated and further transmit a signal. The neurons are interconnected and form a network that collectively steers the brain mechanism. ANN is an abstraction of an interconnected network of neurons with layers of nodes, and it consists of an input layer aggregating the input signal from other connected neurons, a hidden layer responsible for training, and an output layer [20]. Each node takes the input from nodes from the previous layer using various weights and computes the activation function, which is relayed onto the next layer of nodes. The activation function approximates the complex process of a physical neuron, which regulates the strength of the neuronal output in a non-linear manner. The mathematical processing in a node can be represented using the following equation:

$$\text{Output} = \varphi(W^T x + b).$$

A node takes in an input value 'x' and multiplies it by weight 'W,' and then a bias of 'b' is added, which is fed to the activation function 'φ.' Differences between the output of ANN and the target value (i.e., ground truth) are called errors or losses. The training of the ANN is a procedure to update weights that interconnect different nodes to explain the training data, which is equivalent to minimizing the loss value. The losses are back-propagated through the network, and they are used to modify the weights and biases [21,22]. As a result, components of ANN are updated to best explain the target values.
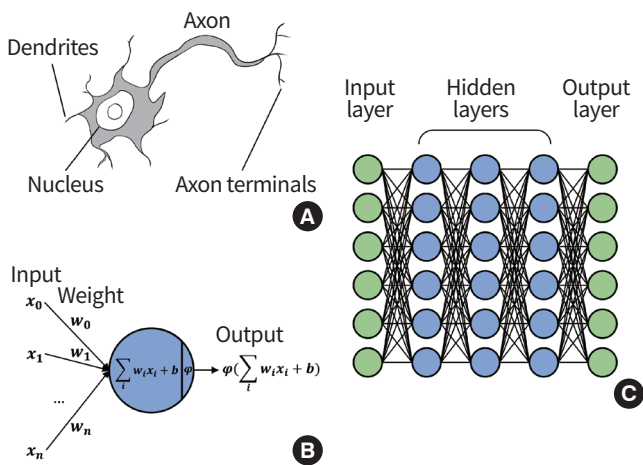

Fig. 2. Overview of artificial neural network. (A) In a biological neuron, the nucleus transforms the chemical inputs from the dendrites into electric signals. The signals are transmitted to the next neurons through the axon terminals. (B) In a node (perceptron), the input values are transformed by the weights, biases, and activation functions. The output values are transmitted to the next perceptron. (C) Multilayer perceptron consists of multiple perceptrons.

A deep neural network (DNN) is an expanded ANN incorporating many hidden layers to increase the flexibility of the model [23-26]. DNN consists of an input layer, many hidden layers, and an output layer. The added layers afford solving more complex problems but can cause other issues. The first issue is the vanishing gradient problem [27]. This occurs because the input information cannot be effectively used to update the weights of deep layers, as the information needs to travel the many layers of the activation function. This leads to under-training of the deep layers. Another issue is over-fitting, which occurs frequently when the complexity of the model increases [25]. A complex model might explain the training data well, but it does not necessarily explain the unseen test data well. The increased complexity of many layers requires a high computational load and thus degrades the training efficiency. The issues of vanishing gradient and over-fitting have been mitigated by enhanced activation function, cost function design, and drop-out approaches. The issue of the high computational load has been dealt with by using highly parallel hardware, such as GPUs and batch normalization.

Modern DL approaches employ dedicated topology of deeply stacked layers to solve a given problem. There are many variants of DL architecture, and an exhaustive review of them is outside the scope this paper. In this paper, we mention three well-known DL architectures: convolutional neural network (CNN), auto-encoder (AE), and recurrent neural network (RNN).

Convolutional neural network

CNN refers to a network architecture composed of many stacked convolutional layers [28]. Convolution is a mathematical operation based on two functions, addition and multiplication. For a given input image, convolution is applied to the input image based on a receptive filed (i.e., extents of convolution operation), which is analogous to the extents of the response region of human vision. The convolution procedure is well-suited for image recognition because it considers locally connected information (i.e., neighboring voxels or pixels). There are pooling layers between convolution layers. The pooling layer is essential to increasing the field of view of the network. It takes a portion of the locally connected nodes of the input layer and results in an output that has a smaller spatial footprint. For example, we might take the maximum value of the outputs of four neighboring nodes. It can control the over-fitting problems and reduce the computational re-
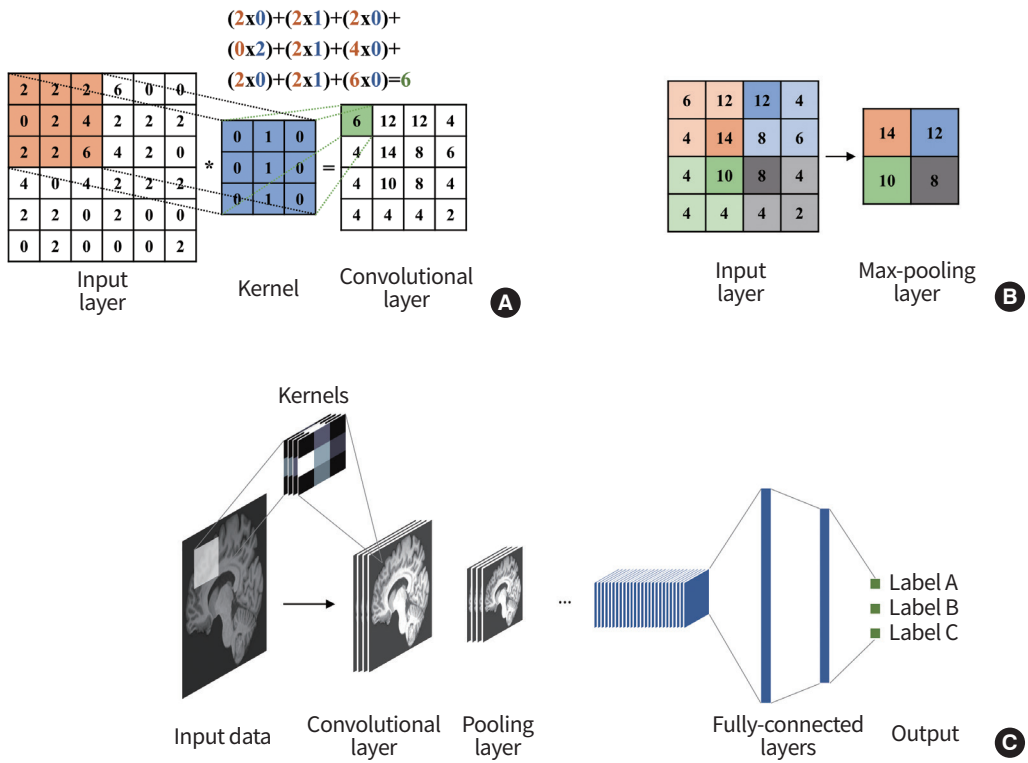
Fig. 3. Overview of convolutional neural network (CNN). (A) The output of the convolutional layer is obtained by convolving its input layer and kernel weights. (B) The max-pooling layer outputs a down-sampled layer by extracting the maximum values from the partially windowed area of the input layer. (C) A general CNN consists of a series of convolutional layers, pooling layers, and fully connected layers.

quirements. In general, the max-pooling algorithm is used for many architectures, extracting the highest values from the receptive sliding window (Fig. 3).

As the last component of CNN, a series of fully connected layers follow. The fully connected layer combines all activations of the previous layers. This layer works the same as in the general ANN. After the series of fully connected layers, the model outputs the final set of feature values appropriate for the given problem. CNN was the initial catalyst for widespread adoption of DL. LeNet-5 is one of the early architectures that drastically improved existing machine learning approaches for recognizing hand-written digits [29], and there are other well-known architectures, including AlexNet, GoogLeNet, VGG-Net, and ResNET [28,30-32]. Further details of the architecture are found in the references.

Auto-encoder and stacked auto-encoder

AE is a network that is used to derive an efficient representation of the input data [33,34]. As the name suggests, this network reduces the input layers to a layer with a small number of nodes (i.e., encoding) and then restores them via the upscaling process (i.e., decoding), as shown in Fig. 4. By this

process, the hidden layer can represent its input layer with a reduced dimension that minimizes the loss of input information. This model generally has an hourglass-shaped structure with a dimension of the hidden layer being smaller than the input layer.

Stacked auto-encoder (SAE) is made by stacking the AEs so that the output of an AE layer is treated as an input layer for another AE (Fig. 4) [35]. Vincent et al. [36] used this SAE structure to obtain a denoised image of input data. The model used the greedy training method to learn several SAEs independently.

Recurrent neural network

General ANN is a feed-forward network where input information starts from the input layer, travels through many hidden layers, and finally arrives at the output layer. RNN has a different topology for its layer configuration. The output of hidden layers is fed back to an input layer using feedback (Fig. 5) [37,38]. RNN considers both current input data and feedback data from previous states and thus is well-suited for modeling sequential data, including both temporal and spatial information.

DEEP LEARNING APPLICATION IN MEDICAL IMAGING

In this section, we review representative applications of DL in medical imaging. First, CADx/CADe tasks are covered with topics of classification, detection, and prediction. Second, image processing tasks are covered with topics of segmentation, registration, and generation. CADx/CADe has many critical usages in radiology, such as finding an altered structural/functional area (detection or localization), predicting the state of the object of interest based on a probabilistic model (prediction), and classifying binary or multiclass categories (classification). In contrast, image processing is the prerequisite for subsequent clinical tasks, such as diagnosis or prog-
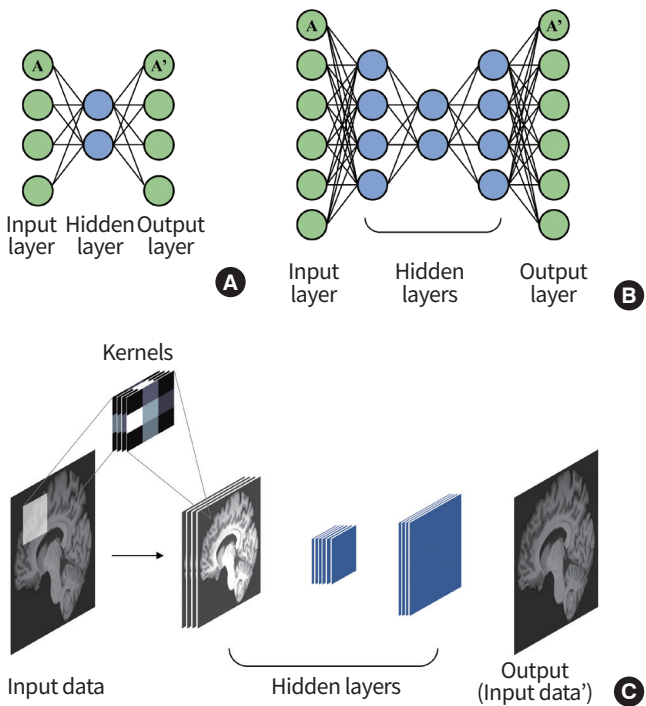

Fig. 4. Overview of auto-encoder (AE). (A) An AE consists of multiple perceptrons that mimic the input layer. (B) A stacked auto-encoder (SAE) can be built by stacking the AEs. (C) A general SAE is trained to reconstruct the output that is similar to the given input.
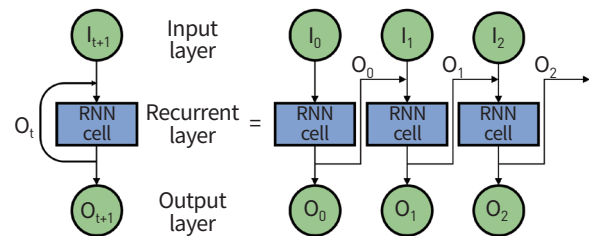

Fig. 5. A recurrent neural network (RNN) cell calculates the output by using its sequential inputs recurrently.

nosis, and they can be categorized into the following tasks: a segmentation task assigns voxels with similar properties into labels, a registration task spatially aligns two or more images onto a common spatial frame, and an enhancement task improves the image quality for the given task. Each topic is described in terms of the network architecture (e.g., CNN, RNN, or SAE) and its clinical application.

Applications using CNN architecture

CNN computes features from locally connected voxels (i.e., neighborhood voxels); hence, it is widely used as the reference model for medical image analysis. In general, a given image goes through many convolution layers, and the features are extracted from different layers. The extracted features are further used for various tasks.

CADx/CADe application using CNN architecture

CADx/CADe applications based on CNN are shown in Table 1.

Detection: The detection task refers to finding the region associated with the diseased condition. The detection task is sometimes referred to as the localization task, as it involves finding voxel positions associated with the disease. The task is

often a pre-requisite for therapy planning. Many studies adopted CNN and reported precise detection performance. Dou et al. [39] applied three-dimensional (3D) CNNs to localize the lesion of cerebral microbleeds from MRI. Zreik et al. [40] proposed an automatic detection method for identifying left ventricle myocardium in coronary CT angiography using CNN. Using both PET/CT imaging, whole body multiple myeloma detection showed better performance than traditional machine learning approaches, such as the random forest classifier, k-NNs, and SVM [41]. Apart from the common imaging modalities (i.e., MRI, CT, and PET), retinal and endoscopy images, as well as histology images, were analyzed for finding the diseased region of interest (ROI) in other studies [17,42,43].

Prediction: The prediction task involves forecasting properties (often not evident to the naked eye) of an object using imaging analysis results. Prediction often occurs in a longitudinal setting, where baseline imaging findings are used to predict properties of the object in the future. Many studies attempted to compute the likelihood of clinical variables, such as drug or therapeutic response, patient's survival, and disease grading, using imaging analysis results. Kawahara et al. [44] proposed a novel CNN framework, BrainNetCNN, to

Table 1. CADx/CADe applications using convolutional neural network architecture

Task	Modalities	Object	Clinical goal
Detection	MRI	Brain	Lesion detection [39]
	CT	Heart	Disease detection [40]
	PET/CT	Bone	Lesion detection [41]
	Retina image	Retina	Lesion detection [17]
	Endoscopic image	Stomach	Disease detection [42]
	Histology image	Breast	Lesion detection [43]
Prediction	MRI	Brain	Disease prediction [44]
	MRI	Liver	Disease prediction [45]
	PET	Esophagus	Treatment response prediction [46]
	OCT	Retina	Treatment response prediction [47]
	MRI	Brain	Survival prediction [48]
	CT	Lung	Survival prediction [49]
Classification	CT	Brain	Disease classification [50]
	CT	Lung	Disease classification [51]
	CT	Liver	Lesion classification [52]
	Mammography	Breast	Lesion classification [53]
	RGB	Skin	Disease classification [54]
	Fluorescent image	Cell	Cell classification [55]

CADx, computer-aided diagnosis; CADe, computer-aided detection; MRI, magnetic resonance imaging; CT, computed tomography; PET, positron emission tomography; OCT, optical coherence tomography; RGB, red green blue.

predict a cognitive and motor developmental outcome score by using structural brain networks of diffusion tensor imaging. Yasaka et al. [45] assessed the performance for predicting the staging of liver fibrosis using a CNN model on gadoxetic acid-enhanced hepatobiliary phase MRI. In the above two published papers, the predicted scores derived from CNN correlated well with the ground-truth. Some studies adopted CNN to predict treatment response [46,47]. Furthermore, CNN-based approaches have shown potential in the prediction of a patient's survival and thus could be promising tools for precision medicine [48,49].

Classification: The classification task determines which classes the imaging data belong to. The classes could be binary (e.g., normal and diseased) or multiple categories (e.g., subclasses within the given diseased condition) depending on the task. Many studies successfully applied CNN to classify the severity of disease for different organs (i.e., brain, lung, and liver) using CT imaging [50-52]. Mohamed et al. [53] applied CNN to mammogram imaging and categorized data into a few classes proportional to breast density. Esteva et al. [54] classified skin cancer photographs into binary classes and even suggested that such operation could be performed using portable devices such as smartphones. In histology, one study showed that cell differentiation could be classified with CNN, and the results have the potential to identify tissue

or organ regeneration [55].

Image processing applications using CNN architecture
Image processing applications based on CNN are shown in Table 2.

Image segmentation: The segmentation task partitions the imaging data into the target and background region. In medical imaging, this typically amounts to finding regions that are related to a disease condition (i.e., finding the tumor region). The task assigns voxels to binary classes using intensity, texture, and other derived information. Many studies successfully used CNN to segment the target region using CT and MRI for various organs, such as the brain, liver, kidney, and prostate [56-61]. Fang et al. [62] proposed an automatic segmentation algorithm of retina layer boundaries using CNN for high-resolution optical coherence tomography imaging. Xu et al. [63] adopted CNN architecture for segmenting epithelial and stromal regions in histology images.

Image registration: The image registration task spatially aligns one image with another so that both can be compared to a common spatial framework through a geometric transform. Many studies adopt multimodal imaging and thus registration is necessary to extract spatially aligned features in those studies. Registration requires a long computational time as it involves a high degree of freedom (DOF) optimiza-

Table 2. Image processing applications using convolutional neural network architecture

Task	Modalities	Object	Clinical goal
Segmentation	MRI	Brain	Tissue segmentation [56]
	MRI	Prostate	Organ segmentation [59]
	CT	Liver	Organ segmentation [57]
	CT	Bladder	Organ segmentation [58]
	MRI	Heart	Lesion segmentation [60]
	CT	Kidney	Organ segmentation [61]
	OCT	Retina	Tissue segmentation [62]
	Histology image	Cell	Tissue segmentation [63]
Registration	X-ray	Knee	Organ registration [64]
	MRI	Abdomen	Region registration [65]
	MRI	Brain	Organ registration [66]
Enhancement, generation, and reconstruction	X-ray	Brain	Image enhancement [67]
	MRI	CT	Image generation [68]
	MRI	CT	Image generation [69]
	Diffusion MRI	Diffusion MRI	Image enhancement [70]
	MRI (3T)	MRI (7T)	Image reconstruction [71]

MRI, magnetic resonance imaging; CT, computed tomography; OCT, optical coherence tomography.

tion problem. The CNN-based approaches could be used as an alternative for the difficult registration problem by finding landmarks or control points through CNN architecture. Miao et al. [64] improved the registration in terms of the computation time and capture range using a real-time 2D/3D registration framework based on CNN. Another study proposed a technique for correcting respiratory motion during free-breathing MRI using CNN [65]. Yang et al. [66] introduced a novel registration framework based on CNN, Quicksilver, which performed fast predictive image registration.

Image enhancement: The image enhancement task aims to improve the image quality of the objects of interest. Depending on the application, enhancement might occur in the context of image generation or reconstruction. Nonetheless, the goal is to improve the image quality. Two studies proposed algorithms to generate a high-quality image (e.g., high dose CT) from a low quality image (e.g., low dose CT) using CNN, which reduced noise and provided more structural information than the low-quality images [67]. Image generation based on the synthesis of different modalities contributed to the improved image quality by combining the distinctive information in each modality [68,69]. Golkov et al. [70] proposed a reconstruction model based on CNN for shorter scanning time in diffusion MRI. This model allowed acceleration of the scan time and obtaining quantitative diffusion MRI. Bahrami et al. [71] explored the generation of 7 Tesla (T)-like MRI from routine 3T MRI based on CNN, where reconstructed images benefited from both resolution and contrast compared to the 3T MRI.

Applications using RNN architecture

RNN integrates results from the previous states and current input and thus is suitable for tasks involving sequential or dynamic data. In medical image analysis, spatial information is exploited using CNN, and temporal information is often exploited using RNN. Studies using dynamic imaging adopted a combination of CNN and RNN so that joint modeling of spatial and temporal information was possible references [72-77].

CAD application using RNN architecture
CADx/CADe applications based on RNN are shown in Table 3.

Detection: Xu et al. [72] proposed an end-to-end DL framework to accurately detect myocardial infarction at the pixel level. In this framework, the location of the heart was found based on CNN and the detected ROI was further analyzed using the dynamic cardiac MRI incorporating temporal motion information using RNN. In the same manner, Chen et al. [73] used the combined CNN and RNN for a fetal ultrasound, where ROI was found using CNN, and temporal information was processed by RNN based on the features of the ROIs in consecutive frames.

Prediction: Han and Kamdar [74] proposed convolutional recurrent neural networks (CRNN) for predicting methylation status in glioblastoma patients. In the CRNN architecture, the convolution layer extracted sequential feature vectors from each MRI slice and RNN was used to predict the methylation status integrating sequential feature information.

Classification: Gao et al. [75] designed a combined architecture of CNN and RNN for classifying the severity of nuclear cataracts in slit-lamp images, where hierarchical feature sets were extracted from CNN, and then the features were hierarchically merged into high-order image-level features. Xue et al. [76] proposed a deep multitask relationship learning network for quantifying the left ventricle in cardiac imaging, where cardiac sequence features were extracted using CNN and then temporally learned by RNN.

Image processing applications using RNN architecture
Imaging processing applications based on RNN are shown in Table 4.

Table 3. CADx/CADe applications using recurrent neural network architecture

Task	Modalities	Object	Clinical goal
Detection	MRI	Heart	Lesion detection [72]
	US	Infancy	Object detection [73]
Prediction	MRI	Brain	Status prediction [74]
Classification	OCT	Eye	Disease classification [75]
	MRI	Heart	Lesion classification [76]

CADx, computer-aided diagnosis; CADe, computer-aided detection; MRI, magnetic resonance imaging; US, ultrasound; OCT, optical coherence tomography.

Table 4. Image processing applications using recurrent neural network architecture

Task	Modalities	Object	Clinical goal
Segmentation	MRI	Brain	Lesion segmentation [77]
	US	Prostate	Organ segmentation [78]
Enhancement, generation, and reconstruction	MRI	MRI	Image reconstruction [79]

MRI, magnetic resonance imaging; US, ultrasound.

Image segmentation: Zhao et al. [77] proposed a tumor segmentation method based on DL using multimodal brain MRI. They used a network composed of the combination of CNN and RNN architectures in which CNNs assigned the segmentation label to each pixel, and then RNNs optimized the segmentation result using both assigned labels and input images. Yang et al. [78] proposed an automatic prostate segmentation method using only the RNN architecture. The static ultrasound image was transformed into an interpretable sequence and then the extracted sequential data was used as input to the RNN.

Image enhancement: Qin et al. [79] suggested a framework to reconstruct high-quality cardiac MRI images from under-sampled k-space data. They adopted CRNN to improve reconstruction accuracy and speed by considering both spatial and temporal information.

Applications using SAE architecture

SAE encodes the input data using a small number of features and then decodes them back to the dimension of the original input data. In medical imaging, SAEs are commonly used to identify a compact set of features to represent the input data. The learned features could be used for specific image enhancement and reconstruction tasks.

CAD application using SAE architecture

CADx/CADe applications based on SAE are shown in Table 5.

Detection: One study showed that a stacked sparse AE could be used for nuclei detection in breast cancer histopathology images [80]. This architecture could capture high-level feature representations of pixel intensity and thus enabled the classification task to be performed effectively for differentiating multiple nuclei.

Prediction: He et al. [81] adopted a variant of SAE to predict cognitive deficits in preterm infants using functional MRI data.

Classification: Chen et al. [82] proposed an unsupervised learning framework based on the convolutional AE neural network in which the patch images from raw CT were used for feature representation and classification for lung nodules. Cheng et al. [83] applied a stacked denoising AE to two different modalities, ultrasound and CT. The aim was to classify benign and malignant lesions from the two modalities.

Image processing applications using SAE architecture

Image processing applications based on SAE are shown in Table 6.

Image segmentation: Zhao et al. [84] applied SAE to learn compact feature representation from cryosection brain imaging and then classified voxels into white matter, gray matter, and cerebrospinal fluid.

Image enhancement: Janowczyk et al. [85] applied sparse AE to reduce the effects of color variation in histology images. Benou et al. [86] adopted SAE to model the spatiotemporal structure in dynamic MRI, which led to denoised MRI. The denoised MRI resulted in a more robust estimation of pharmacokinetic parameters for blood-brain barrier quantification. Some adopted SAE to generate pseudo-CT scans from MRI, which were used to reconstruct PET imaging with better tissue characteristics [87].

Among the discussed three network architectures, CNN is the most widely used. Many studies adopted CNN for various target organs (brain, lung, breast, etc.) using various imaging modalities (MRI, CT, etc.). This is partly due to the fact that CNN processes information in a hierarchical manner similar to visual processing in the human brain. RNN is typically used in conjunction with CNN, where the learned CNN features were sequentially handled through the RNN. SAE has strengths in compact feature representation, which is useful in denoising and reconstruction tasks. In summary, there is no one network structure that solves all the medical imaging

Table 5. CADx/CADe applications using stacked auto-encoder architecture

Task	Modalities	Object	Clinical goal
Detection	Histology image	Breast	Lesion detection [80]
Prediction	MRI	Brain	Risk prediction [81]
Classification	CT	Lung	Lesion classification [82]
	US, CT	Breast	Lesion classification [83]

CADx, computer-aided diagnosis; CADe, computer-aided detection; MRI, magnetic resonance imaging; CT, computed tomography; US, ultrasound.

Table 6. Image processing application using stacked auto-encoder architecture

Task	Modalities	Object	Clinical goal
Segmentation	RGB	Brain	Tissue segmentation [84]
Enhancement, generation, and reconstruction	Histology image	Tissue	Image enhancement [85]
	MRI	Brain	Image enhancement [86]
	MRI	CT	Image generation [87]

RGB, red green blue; MRI, magnetic resonance imaging; CT, computed tomography.

problems, and thus, researchers should choose the appropriate network architecture suitable for a given problem.

SOFTWARE TOOLS FOR DEEP LEARNING

Researchers need dedicated software tools to perform DL research. Fortunately, much of the DL software is open-source; thus, the studies are more reproducible and cost effective. We review commonly used open-source DL software tools in this section.

Caffe is one of the early DL tools developed by Berkeley Vision and Learning Center [87]. The tool emphasized modularity and speed based on C++ and Python. Many early CNN works have been performed with Caffe. Researchers could easily transfer the learned models (i.e., layer structure and weights) from other research projects and use them to initialize their own DL model. Such procedures are often referred to as transfer learning, and they typically save computation time in training the model. Theano is another early DL tool [88]. Theano is Python-based and is efficient in computation using a GPU and central processing unit together.

Tensorflow is one of the most widely used tools for DL and it is backed by the internet search giant Google [89]. The Google Brain Team developed the software and it is based on Python. Users are free to leverage the vast software library available in Python. The software is well-maintained and updated frequently for additional functionality.

Torch is a DL tool that aims to facilitate the procedures to build complex models. It was originally based on non-Python Lua, but recently it added Python support via PyTorch for improved user friendliness [90].

LIMITATIONS AND FUTURE DIRECTIONS

Modern DL methods learn very high DOF models (i.e., millions of weights) from the training data. Hence, the sheer number of required training data is very high compared to conventional machine learning methods. Recent DL applications in brain MRI learned models from more than 1.2 million training data [91]. There are algorithm enhancements, such as augmentation, to artificially boost the number of training samples, but the quality of the DL methods directly rely on the number and quality of training samples. This is one of the biggest hurdles of DL research in medical imaging and is one of the reasons why large corporations such as Google can produce high-quality DL models, as they already have a huge number of training samples in-house. In medical imaging research, the large sample requirement could be partly alleviated by multisite data acquisition. Researchers need to apply a standardized protocol to acquire data so that they can be effectively combined into one set of training data for DL. Another way to boost training samples is to use an open research database. The Human Connectome Project houses thousands of high-quality brain MRI that are open to the research community. A potential issue with mixing data from different sites or research databases is the heterogeneity in image quality. One can render high-quality data into low-quality data so that all data are of similar quality. This is the practical approach because rendering low-quality data into high-quality data is difficult.

Many DL methods belong to the supervised approach; hence, they require manual labeling. Labeling thousands of imaging data by human experts is cumbersome. On top of that, inter- and intraobserver variability need to be considered, which makes the labeling procedure even more problematic. One possible way to tackle this issue is to apply a two-tiered semiautomatic approach [43]. An automatic algorithm would perform the first round of labeling and then the human experts can either accept or modify the results of the first round. Recently, one study proposed a DL method to automatically retrieve images from a large database that matched human set criterion [92-94]. This approach has the potential to be used as the initial labeling for further refinement.

Another issue related to sample size is the lack of balance in comparison groups. Many studies compare two groups: the normal control group and the diseased group. Many clinical sites have many samples of the disease group, while matched controls could be lacking. Ideally, DL algorithms require an equal number of samples in the comparison groups. If there is a large imbalance between the comparison groups, the DL algorithm would not be able to fully learn the under-represented group. A good practice is to prospectively plan research projects to avoid imbalance in comparison groups as much as possible. Many diseased conditions we explored are in fact quite heterogeneous. A given disease condition might consists of many subtypes, and more studies are aiming to dissect the differences among the subtypes. Researchers should make sure that they have appropriate labels for the subtypes and enough samples of each subtype so that DL models could be applied effectively.

DL is a highly flexible modeling approach to learn an inherent representation of the input data. It optimizes a loss function to find millions of weights that can best explain the in-

put training data. It is very difficult to interpret how certain weights contribute to the final model. In many cases, we are left with a black box that performs the given task really well [95,96]. In traditional machine learning, we could quantify how each semantic feature contributes to the final model and better understand or improve the model as necessary. For example, we know that an irregular tumor boundary is strongly linked with malignant tumors [97,98]. However, for DL, we cannot say the n-th weight in the j-th layer has a strong link with the final outcome. Such interpretability is largely lacking in DL. Ribeiro et al. [99] proposed a novel explanation technique, local interpretable model-agnostic explanations, which performs a local approximation of the complex DL model. This algorithm carries out linear classification based on distinctive features from the simplified model, which might be interpretable. Some studies surveyed visual interpretability of various DL architectures and showed that the DL model is reliable across many domains [100-102]. There is active research tackling this interpretability issue and researchers should pay attention to future developments.

A DL algorithm requires many tuning parameters to properly train the model. The hyperparameters include learning rate, dropout rate, and kernel functions. A slight modification of these parameters might lead to drastically different models with varying performances. So far, these parameters have been chosen based on heuristic approaches relying on the experience of human experts. Active research is being performed regarding methods to automatically find optimal hyperparameters for DL. Domhan et al. [103] developed a method for automatically optimizing the hyperparameters, where the process was based on the inference of performance using a learning curve from the first few iterations. They would terminate the model, leading to bad performance based on the few initial data points. Shen et al. [104] proposed an intelligent system, including a parameter-tuning policy network, which determines the direction and magnitude for fine-tuning hyperparameters.

CONCLUSION

DL is already widespread, and it will continue to grow in the near future in all fields of science. The advent of DL is fostering new interactions among different fields. Experts in machine learning (mostly from computer science) are actively embedded in research teams to solve critical medical problems. Medical image processing will benefit immensely from DL approaches, as DL has shown remarkable performance in non-medical regular imaging research compared to conventional machine learning approaches. In this review paper, we touched on a brief history from conventional machine learning to DL, surveyed many DL applications in medical imaging, and concluded with limitations and future directions of DL in medical imaging. Despite the limitations, the advantages of DL far outweigh its shortcomings, and thus, it will be an essential tool for diagnosis and prognosis in the era of precision medicine. Future research teams in medical imaging should integrate DL experts in addition to clinical scientists in their teams so as to fully harness the potential of DL.

CONFLICTS OF INTEREST

ACKNOWLEDGMENTS

REFERENCES

1. Wang S, Summers RM. Machine learning and radiology. Med Image Anal 2012;16:933-51.
2. de Bruijne M. Machine learning approaches in medical image analysis: from detection to diagnosis. Med Image Anal 2016;33:94-7.
3. Wernick MN, Yang Y, Brankov JG, Yourganov G, Strother SC. Machine learning in medical imaging. IEEE Signal Process Mag 2010;27:25-38.
4. Gillies RJ, Kinahan PE, Hricak H. Radiomics: images are more than pictures, they are data. Radiology 2016;278: 563-77.
5. Aerts HJ, Velazquez ER, Leijenaar RT, Parmar C, Grossmann P, Carvalho S, et al. Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. Nat Commun 2014;5:4006.
6. Lambin P, Leijenaar RTH, Deist TM, Peerlings J, de Jong EEC, van Timmeren J, et al. Radiomics: the bridge between medical imaging and personalized medicine. Nat Rev Clin Oncol 2017;14:749-62.
7. van Ginneken B, Schaefer-Prokop CM, Prokop M. Computer-aided diagnosis: how to move from the laborato-

ry to the clinic. Radiology 2011;261:719-32.

8. Doi K. Computer-aided diagnosis in medical imaging: historical review, current status and future potential. Comput Med Imaging Graph 2007;31:198-211.

9. Park SH, Han K. Methodologic guide for evaluating clinical performance and effect of artificial intelligence technology for medical diagnosis and prediction. Radiology 2018;286:800-9.

10. McCulloch WS, Pitts W. A logical calculus of the ideas immanent in nervous activity. Bull Math Biophys 1943;5:115-33.

11. Van Der Malsburg C. Frank Rosenblatt: principles of neurodynamics: perceptrons and the theory of brain mechanisms. In: Palm G, Aertsen A, editors. Brain theory; 1984 Oct 1-4; Trieste (IT): Springer Berlin Heidelberg; 1986. p. 245-8.

12. Rumelhart DE, Hinton GE, Williams RJ. Learning internal representations by error propagation. In: Anderson JA, Rosenfeld E, editors. Neurocomputing: foundations of research. Cambridge (MA): MIT Press; 1988. p. 696-9.

13. Nasrabadi NM. Pattern recognition and machine learning. J Electron Imaging 2007;16:49901.

14. Kotsiantis SB. Supervised machine learning: a review of classification techniques. Informatica 2007;31:249-68.

15. Kuang D, Guo X, An X, Zhao Y, He L. Discrimination of ADHD based on fMRI data with deep belief network. In: Huang DS, Han K, Gromiha M, editors. Intelligent computing in bioinformatics; 2014 Aug 3-6; Taiyuan, CN. Cham (CH): Springer International Publishing; 2014. p. 225-32.

16. Kuang D, He L. Classification on ADHD with deep learning. Proceedings of the 2014 International Conference on Cloud Computing and Big Data (CCBD); 2014 Nov 12-14; Wuhan, CN. Washington, DC: IEEE Computer Society; 2014. p. 27-32.

17. Lam C, Yu C, Huang L, Rubin D. Retinal lesion detection with deep learning using image patches. Invest Ophthalmol Vis Sci 2018;59:590-6.

18. Liu M, Zhang J, Adeli E, Shen D. Deep multi-task multi-channel learning for joint classification and regression of brain status. In: Descoteaux M, Maier-Hein L, Franz A, Jannin P, Collins DL, Duchesne S, editors. Medical Image Computing and Computer-Assisted Intervention: MICCAI 2017; 2017 Sep 11-13; Quebec City, QC. Cham (CH): Springer International Publishing; 2017. p. 3-11.

19. Hagan MT, Demuth HB, Beale MH, De Jesus O. Neural network design. 2nd ed. Boston (MA): PWS Publishing; 2014.

20. Rosenblatt F. Report no. 85-460-1. The perceptron, a perceiving and recognizing automaton. Buffalo (NY): Cornell Aeronautical Laboratory; 1957.

21. Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. Nature 1986; 323:533-6.

22. Le QV, Ngiam J, Coates A, Lahiri A, Prochnow B, Ng AY. On optimization methods for deep learning. In: Getoor L, Scheffer T, editors. Proceedings of the 28th International Conference on Machine Learning; 2011 Jun 28-Jul 2; Bellevue, WA. Madison (WI): Omnipress; 2011. p. 265-72.

23. Schmidhuber J. Deep learning in neural networks: an overview. Neural Netw 2015;61:85-117.

24. Rumelhart DE, McClelland JL; University of California, San Diego, PDP Research Group. Parallel distributed processing. Cambridge (MA): MIT press; 1986. p. 567.

25. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature 2015; 521:436-44.

26. Bengio Y, Courville A, Vincent P. Representation learning: a review and new perspectives. IEEE Trans Pattern Anal Mach Intell 2013;35:1798-828.

27. Glorot X, Bengio Y. Understanding the difficulty of training deep feedforward neural networks. Proc Mach Learn Res 2010;9:249-56.

28. Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. In: Pereira F, Burges CJC, Bottou L, Weinberger KQ, editors. Advances in Neural Information Processing Systems 25 (NIPS 2012); 2012 Dec 3-6; Lake Tahoe, NV. Red Hook (NY): Curran Associates Inc.; 2013. p. 1097-105.

29. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. Proc IEEE Inst Electr Electron Eng 1998;86:2278-324.

30. Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, et al. Going deeper with convolutions. In: IEEE Computer Society, editor. 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2015 Jun 7-12; Boston, MA. Los Alamitos (CA): IEEE Computer Society; 2015. p. 1-9.

31. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. International Conference on Learning Representations 2015; 2015 May 7-9; San Diego, CA. Computational and Biological Learning Society; 2015.

32. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. Proc IEEE Comput Soc Conf Comput

Vis Pattern Recognit 2016;2016:770-8.

33. Hinton GE, Zemel RS. Autoencoders, minimum description length and helmholtz free energy. In: Cowan JD, Tesauro G, Alspector J, editors. Advances in Neural Information Processing Systems 6 (NIPS 1993); 1993 Nov 29-Dec 2; Denver, CO. San Francisco (CA): Morgan Kaufmann Publishers Inc.; 1994. p. 3-10.

34. Bourlard H, Kamp Y. Auto-association by multilayer perceptrons and singular value decomposition. Biol Cybern 1988;59:291-4.

35. Bengio Y, Lamblin P, Popovici D, Larochelle H. Greedy layer-wise training of deep networks. In: Scholkopf B, Platt JC, Hofmann T, editors. Advances in Neural Information Processing Systems 19 (NIPS 2006); 2006 Dec 4-7; Vancouver, BC. Cambridge (MA): MIT Press; 2006. p. 153-60.

36. Vincent P, Larochelle H, Lajoie I, Bengio Y, Manzagol PA. Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. J Mach Learn Res 2010;11:3371-408.

37. Graves A, Jaitly N. Towards end-to-end speech recognition with recurrent neural networks. In: International Machine Learning Society, editor. Proceedings of the 31st International Conference on International Conference on Machine Learning. Volume 32; 2014 Jun 21-26; Beijing, CN. Stroudsburg (PA): International Machine Learning Society; 2014. p. II-1764-72.

38. Graves A. Supervised sequence labelling. In: Graves A, editor. Supervised sequence labelling with recurrent neural networks. Berlin (DE): Springer Berlin Heidelberg; 2012. p. 5-13.

39. Dou Q, Chen H, Yu L, Zhao L, Qin J, Wang D, et al. Automatic detection of cerebral microbleeds from MR images via 3D convolutional neural networks. IEEE Trans Med Imaging 2016;35:1182-95.

40. Zreik M, Lessmann N, van Hamersvelt RW, Wolterink JM, Voskuil M, Viergever MA, et al. Deep learning analysis of the myocardium in coronary CT angiography for identification of patients with functionally significant coronary artery stenosis. Med Image Anal 2018;44:72-85.

41. Xu L, Tetteh G, Lipkova J, Zhao Y, Li H, Christ P, et al. Automated whole-body bone lesion detection for multiple myeloma on (68)Ga-pentixafor PET/CT imaging using deep learning methods. Contrast Media Mol Imaging 2018; 2018:2391925. https://doi.org/10.1155/2018/2391925.

42. Hirasawa T, Aoyama K, Tanimoto T, Ishihara S, Shichijo S, Ozawa T, et al. Application of artificial intelligence using a convolutional neural network for detecting gastric cancer in endoscopic images. Gastric Cancer 2018;21:653-60.

43. Albarqouni S, Baur C, Achilles F, Belagiannis V, Demirci S, Navab N. AggNet: deep learning from crowds for mitosis detection in breast cancer histology images. IEEE Trans Med Imaging 2016;35:1313-21.

44. Kawahara J, Brown CJ, Miller SP, Booth BG, Chau V, Grunau RE, et al. BrainNetCNN: convolutional neural networks for brain networks: towards predicting neurodevelopment. Neuroimage 2017;146:1038-49.

45. Yasaka K, Akai H, Kunimatsu A, Abe O, Kiryu S. Liver fibrosis: deep convolutional neural network for staging by using gadoxetic acid-enhanced hepatobiliary phase MR images. Radiology 2018;287:146-55.

46. Ypsilantis PP, Siddique M, Sohn HM, Davies A, Cook G, Goh V, et al. Predicting response to neoadjuvant chemotherapy with PET imaging using convolutional neural networks. PLoS One 2015;10:e0137036.

47. Prahs P, Radeck V, Mayer C, Cvetkov Y, Cvetkova N, Helbig H, et al. OCT-based deep learning algorithm for the evaluation of treatment indication with anti-vascular endothelial growth factor medications. Graefes Arch Clin Exp Ophthalmol 2018;256:91-8.

48. van der Burgh HK, Schmidt R, Westeneng HJ, de Reus MA, van den Berg LH, van den Heuvel MP. Deep learning predictions of survival based on MRI in amyotrophic lateral sclerosis. Neuroimage Clin 2016;13:361-9.

49. Oakden-Rayner L, Carneiro G, Bessen T, Nascimento JC, Bradley AP, Palmer LJ. Precision radiology: predicting longevity using feature engineering and deep learning methods in a radiomics framework. Sci Rep 2017;7:1648.

50. Gao XW, Hui R, Tian Z. Classification of CT brain images based on deep learning networks. Comput Methods Programs Biomed 2017;138:49-56.

51. Anthimopoulos M, Christodoulidis S, Ebner L, Christe A, Mougiakakou S. Lung pattern classification for interstitial lung diseases using a deep convolutional neural network. IEEE Trans Med Imaging 2016;35:1207-16.

52. Yasaka K, Akai H, Abe O, Kiryu S. Deep learning with convolutional neural network for differentiation of liver masses at dynamic contrast-enhanced CT: a preliminary study. Radiology 2018;286:887-96.

53. Mohamed AA, Berg WA, Peng H, Luo Y, Jankowitz RC, Wu S. A deep learning method for classifying mammographic breast density categories. Med Phys 2018;45:314-21.

54. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau

HM, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature 2017;542:115-8.

55. Niioka H, Asatani S, Yoshimura A, Ohigashi H, Tagawa S, Miyake J. Classification of C2C12 cells at differentiation by convolutional neural network of deep learning using phase contrast images. Hum Cell 2018;31:87-93.

56. Zhang W, Li R, Deng H, Wang L, Lin W, Ji S, et al. Deep convolutional neural networks for multi-modality iso-intense infant brain image segmentation. Neuroimage 2015;108:214-24.

57. Hu P, Wu F, Peng J, Liang P, Kong D. Automatic 3D liver segmentation based on deep learning and globally optimized surface evolution. Phys Med Biol 2016;61:8676-98.

58. Cha KH, Hadjiiski L, Samala RK, Chan HP, Caoili EM, Cohan RH. Urinary bladder segmentation in CT urography using deep-learning convolutional neural network and level sets. Med Phys 2016;43:1882.

59. Tian Z, Liu L, Zhang Z, Fei B. PSNet: prostate segmentation on MRI based on a convolutional neural network. J Med Imaging (Bellingham) 2018;5:021208. https://doi.org/10.1117/1.jmi.5.2.021208.

60. Tan LK, Liew YM, Lim E, McLaughlin RA. Cardiac left ventricle segmentation using convolutional neural network regression. In: IEEE Engineering in Medicine and Biology Society, editor. 2016 IEEE EMBS Conference on Biomedical Engineering and Sciences (IECBES); 2016 Dec 4-8; Kuala Lumpur, ML. Piscataway (NJ): IEEE; 2016. p. 490-3.

61. Sharma K, Rupprecht C, Caroli A, Aparicio MC, Remuzzi A, Baust M, et al. Automatic segmentation of kidneys using deep learning for total kidney volume quantification in autosomal dominant polycystic kidney disease. Sci Rep 2017;7:2049.

62. Fang L, Cunefare D, Wang C, Guymer RH, Li S, Farsiu S. Automatic segmentation of nine retinal layer boundaries in OCT images of non-exudative AMD patients using deep learning and graph search. Biomed Opt Express 2017;8:2732-44.

63. Xu J, Luo X, Wang G, Gilmore H, Madabhushi A. A deep convolutional neural network for segmenting and classifying epithelial and stromal regions in histopathological images. Neurocomputing 2016;191:214-23.

64. Miao S, Wang ZJ, Liao R. A CNN regression approach for real-time 2D/3D registration. IEEE Trans Med Imaging 2016;35:1352-63.

65. Lv J, Yang M, Zhang J, Wang X. Respiratory motion correction for free-breathing 3D abdominal MRI using CNN-based image registration: a feasibility study. Br J Radiol 2018;91:20170788. https://doi.org/10.1259/bjr.20170788.

66. Yang X, Kwitt R, Styner M, Niethammer M. Quicksilver: fast predictive image registration: a deep learning approach. Neuroimage 2017;158:378-96.

67. Yang X, De Andrade V, Scullin W, Dyer EL, Kasthuri N, De Carlo F, et al. Low-dose X-ray tomography through a deep convolutional neural network. Sci Rep 2018;8:2575.

68. Han X. MR-based synthetic CT generation using a deep convolutional neural network method. Med Phys 2017;44:1408-19.

69. Roy S, Butman JA, Pham DL. Synthesizing CT from ultrashort echo-time MR images via convolutional neural networks. In: Tsaftaris SA, Gooya A, Frangi AF, Prince JL, editors. Simulation and synthesis in medical imaging. Cham (CH): Springer International Publishing; 2017. p. 24-32.

70. Golkov V, Dosovitskiy A, Sperl JI, Menzel MI, Czisch M, Samann P, et al. Q-space deep learning: twelve-fold shorter and model-free diffusion MRI scans. IEEE Trans Med Imaging 2016;35:1344-51.

71. Bahrami K, Rekik I, Shi F, Shen D. Joint reconstruction and segmentation of 7T-like MR images from 3T MRI based on cascaded convolutional neural networks. In: Descoteaux M, Maier-Hein L, Franz A, Jannin P, Collins DL, Duchesne S, editors. Medical Image Computing and Computer-Assisted Intervention: MICCAI 2017; 2017 Sep 11-13; Quebec City, QC. Cham (CH): Springer International Publishing; 2017. p. 764-72.

72. Xu C, Xu L, Gao Z, Zhao S, Zhang H, Zhang Y, et al. Direct detection of pixel-level myocardial infarction areas via a deep-learning algorithm. In: Descoteaux M, Maier-Hein L, Franz A, Jannin P, Collins DL, Duchesne S, editors. Medical Image Computing and Computer-Assisted Intervention: MICCAI 2017; 2017 Sep 11-13; Quebec City, QC. Cham (CH): Springer International Publishing; 2017. p. 240-9.

73. Chen H, Dou Q, Ni D, Cheng JZ, Qin J, Li S, et al. Automatic fetal ultrasound standard plane detection using knowledge transferred recurrent neural networks. In: Navab N, Hornegger J, Wells WM, Frangi A, editors. Medical Image Computing and Computer-Assisted Intervention: MICCAI 2015; 2015 Oct 5-9; Munich, DE. Cham (CH): Springer International Publishing; 2015. p. 507-14.

74. Han L, Kamdar MR. MRI to MGMT: predicting methylation status in glioblastoma patients using convolutional recurrent neural networks. Pac Symp Biocomput 2018;23:331-

42.

75. Gao X, Lin S, Wong TY. Automatic feature learning to grade nuclear cataracts based on deep learning. IEEE Trans Biomed Eng 2015;62:2693-701.

76. Xue W, Brahm G, Pandey S, Leung S, Li S. Full left ventricle quantification via deep multitask relationships learning. Med Image Anal 2018;43:54-65.

77. Zhao X, Wu Y, Song G, Li Z, Zhang Y, Fan Y. A deep learning model integrating FCNNs and CRFs for brain tumor segmentation. Med Image Anal 2018;43:98-111.

78. Yang X, Yu L, Wu L, Wang Y, Ni D, Qin J, et al. Fine-grained recurrent neural networks for automatic prostate segmentation in ultrasound images. In: Singh S, Markovitch S, editors. Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17); 2017 Feb 4-9; San Francisco, CA. Palo Alto (CA): AAAI Press; 2017. p. 1633-9.

79. Qin C, Schlemper J, Caballero J, Price A, Hajnal JV, Rueckert D. Convolutional recurrent neural networks for dynamic MR image reconstruction [Internet]. Ithaca (NY): Cornell University; 2017 [cited 2018 Apr 23]. Available from: https://arxiv.org/abs/1712.01751.

80. Xu J, Xiang L, Liu Q, Gilmore H, Wu J, Tang J, et al. Stacked Sparse Autoencoder (SSAE) for nuclei detection on breast cancer histopathology images. IEEE Trans Med Imaging 2016;35:119-30.

81. He L, Li H, Holland SK, Yuan W, Altaye M, Parikh NA. Early prediction of cognitive deficits in very preterm infants using functional connectome data in an artificial neural network framework. NeuroImage Clin 2018;18:290-7.

82. Chen M, Shi X, Zhang Y, Wu D, Guizani M. Deep features learning for medical image analysis with convolutional autoencoder neural network. IEEE Trans Big Data. 2017 Jun 20 [Epub]. https://doi.org/10.1109/TBDATA.2017.2717439.

83. Cheng JZ, Ni D, Chou YH, Qin J, Tiu CM, Chang YC, 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;6:24454.

84. Zhao G, Wang X, Niu Y, Tan L, Zhang SX. Segmenting brain tissues from Chinese visible human dataset by deep-learned features with stacked autoencoder. Biomed Res Int 2016;2016:5284586. http://dx.doi.org/10.1155/2016/5284586.

85. Janowczyk A, Basavanhally A, Madabhushi A. Stain Normalization using Sparse AutoEncoders (StaNoSA): application to digital pathology. Comput Med Imaging Graph 2017;57:50-61.

86. Benou A, Veksler R, Friedman A, Riklin Raviv T. Ensemble of expert deep neural networks for spatio-temporal denoising of contrast-enhanced MRI sequences. Med Image Anal 2017;42:145-59.

87. Liu F, Jang H, Kijowski R, Bradshaw T, McMillan AB. Deep learning MR imaging-based attenuation correction for PET/MR imaging. Radiology 2018;286:676-84.

88. The Theano Development Team, Al-Rfou R, Alain G, Almahairi A, Angermuller C, Bahdanau D, et al. Theano: a Python framework for fast computation of mathematical expressions [Internet]. Ithaca (NY): Cornell University; 2016 [cited 2018 Apr 23]. Available from: https://arxiv.org/abs/1605.02688.

89. Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, et al. TensorFlow: a system for large-scale machine learning. In: Keeton K, Roscoe T, editors. OSDI'16 Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation; 2016 Nov 2-4; Savannah, GA. Berkeley (CA): USENIX Association; 2016. p. 265-83.

90. Collobert R, van der Maaten L, Joulin A. Torchnet: an open-source platform for (deep) learning research. In: Balcan MF, Weinberger KQ, editors. Proceedings of the 33rd International Conference on Machine Learning (ICML-2016); 2016 Jun 19-24; New York, NY. Stroudburg (PA): International Machine Learning Society; 2016. p. 19-24.

91. Wachinger C, Reuter M, Klein T. DeepNAT: deep convolutional neural network for segmenting neuroanatomy. Neuroimage 2018;170:434-45.

92. Cao Y, Steffey S, He J, Xiao D, Tao C, Chen P, et al. Medical image retrieval: a multimodal approach. Cancer Inform 2015;13(Suppl 3):125-36.

93. Qayyum A, Anwar SM, Awais M, Majid M. Medical image retrieval using deep convolutional neural network. Neurocomputing 2017;266:8-20.

94. Sun Q, Yang Y, Sun J, Yang Z, Zhang J. Using deep learning for content-based medical image retrieval. In: Cook TS, Zhang J, editors. Medical imaging 2017: imaging informatics for healthcare, research, and applications; 2017 Feb 11-16; Orlando, FL. Bellingham (WA): SPIE; 2017. p. 1013811-2. https://doi.org/10.1117/12.2251115.

95. Castelvecchi D. Can we open the black box of AI? Nature 2016;538:20-3.

96. Rodvold DM, McLeod DG, Brandt JM, Snow PB, Murphy GP. Introduction to artificial neural networks for physicians: taking the lid off the black box. Prostate 2001;46:39-44.

97. Li F, Sone S, Abe H, Macmahon H, Doi K. Malignant versus benign nodules at CT screening for lung cancer: compari-

son of thin-section CT findings. Radiology 2004;233:793-8.

98. You X, Sun X, Yang C, Fang Y. CT diagnosis and differentiation of benign and malignant varieties of solitary fibrous tumor of the pleura. Medicine (Baltimore) 2017;96:e9058.

99. Ribeiro MT, Singh S, Guestrin C. "Why should I trust you?": explaining the predictions of any classifier. In: Krishnapuram B, editor. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016 Aug 13-17; San Francisco, CA. New York (NY): Association for Computing Machinery; 2016. p. 1135-44.

100. Hohman F, Kahng M, Pienta R, Chau DH. Visual analytics in deep learning: an interrogative survey for the next frontiers [Internet]. Ithaca (NY): Cornell University; 2018 [cited 2018 Apr 23]. Available from: https://arxiv.org/abs/1801.06889.

101. Zhang Q, Zhu S. Visual interpretability for deep learning:
a survey. Front Inf Technol Electron Eng 2018;19:27-39.

102. Liu X, Wang X, Matwin S. Interpretable deep convolutional neural networks via meta-learning [Internet]. Ithaca (NY): Cornell University; 2018 [cited 2018 Apr 23]. Available from: https://arxiv.org/abs/1802.00560.

103. Domhan T, Springenberg JT, Hutter F. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In: Yang Q, Wooldridge MJ, editors. Proceedings of the 24th International Conference on Artificial Intelligence; 2015 Jul 25-31; Buenos Aires, AR. Palo Alto (CA): AAAI Press; 2015. p. 3460-8.

104. Shen C, Gonzalez Y, Chen L, Jiang SB, Jia X. Intelligent parameter tuning in optimization-based iterative CT reconstruction via deep reinforcement learning [Internet]. Ithaca (NY): Cornell University; 2017 [cited 2018 Apr 23]. Available from: https://arxiv.org/abs/1711.00414.