Machine Learning and Imaging Informatics in Oncology

Huan-Hsin Tseng Lise Wei Sunan Cui Yi Luo Randall K. Ten Haken
Issam El Naqa

Department of Radiation Oncology, University of Michigan, Ann Arbor, MI, USA

Keywords

Oncology · Machine learning · Imaging informatics

Abstract

In the era of personalized and precision medicine, informatics technologies utilizing machine learning (ML) and quantitative imaging are witnessing a rapidly increasing role in medicine in general and in oncology in particular. This expanding role ranges from computer-aided diagnosis to decision support of treatments with the potential to transform the current landscape of cancer management. In this review, we aim to provide an overview of ML methodologies and imaging informatics techniques and their recent application in modern oncology. We will review example applications of ML in oncology from the literature, identify current challenges and highlight future potentials.

© 2018 S. Karger AG, Basel

Introduction

Machine learning (ML) is an interdisciplinary field from artificial intelligence that draws upon advances in computer science, neuroscience, psychology and statistics for developing computer algorithms that can learn

tasks from data, without being explicitly programmed for this purpose [1]. The application of ML is currently prevalent in a wide range of diverse fields (e.g., banking, sports, politics and advertising) producing reliable guidance to decision making and reducing manual labor [2]. A subarea of ML called deep learning, allowing abstract representation of data via deep neural networks (also known as multilayer neural networks), has recently shown its potential in mimicking human cognition and challenging human intellectual abilities from video/board games to medicine. Recently, many high-profile companies have implemented ML techniques in their practice. For example, the Google Cloud [3] can convert audio to text and translate an arbitrary string into any supported language. Spotify music utilizes convolutional and recurrent neural networks for recognizing music genres and making recommendations for its users [4]. Rider sharing Apps like Uber and Lyft can predict rider demands using ML algorithms to minimize the waiting time of their customers [5].

In the field of medicine, information technology ventures are actively developing and seeking applications for their ML tools. For instance, Google DeepMind released mobile applications for diagnosis of eye disease, kidney injury and management of electronic patient records. In the field of oncology, there has been growing interest in

applying ML for diagnosis, prognosis and treatment queries. For example, the IBM Watson for Oncology (WFO) system has demonstrated its effectiveness in making treatment recommendations for specific cancer patients. In breast cancer, WFO was able to learn an extensive corpus of medical journal, textbook and treatment guidelines information at Memorial Sloan Kettering Cancer Center (MSKCC) by natural language processing to identify articles that are well matched to the characteristics of specific patients. It also incorporated data from over 550 breast cancer cases in MSKCC, including variables like patient characteristics, comorbidities, functional status, tumor characteristics, stage, imaging and other laboratory findings [6]. WFO can further refine its analytical process according to feedback given by experts. The system can finally provide treatment planning recommendations (surgery, chemotherapy, immunotherapy, radiotherapy) and alternative options within each treatment plan (e.g., drugs or doses) to a specific patient. It has been tested at the Manipal Comprehensive Cancer Center, showing a high concordance (93%) with a multidisciplinary tumor board [7].

Algorithms based on deep learning, such as a convolutional neural network (CNN), have been applied in imaging diagnosis of a wide variety of cancers showing high accuracy comparable or superior to human experts. For instance, a CNN was able to distinguish between the most common and deadliest types of skin cancers by learning from a data set consisting of 12,940 clinical images, outperforming two dermatologists on a subset of the same validation set [8]. In a challenge competition to develop automated solutions for detecting lymph node metastases in breast cancers from pathology images, the top-performing algorithm was also based on a deep learning CNN [9], which achieved a better diagnostic performance than a panel of 11 pathologists under a simulated exercise designed to mimic routine clinical workflow. Besides the high prediction accuracy offered by ML algorithms, they also enjoy high efficiency and can be cost effective. For instance, a well-trained CNN from previous clinical examples can achieve accurate diagnosis in a fraction of a second at any time offering the possibility of a universal access to diagnostic care anywhere, anytime.

ML can also be an effective tool for molecular targeting by unveiling the complex relationships of underlying genetics and other biological information. A blood test called CancerSEEK [10] was reported to be able to detect 8 common cancer types from very early stages of disease and localize the origin of cancer to a small number of anatomical sites by assessing the levels of circulating pro-

Fig. 1. A schematic of the relation between artificial intelligence, machine learning, deep learning, big data and data science. It is noted that machine learning is a computational branch from artificial intelligence that aims to provide computers with the ability to perform tasks beyond their original programming such as data mining and big data analytics.

teins and mutations in the DNA. This test can be applied in early detection of cancers, and it can conceivably reduce deaths.

In the future, cognitive learning systems such as WFO may potentially offer physicians the necessary tools to tailor their treatments to an individual patient based on the synthesized knowledge by ML algorithms from the existing literature and/or interactive learning from clinical oncology experts. Due to the sophisticated technical advances and the tremendous growth of genetic and clinical data, it is becoming harder for busy physicians to stay current about every emerging new finding. On the other hand, ML-based systems can acquire such knowledge from a large volume of unstructured and structured data sets, aggregate and effectively present such synthesized knowledge to practicing physicians as a second opinion to aid and support their decision making and improve cancer patients' management.

In this article, we will review some of the basic concepts and methods commonly applied in ML. Then, we will present several examples of ML and imaging informatics applications in diagnosis, prognosis and treatment of cancers. Finally, we will discuss current technical and administrative barriers for a more comprehensive and wider incorporation of ML techniques into clinical practice and offer some tentative recommendations to realize the tremendous potentials of ML for oncology and cancer patients' care.

Fig. 2. a An illustration of supervised learning using neural networks **(b)** classifying synthetic data of binary labels (blue and red scatter dots), where the nonlinear decision boundary is shown in white. **b** A multilayer (deep) neural network with 2 hidden layers. The so-called deep learning usually refers to learning algorithms heavily relying on such computational units.

What Is ML?

ML is broadly referred to as computer algorithms that can provide computers with the ability to learn patterns from data or make predictions based on prior examples. ML, a term first coined by Arthur Samuel, is considered as a major branch of artificial intelligence (Fig. 1), as proposed by John McCarthy [11], and was defined as “involves machines that can perform tasks that are characteristic of human intelligence.” ML is generally designed to learn analytical patterns from data and making gener-

alizations (predictions) based on its exposition to previous samples [12]. Thus, it is a field strongly tied to cognitive psychology, neuroscience, computational and statistical principles that also aim at data mining and performance predictions. To explain the concept of ML more concretely, the following provides more details.

Definition

A technical definition of ML is, quoted from Michalski et al. [13], “a computer program is said to learn from experience E with respect to some class of tasks T and

performance measure P if its performance at tasks in T , as measured by P , improves with experience E .” Tom Mitchell further illustrated this definition by showing an example of playing checkers, where E = playing checkers, P = ability to win and T = game rules. In other words, ML provides computers with the ability to perform tasks beyond what they were originally programmed for, i.e., they learn how to perform tasks in a more or less similar fashion to an autonomous human operator.

Categories

There are mainly 3 categories recognized in ML: supervised learning, unsupervised learning and reinforcement learning. Supervised learning requires a data set containing input and output labels, which are the desired outputs or outcomes, so that a computer is trained by a labeled data set as if it were learning under the supervision of a teacher. Technically, supervised learning aims to find a mathematical function that can map input data pairs into output labels. On the other hand, unsupervised learning can operate on a data set without given labels. In such a case a computer algorithm is tasked to figure out (intrinsic) structures within the data (e.g., Fig. 3), where these intrinsic structures could mean clusters or support regions of data.

Some typical supervised learning algorithms include (logistic, LASSO, Ridge, etc.) regressions, support vector machines (SVMs), random forests, neural networks (NNs), etc. Examples of unsupervised learning are principal component analysis (PCA) [14], Laplacian eigenmaps [15], t-SNE [16], p-SNE [17], autoencoders [18], etc. Illustrations of a supervised learning and an unsupervised learning are given in Figure 2a and Figure 3a, respectively. A clinical application by Dawson et al. [19] utilized an unsupervised PCA to indicate the presence of linear separability in xerostomia (dry mouth) data of patients at high or low risks after radiotherapy exposure of the parotid gland to radiation (Fig. 3b). Intuitively, supervised learning can usually perform more effectively in classifying data due to the additional guidance of known answers (labels) provided to it. Thus, in this respect, unsupervised learning is generally considered a harder computational problem, where cognitive learning is assumed to be implicit.

From a probabilistic perspective, supervised learning algorithms can also be categorized into *discriminant* or *generative* models. It is usually assumed that inputs (x) and their labels (y) in supervised classification arise from a joint probability $p(x, y)$. A discriminant classifier model, defined by the posterior probability $p(x | y)$ can be used

Fig. 3. a An illustration of an unsupervised learning using p-SNE with open image database Olivetti faces, where similar images (same person) are clustered automatically without providing any identity information. **b** Dawson et al. [19] (reprint permission granted) demonstrated that principal component analysis can be used to observe clinical data structure. In this case, the data describing the xerostomia occurrences due to parotid gland dose distributions are linearly separable.

(Figure continued on next page.)

to map inputs (x) to class labels (y) without necessarily knowing the underlying joint probability function. Whereas a generative classifier attempts to learn the architecture of such joint probability $p(x, y)$ first and then make their predictions by using Bayes' rule to calculate conditional probabilities $p(y | x)$ and choosing the most likely label y [20]. The advantage of the generative approach is that we can use the algorithm to generate new synthetic data similar to the existing ones, while the discriminant algorithm generally offers better performance for classification tasks [16, 21].

If a given data set contains temporal information, one may utilize dynamic ML algorithms that can take time information into account such as a recurrent neural network (RNN) or long short-term memory network [22]. In addition, classical Bayesian network techniques are able to

perform both *dynamic* and *static* predictions, and recursive Bayesian methods can estimate an unknown probability density function recursively over time based on incoming information. If the variables are linear and normally distributed, the recursive Bayesian method becomes equivalent to the well-known Kalman filter widely used in control and signal processing applications [23, 24].

As for the third category, reinforcement learning, it is designed to embody a software agent (which may represent a clinician in our case) to take actions when interacting with a given environment (e.g., clinical treatment). Usually there is a definite goal for the agent to reach, via a so-called *reward function* (e.g., better treatment out-

come). Winning a chess/Go game by an agent can be a goal of a reinforcement learning algorithm in a board game, such as in the example AlphaGo of DeepMind [25]. It is worth noting that reinforcement learning is a modern extension of classical statistical decision-making schemes known as Markov decision processes, which originally appeared in the 1950s [26] and are currently empowered by advanced computing technologies.

Deep Learning

Recently, an interesting and powerful branch in ML called deep learning is demonstrating tremendous success in solving pattern recognition and computer vision

Fig. 4. The structure of an convolutional neural network, usually consisting of 3 distinct layers: the convolution layer, the pooling layer and a final fully connected layer (Fig. 2b), where the convolution layer and pooling (sub-sampling) layers may be connected several times before a final fully connected layer is encountered. An image mapped by a convolution layer is called a feature map, which triggers attention of many computer scientists. Figure created by Aphex34 distributed under a CC BY-SA 4.0 license (from Wikimedia Commons).

problems compared to classical ML techniques. These learning algorithms are generally based on NN architectures with more than 2 hidden layers (Fig. 2b), and hence the qualifier “deep.” Deep learning has empirically proven to be capable of efficient learning from complex tasks [27], which is essentially due to an inherent characteristic called the universal approximation property [28]. In CNNs, such a property can be interpreted as learning data representation in a hierarchal manner while optimizing prediction for the task at hand without the risk of overfitting [29] and avoiding the feature selection problem inherent in classical ML problems [27]. Some specialized (deep) NNs of a desired purpose are developed subsequently to perform different tasks such as CNNs (Fig. 4), for image recognition/classification and RNNs for sequential learning such as text captions of images.

CNNs are said to be inspired by the work of Hubel and Wiesel on the animal visual cortex [30]. In practice, a CNN is one particular type of NN (Fig. 4) usually consisting of 3 parts: (1) a convolutional layer, (2) a pooling (down-sampling) layer and (3) one final fully connected layer for classification purposes. The distinguished convolutional part is generally the most important piece (hence the name CNN) and it aims at learning and extracting features from the input data (e.g., edges, textures, etc., from images) [29]. The pooling layer serves as a data reduction mechanism, and the last fully connected layer makes a final judgment as to which class the input images may belong to.

A CNN can go deep by consecutively repeating the convolution and the pooling layers. With the ever-in-

creasing computing power to deal with the growing size of data, deep learning can be efficiently trained and applied. As mentioned in LeCun et al. [27], deep learning outperforms conventional shallow NNs and other ML algorithms by capturing complex structures in high-dimensional data. Each feature map via a convolutional kernel is considered a representation of a higher or more abstract level [29]. Thus, representations of deeply transformed layers emphasize quantities that are crucial for computers to discriminate, such that they can be considered as latent (hidden) features in the data that are automatically learned by a CNN.

Certain CNN architectures have proven to be effective: LeNet is the first application of a CNN to read zip codes and digits by Yann LeCun et al. in the 1990s [31]; AlexNet, the champion of the ImageNet pattern recognition challenge 2012, demonstrated surprising performance compared to other common methods and is considered a watershed moment in modern ML applications [32]. GoogLeNet, the laureate of ILSVRC 2014, introduced an inception module into the architecture to reduce the huge number of parameters by up to 12 times over that of AlexNet and found out an optimal local sparse structure [33]. Also, VGGNet, introduced in 2014, used filters only stacked to increase the learning depth [34]. There are two other variants widely used in the literature, VGG16 and VGG19, of the original VGGNet, where the “16” and “19” refer to the layer depth, respectively. Notably, in general the depth of a layer could improve the classification accuracy, but this is not always the case. The learning accuracy empirically ceases to grow after a certain depth has

been reached. This is due to the vastly increasing amount of weights when adding more layers. In the case where using weights free of concern is a luxury, alternatives should be considered. Particularly, in the medical field where data sets are typically of small size, U-Net [35] may be a good choice for such scenarios.

The success of CNNs and RNNs owes to their design for taking related information into account, where a CNN considers multidimensional neighboring data (e.g., image pixels/MRI voxels) and an RNN focuses on (1-dimensional) sequential (temporal) relations such as human voice recognition or texts. These features make them naturally suitable for decision support by incorporating both the spatial and the temporal information. Hence, making them strong candidates for aiding treatment planning and dose adaptation. Another possible way to achieve automatic decision support in radiotherapy, for instance, is via reinforcement learning, where desired benefits are maximally pursued by a software agent. It is the same principle that drove AlphaGo into winning the Chinese Go game. Thus, deep learning has the potential to optimize prediction of outcomes and identify optimal strategies for precision treatment in oncology. Later in the section “Handcrafted Feature Extraction,” we shall discuss how these technologies are currently applied to help extract useful cues from imaging data, as a valuable resource for precision oncology.

ML for Imaging Informatics in Oncology

Medical imaging has been widely applied clinically over the past decades. Recently, it has shown even more potential and utility due to the vast development of quantitative imaging techniques and recent breakthroughs in the ML community [29].

In general, there are 2 main types of imaging acquisitions: (1) anatomical imaging, including conventional X-ray, ultrasound, computed tomography (CT) and magnetic resonance imaging (MRI), etc.; (2) functional (molecular) imaging, including positron emission tomography (PET), single-photon emission computed tomography (SPECT) and diffusion-weighted MRI, etc. In order to combine the advantage of anatomical resolution and functional information of tissues, multimodality imaging techniques such as SPECT/CT, PET/CT or PET/MRI were also developed. With all these techniques available, images can provide valuable data encoded with patient individual information about tumor, tumor environment and genotype that can be data mined to help with the di-

agnosis, prognosis and prediction of oncology outcomes [36].

This raises the question: “How can we discover underlying biological relationships in these huge amounts of imaging data?” Beyond the already complex procedures clinicians use to read and interpret medical images for making routine decisions, they are also extracting other information, however in a relatively qualitative way (for example, the boundary of the tumor or its heterogeneity, etc.). Fortunately, with the development of advanced pattern recognition techniques and statistical learning tools, digital medical images can now be converted into mineable high-dimensional data via high-throughput extraction of quantitative features. This has shown great potential for precision medicine in oncology. The conversion of medical images into a large number of advanced features and the subsequent analysis that relates these features with biological end points and clinical outcomes give rise to the field of *radiomics*.

Radiomics

Medical imaging plays an important role in oncological practice from diagnosis, staging of tumors, treatment guidance, evaluation of treatment outcomes and follow-up of patients. Being a noninvasive method, images are able to provide both spatial and temporal information of the tumor. Extraction of quantitative features from medical images together with subsequently relating these features to biological end points and clinical outcomes is referred to as the field of radiomics. “Radio” comes from radiology, which refers to radiology images, e.g. CT, MRI and PET, while “-omics” stands for the technologies that aim at providing collective and quantitative features for the entire system and explore the underlying mechanisms. It is widely used in biology, such as in the study of genes (genomics), proteins (proteomics) and metabolites (metabolomics) [37]. The origin of radiomics is medical image analysis and understanding. The goal of radiomics is to take advantage of the digital data stored in those images to develop diagnostic, predictive or prognostic radiomic models to help understand the underlying biological/clinical processes, support personalized clinical decisions and optimize individualized treatment planning. The core of radiomics is the extraction of quantitative features, with which we can apply all the advanced ML algorithms and build models to bridge between images and biological and clinical end points. A central hypothesis of radiomic analysis is that the imaging features are able to capture distinct phenotypic differences, like genetics and proteomics patterns or other clinical outcomes, so that we can infer these

end points. This hypothesis has recently been proven by many researches. Segal et al. [38] showed that the dynamic imaging traits (features) in CT systematically correlate with the global gene expression programs of primary human cancer. Aerts et al. [39] found that a large number of radiomic features extracted from CT images have prognostic power in independent data sets of lung and head/neck cancer patients. It is interesting to note that the study identified a general prognostic phenotype for both lung and head/neck cancers. Vallières et al. [40] extracted features from ^{18}F -labeled fluorodeoxyglucose (FDG) PET and CT images and performed risk assessment for locoregional recurrences, distant metastases and overall survival in head and neck cancer. These studies assured the potential of radiomic features for analyzing the properties of specific tumors. A central component of all the examples stated above is the ability to obtain or infer the hidden information from the pixels (voxels) in the digital images. Generally, there are 2 main processes to help relate raw images to the end points: (1) feature extraction; (2) classification or regression using the extracted features. There are 2 general ways to extract useful features: (1a) handcrafted techniques or directly using existing radiomic signatures; (1b) automatic learning of image representation by deep NNs (e.g., CNNs). For oncology classification problems, methods and tools for feature extraction via both conventional ML algorithms (e.g., SVMs, random forests) and newly emerging deep learning algorithms are growing at a rapid pace.

Handcrafted Feature Extraction (Conventional Radiomics)

The regions of interest (ROIs) in cancer diagnosis/prognosis are generally spatially and temporally heterogeneous. It has been shown that radiomics features for capturing these intratumor heterogeneities were powerful in prognostic modeling of aggressive tumors. These changes in 4D (3D space + 1D time) space play an important role in analyzing and monitoring the disease status [39]. Thus, it is natural to divide radiomic features further into 2 types: spatial (static) and temporal (dynamic) [41]. Static features are based on intensity, shape, size (volume), texture and wavelet, while dynamic features are based on kinetic analysis using time-varying acquisition protocols, such as dynamic PET or MRI. Both of these features offer information on the tumor phenotype and its microenvironment (habitat). Examples of static features are: (a) morphological (shape descriptors), (b) first-order and (c) second-order (texture) features. Texture features, specifically, can provide statistical interrelationships between

voxels and capture special patterns in the ROIs to compensate for the loss of information from the first-order features due to the spatial information associated with the relative positions of the intensity levels of the voxels.

These features can then provide a quantitative representation that to some degree can mimic the features that clinicians may pay attention to, while also offer the potential of obtaining more information invisible to the human eye. After producing isotropic interpolated voxel sizes and discretize gray-level images, features can be calculated from varying common texture matrices such as: the gray-level co-occurrence matrix [42], gray-level run length matrix [43], gray-level size zone matrix [44], gray-level distance zone matrix [45], neighborhood gray tone difference matrix [46] and the neighboring gray-level dependence matrix [47]. These gray-level matrices provide statistical methods to capture the spatial dependence of gray-level intensities that constitute the textures of an image. For a more detailed introduction to radiomics analysis, one may refer to Zwanenburg et al. [48].

For time-varying acquisition protocols, such as dynamic PET and MRI, radiomic features are extracted based on kinetic analysis of the dynamic images. Compartment models are widely used for tracer transport, its binding rates and metabolism modeling. For example, FDG-PET imaging has shown great success in tumor detection and cancer staging using FDG as the tracer of choice to visualize intratumoral glucose metabolism. Beyond these common radiomic features (mostly statistical texture features), one can also apply other advanced pattern recognition features like fractal features, which are based on the concept of fractional Brownian motion and which represent the normalized average absolute intensity difference of pixel pairs on a surface at different scales, scale-invariant features, which are invariant to image spatial scaling and rotations, and are able to provide robust matching across a large range of affine transformation [49], or histograms of oriented gradient features, which are obtained by counting occurrence of gradient orientation in localized parts of an image [50]. In addition, one may also develop their own new ad hoc features based on the understanding of the specific task at hand.

Machine-Engineered Feature Extraction (Deep Radiomics)

An alternative approach for feature extraction is driven purely by the data themselves using ML techniques such as CNNs and is usually referred to as “feature engineering.” Unlike obtaining handcrafted features as mentioned above, CNNs are able to engineer important fea-

tures considered critical for computers to learn characteristics from various image data automatically, although these latent features are not always human-recognizable. But, owing to the powerful performance of CNNs, computer scientists have managed to unveil the feature maps a computer can recognize, making such features more interpretable [29].

The idea of CNNs has been applied to medical image processing as early as 1993, El Naqa [36] used a shift-invariant NN to detect clustered microcalcifications in digital mammograms. Gillies et al. [37] investigated the classification of ROIs on mammograms using a CNN with spatial domains and texture images. Recently, a large number of inspiring works applying deep CNNs to medical image analysis have been presented. For example, Segal et al. [38] used 3 CNN architectures, namely CifarNet, AlexNet and GoogLeNet with transfer learning for computer-aided detection problems. They reported that the applications of CNN image features can be improved by either exploring the handcrafted features or by transfer learning (i.e., using information from other domains such as natural images to inform the medical application at hand) [38]. Although CNN methods require little engineering by hand and learn their features automatically, they are also limited by the available data size. In the medical field, it is relatively difficult to collect such large amounts of data comparable to that in other fields such as computer vision or board games. Another limitation is the lack of labeling for the data (clinical outcomes). Even if the former two issues are resolved, the features obtained from CNNs may be hard to interpret in the clinical sense, which may not be reassuring for medical practitioners and patient care. Transfer learning, data augmentation, generative adversarial nets [39], semisupervised learning among others have been proposed to address the data limitation problem by providing additional data for training, while deep learning interpretation is still in its infancy.

Feature Selection, Model Construction and Validation

In imaging tasks for radiation oncology such as segmentation or tumor contouring, supervised learning is typically useful. As mentioned earlier, in supervised learning one finds a “good fit” for the labeled data among several ML and statistical models, where such a “good fit” is usually determined by training error and internal/external validation performance. It can be challenging to find an accurate and stable model fitting data with a large number of features extracted from medical images, especially when the sample size (patients’ number) is much less than the data features. In these situations, overfitting

and high variance are the major concerns due to the so-called *curse of higher dimensionality*. One viable method is to trim the data with feature selection, namely selecting a subset of variables that are indicative for one’s classification purpose. Such feature selection of data may be useful when features are redundant, highly correlated or sometimes irrelevant with respect to the classification task.

By performing feature selection, the variable space is reduced to mitigate the tendency of overfitting such that it helps building a more robust predicting model. Furthermore, computation cost and storage of data can be reduced. It may also be beneficial for interpreting and understanding the underlying data mechanism if selected structures properly reflect the classification purpose. Common feature selection methods are filtering methods and wrapper methods [51], where the former selects variables by ranking them with correlation coefficients and the latter searches for an optimal set of features by evaluating “usefulness” towards a given predictor using different combinations of features within a learning scheme.

After feature selection of data, one proceeds to build a classification model by evaluating the performance on training data and generalization error on an independent test (validation) set. A complex model can fit the training data well (low bias), while its generalization performance to out-of-sample may be poor (high variance).

Such a bias-variance tradeoff is due to the complexity of a model. A fundamental quantity called *Vapnik-Chervonenkis dimension* (VC dimension) characterizes such model complexity for a class of models, e.g., NNs, linear classifiers and SVMs. It is known that linear classifiers on a 2D plane have VC dimension = 3, while NNs have VC dimension = $O(n \log n)$, where n is the total number of parameters (weights) in the network [52]. One then easily sees that NNs are equipped with stronger capacity to fit (training) data, which meets our empirical understanding. Therefore, the purpose of validation and model selection aims to choose a proper class of classifiers with suitable VC dimension to characterize the data. In fact, Vapnik has proved a useful formula describing the relation between training/testing error and VC dimension [53]:

$$\Pr \left(\text{test error} \leq \text{training error} + \sqrt{\frac{D \left(\log \left(\frac{2N}{D} \right) + 1 \right) - \log \left(\frac{\eta}{4} \right)}{N}} \right) = 1 - \eta,$$

with $0 \leq \eta \leq 1$ and $D \ll N$ when D is the VC dimension and N is the size of the training set. This equation exactly describes that when the probability of test error is likely

to be larger than training error and by how much (the square root term) determined by D . In particular, when the VC dimension is large, the test error is most likely to be larger than the training error, and hence meets our intuition of overfitting. With the knowledge where overfitting/underfitting comes from, one is then dedicated to select a proper model describing data.

To perform model assessment and overcome model selection, 2 major methods of validation may be utilized: K -fold cross-validation (CV) and the bootstrap method, where K -fold (K : a positive integer, e.g., $K = 5, 8, 10, \dots$) CV randomly splits data K times into approximately K equal-sized and mutually exclusive parts with one part reserved as validation data and the other $K - 1$ ones serving as the training set. Choice of K is arbitrary; however, $K = 5, 10, N$ are commonly used with N as the total sample size. In general, large K leads to low bias and overestimates the true prediction error since the training sets will be approaching the entire data set but tend to give higher variance [54]. An extreme case where $K = N$, called leave-one-out CV, is almost unbiased but will have higher variance [55], and thus it is suitable for smaller data sets. Other choices like $K = 5, 10$ will provide good compromise between bias and variance tradeoff [56]. Some studies further reported that 10-fold CV yielded the best results with experiments on real-world data of certain types [57].

Another validation method called bootstrapping has a basic idea of randomly drawing samples with replacement from the training data, called *bootstrap samples*, with each bootstrap sample having the same size as the original training set. To apply the bootstrap idea for assessing models, several statistical estimators have been proposed: in particular, *leave-one-out bootstrap*, “0.632 estimator” [55] and “0.632+ estimator” [58] are commonly used. Leave-one-out bootstrap mimicking leave-one-out CV keeps tracks of predictions from bootstrap samples not containing certain observation i , such that its error estimator can overcome overfitting problems when compared to the pure bootstrap method, where the estimator is given as the following:

$$Err = \frac{1}{N} \sum_{i=1}^N \frac{1}{|C^{-i}|} \sum_{b \in C^{-i}} L(y_i, f^{*b}(x_i)),$$

where the “0.632 estimator” [55] further alleviates bias towards estimates of prediction error. Another technique, the “0.632+ estimator” [58], further improves the 0.632 estimator by considering the amount of overfitting. It is known that bootstrap can be shown to fail in certain exquisitely designed statistical examples as well as in a

case when a memorizer module is added [57]. However, such artificial counterexamples require more mathematical labor that is beyond the scope of this paper; therefore interested readers are encouraged to view the construction in Hastie et al. [54]. On many occasions, both CV and bootstrap methods are shown to be valuable and have compatible results. As in Hastie et al. [54], a comparison of CV and bootstrap is demonstrated for particular problems and fitting models. They found that either CV or bootstrap yields a model fairly close to the best available. Therefore, one usually needs to determine which validation method actually provides the best description based on the field test with one’s data at hand.

As a final remark, model selection and feature selection should not exhaust all samples, since the features or model selected by exhausting all samples may derive more optimistic performance estimation. To avoid this issue, in small-sample-size problems, nested CV techniques can be utilized to make full use of the data as well as to give an unbiased prediction estimation, where typically an outer loop is created for performance estimation, and an inner loop (in contrast to the outer loop) is established for searching optimal hyperparameters, model training, or feature selection. In any case, the gold standard for validating a model is performing external validation on independent data sets. This is highlighted in the Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis (TRIPOD) recommendations [59].

Examples of ML Application in Oncology

In this section, we illustrate how ML can extract relevant imaging information for oncology applications. Three examples shall be investigated, where the first two compare the radiomics using classical analysis and the modern deep CNN method; the last one directly utilizes deep learning for identifying metastatic breast cancer motivated by a recent grand challenge [9, 60].

Radiomics signatures have been built with the power to detect tumor heterogeneity, predict outcomes (e.g., recurrence, distant metastases, response to certain treatment) and conduct survival analysis of various cancers. We will review several results obtained with radiomics, using conventional ML methods (such as random forests, SVMs), and compare them with deep learning methods, e.g. CNNs, and present one example that fuses both conventional and deep models (i.e., deep radiomics).

Example 1: Conventional Radiomics for Predicting Failure Risks in Head and Neck Cancers

Vallires et al. [40] built a radiomics model based on pretreatment FDG-PET and CT images to analyze the risk assessment of locoregional recurrences, distant metastases and overall survival in a head and neck cancer. 1,615 radiomic features from 300 patients from 4 institutions were extracted and divided into training and testing data sets. Considering the large number of features, feature set reduction was first applied by ranking the features based on an information gain criterion, where the correlation of the feature with end points and features of their intercorrelations were considered. The goal was to maximize the relevance of features with outcomes and minimize the redundancy. Subsequently, a forward feature selection method was applied to determine the final model. Feature selection, prediction performance estimation, choice of model complexity and final model computation processes were carried out using logistic regression (classification), Cox regression (survival analysis) and bootstrap resampling. Prediction models consisting of radiomic information only were first constructed for each of the three head and neck cancer outcomes. Meanwhile, clinical factors (e.g., tumor volume, age, T stage) were analyzed by stratified random subsampling in the training set and resulted in 3 groups of clinical variables for each outcome. Final prediction models were constructed by combining the selected radiomic and clinical features via random forests. The performance of prediction models was estimated using receiver operating characteristic metrics, the concordance index (CI) and the *p* value obtained from Kaplan-Meier analysis using the log-rank test between 2 risk groups (locoregional recurrence: AUC = 0.69, CI = 0.67; distant metastases: AUC = 0.86, CI = 0.88). The analysis showed that radiomics can provide important prognostic information for the risk assessment of the 3 outcomes in head and neck cancer in this study (Fig. 5).

Example 2: Deep Radiomics for Breast Cancer Diagnosis

Antropova et al. [61] devised a method that extracted low- to mid-level features using a pretrained CNN and combined them with handcrafted radiomic features for breast cancer diagnosis. Full-field digital mammography, breast ultrasound and dynamic contrast-enhanced MRI images were used in this study. Three data sets were used separately to test the methodology. CNN features were extracted with VGG19 architecture, pretrained on ImageNet. There are 5 stacks in the architecture, with each stack containing 2 or 4 convolutional layers and a maximum pooling layer, followed by 3 fully connected layers. Images were duplicated to be input to the 3-color channels for full-field digital mammography and ultrasound data extracted at precontrast and the first two postcontrast time points for the dynamic contrast-enhanced MRI data set. CNN features were then obtained from each of the 5 maximum pool layers. Five-feature vectors were obtained by average pooling along spatial dimensions. This method avoided the preprocessing step for images with varying sizes by preserving the spatial structure. Radiomic features describing lesion properties such as size, shape, texture and morphology were also extracted from ROIs. A nonlinear SVM with a radial basis function kernel was used to build models for both CNN and radiomic features. The two classifiers were fused by averaging the outputs to give the final model. From receiver operating characteristic analysis, they claimed that the fusion-based method, on all imaging modalities, performed significantly better than conventional radiomic models in the task of distinguishing benign and malignant lesions, with an AUC = 0.89 for dynamic contrast-enhanced MRI, with AUC = 0.86 for full-field digital mammography and AUC = 0.90 for ultrasound. In summary, their analysis showed the feasibility of combining deep learning and conventional feature extraction methods for breast cancer diagnosis (Fig. 6).

Fig. 5. a The workflow of the model built by Vallires et al. [40]. The best combinations of radiomic features were selected in the training set, where these radiomic features were then combined with selected clinical variables in the training set. Independent prediction analysis was later performed in the testing set for all classifiers fully constructed in the training set. ROC, receiver operating characteristic. **b** Risk assessment of tumor outcomes in Vallires et al. [40]. **b1** Probability of occurrence of events for each patient of

the testing set. The output probability of occurrence of events of random forests allows for risk stratification. **b2** Kaplan-Meier curves of the testing set using a risk stratification into 2 groups as defined by a random forest output probability threshold of 0.5. All curves show significant prognostic performance. **b3** Kaplan-Meier curves of the testing set using a risk stratification into 3 groups as defined by random forest output probability thresholds of 1/3 and 2/3.

(For figure see next pages.)

b

1

2

3

5

Fig. 6. Lesion classification pipeline based on diagnostic images. CNN, convolutional neural network; ROI, region of interest; SVM, support vector machine; ROC, receiver operating characteristics; FFDM, full-field digital mammography. Two types of features are extracted from a medical image: CNN features with pretrained CNN (**a**) and handcrafted features with conventional CADx (**b**). High- and low-level features extracted by pretrained CNN are evaluated in terms of their classification performance and preprocessing requirements. Furthermore, the classifier outputs from the pooled CNN features and the handcrafted features are fused in the evaluation of a combination of the two types of features.

(Figure continued on next page.)

Example 3: A Deep CNN for Detection of Lymph Node Metastases

An automated detection of cancer challenge, Camelyon16, was set up [9] to develop algorithms for lymph node detection of metastases in women with breast cancer, where a training data set of 399 whole slide images collected from two institutions, Radboud University Medical Center and University Medical Center Utrecht, was provided. The competition was intended to motivate ML algorithms in the application of medical cancer imaging, where two independent tasks were asked to be evaluated. In task 1, the participants were asked to demonstrate the ability of localizing tumor; in task 2, the participants were asked to discriminate images with or without sentinel axillary lymph nodes, i.e., an image classification problem. During the competition, a panel of 11 pathologists participated and independently reviewed the same data set. Of all the 23 teams participating, the best results (algorithms) were derived by a joint team led by Harvard Medical School and Massachusetts Institute of Technology, where Wang et al. [60] designed a composite image classifier using a deep CNN, GoogLeNet of 27 layers and in total more than 6 million parameters. Essentially, their

NNs were trained by millions of patches (out of the whole slide image, see Fig. 7a) for patch level predictions to discriminate tumor patches from normal patches. These patch-by-patch predictions were then gathered to form a complete tumor probability heat map for one whole slide, as in Figure 7b. Furthermore, to decrease the computational time, several techniques were applied, including transforming images from RGB (red, blue, green) color into HSV (hue, saturation, value) color and identifying tissues via meaningless white background removed.

In fact, before Wang et al. [60] presented their final winning model, several existing advanced architectures were tested for selecting strong candidates, such as GoogLeNet [33], AlexNet [62], VGG16 [63] and a face-orientated deep network [64] giving the following testing performance:

Patch-wise classification accuracy:

GoogLeNet	98.4%
AlexNet	92.1%
VGG16	97.9%
FaceNet	96.8%

Naturally, due to this experiment, GoogLeNet appeared to be the strongest candidate for their task, and

thus was utilized as the winning model. By their delicate design of data preprocessing, model selection between several viable CNN architectures and final data post-processing, their classifiers were able to achieve an AUC = 0.925 for task 2 of the whole slide image classification and a score of 0.7051 for the tumor localization of task 1, where a human pathologist independently cross-validated the image data, obtaining a whole slide image classification AUC = 0.966 and a tumor localization score of 0.733. In terms of AUC score, it can be concluded that the designed CNN classifiers reached similar accuracy as a board-certified expert. Interestingly, they found that the

errors made by the pathologist and the deep learning system were not strongly correlated. Therefore, by combining their deep learning classifiers with the human pathologist's diagnoses the pathologist's accuracy was increased to AUC = 0.995, reducing in an approximately 85% human error rate. It was also noted in Ehteshami Bejnordi et al. [9] that this was the first study to recognize that ML algorithms can rival human pathologists' performance. Another interesting observation was that deep learning-based algorithms significantly outperformed other conventional ML methods, where the top 19 out of all 23 teams utilized deep CNNs.

Fig. 7. One proposed framework for cancer metastases detection by Wang et al. [60] who won the first prize in the Camelyon16 cancer detection competition [9]. The model was based on deep convolutional neural networks, GoogLeNet of 27 layers.

Discussion

Modern ML algorithms serve as powerful tools to improve medical practice by reducing human labor and possible errors. They can potentially improve a patient's diagnosis and treatment precision by complementing human perception. With the latest innovations in ML techniques, we are looking towards an exciting data revolution in the medical field in general and in oncology in particular. ML is expected to allow efficient utilization of resources and to save time and unnecessary medical expenses to patients, their doctors and the society at large.

However, and before having a full-fledged embracement of this new digital revolution, one needs to validate whether these innovative technologies can be widely adopted in medicine. For example, it may still be unclear clinically how to decide which kind of features is better suited for solving a specific diagnostic or therapeutic task using an ML algorithm (e.g., handcrafted features vs. machine-learned features). In the example of identifying

metastatic breast cancer, researchers found that the recognition of deep learning systems and human pathologists can be complementary, where the system helped reduce the pathologist error rate from over 3% to less than 1%. Therefore, at the current stage deep learning systems are more suited as a secondary opinion to aid in decision support or quality checks, rather than a stand-alone system.

One key factor for improving ML performance is the available quality and quantity of data. Before full development of artificial intelligence (as in Fig. 1), computers are unlikely able to comprehend complex physical laws out of only a few examples, but rather are only capable of deducing empirical relations based on large observations by statistical inference algorithms. In the medical field where data sets tend to be small, partially observed or labeled, and sometimes noisy, the full optimization of ML can be a challenge. In the medical imaging scope, if a given data set is too small and/or noisy, although handcrafted features can be time-consuming to generate, they can take advantage of prior domain knowledge and may outper-

form undertrained deep NN/CNN methods. However, there are a few steps that may mitigate the small data problem: (1) data preprocessing (e.g., PCA or autoencoders) for reducing the number of fitting parameters a priori; (2) data augmentation techniques such as transfer learning and generative adversarial nets in an attempt to overcome sample size limitations, as in certain experiments such as image segmentation, diagnosis and end point prediction tasks [65], data augmentation techniques have demonstrated promising results, and (3) considering the combination of traditional features and CNNs features applied to images.

Gathering and sharing data sets across institutions certainly serves as another viable way to increase data size and improve ML utility. However, certain problems may also arise: e.g., how to train models from heterogeneous data sets from different institutes that may also have varying data formats. The harmonization and standardization of naming conventions (abbreviations, code names, etc.) alone can easily confuse computers and even lead to miscalculation by various feature definitions pertaining to one's institute. Another important challenge is maintaining the confidentiality and the privacy of patient information in such data-sharing processes where administrative (institutional review boards) regulations and laws (e.g., HIPAA privacy rule) could be at risk of violation. For the purpose of privacy protection, a newly developing cryptography technology especially for datamining and statistical data queries called *differential privacy* [66, 67] can be applied to shared medical data sets, where the basic idea is the injection of a proper noise level, hashing or subsampling to scramble the original data set from possible unwarranted probing. Large companies such as Apple, Google, Facebook and Microsoft are applying and promoting such technologies for their customer privacy concerns [68]. Another approach is the utilization of distributed (rapid) learning presented in Eurocat [69, 70], where algorithms instead of data are shared across the different institutions. With all these exciting breakthroughs in ML and their potential in oncology, one still needs to be careful when wielding such methods and meticulously design the data validation experiments to avoid pitfalls of overfitting and misinformation [71].

Conclusion

The past few years have observed a tremendous rise in ML applications to a wide range of areas in oncology, including building predictive models of disease diagnosis,

treatment response, and automation of workflow and decision support. But as methods and techniques in ML keep evolving, one can expect the role of ML to continue reshaping the field of oncology and cancer management. ML is expected to alter the way patients receive treatments and doctors reach their clinical decisions. Diagnostics will be faster, cheaper and more accurate than ever. However, to usher the advent of the ML era, one has to be mindful of the characteristics and the limitations of this technology too. ML methods require large amounts of data for their training and validation, which also beg the questions of computerized trust, data sharing and privacy concerns. With the pre-existing domain knowledge, the merits of man-crafted features standing for accumulative knowledge based on numerous observations should be incorporated and inherently infused with modern ML architectures. With the assistance of the state-of-the-art ML algorithms, imaging informatics holds the potential to provide better precision health care for cancer patients as well as revealing underlying biological patterns. The application of ML algorithms in the medical realm is promising, yet there remain many challenges before they can realize their potential into routine clinical oncology practice.

Acknowledgments

This work was supported in part by the National Institutes of Health P01 CA059827.

Disclosure Statement

The authors have no conflicts of interest to declare.

References

- 1 Mitchell TM, et al. Machine learning. wcb, 1997.
- 2 Huang SH, Pan YC. Automated visual inspection in the semiconductor industry: A survey. *Comput Ind.* 2015;66:1–10.
- 3 Google LLC. Google cloud speech-to-text enables developers to convert audio to text by applying powerful neural network models in an easy to use api [Online]. 2018. Available from: <https://cloud.google.com/speech-to-text/>
- 4 Ciocca S. How does spotify know you so well? [Online]. 2017. Available from: <https://medium.com/s/story/spotify-finds-your-new-music-19a41ab76efe>.

- 5 Kooti F, Grbovic M, Aiello LM, Djuric N, Radosavljevic V, Lerman K. Analyzing Uber's ride-sharing economy. Proceedings of the 26th International Conference on World Wide Web Companion; 2017 April 3–7; Perth: International World Wide Web Conferences Steering Committee; 2017. p. 574–582.
- 6 Somashekhar SP, Sepúlveda MJ, Puglielli S, Norden AD, Shortliffe EH, Rohit Kumar C, et al. Watson for Oncology and breast cancer treatment recommendations: agreement with an expert multidisciplinary tumor board. *Ann Oncol*. 2018 Feb;29(2):418–23.
- 7 Watson for oncology and breast cancer treatment recommendations: agreement with an expert multidisciplinary tumor board. *Ann Oncol*. 2018 Feb;29(2):418–23.
- 8 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 Feb;542(7639):115–8.
- 9 Ehteshami Bejnordi B, Veta M, Johannes van Diest P, van Ginneken B, Karssemeijer N, Litjens G, et al.; the CAMELYON16 Consortium. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA*. 2017 Dec;318(22):2199–210.
- 10 Cohen JD, Li L, Wang Y, Thoburn C, Afsari B, Danilova L, et al. Detection and localization of surgically resectable cancers with a multi-analyte blood test. *Science*. 2018 Feb; 359(6378):926–30.
- 11 Wikipedia contributors. John McCarthy (computer scientist) — Wikipedia, the free encyclopedia [Online]. 2018 [cited 2018 March 16]. Available from: [https://en.wikipedia.org/w/index.php?title=John_McCarthy_\(computer_scientist\)&oldid=823313442](https://en.wikipedia.org/w/index.php?title=John_McCarthy_(computer_scientist)&oldid=823313442).
- 12 Christopher MB. *Pattern recognition and machine learning*. New York: Springer; 2016.
- 13 Michalski RS, Carbonell JG, Mitchell TM. *Machine learning: An artificial intelligence approach*. Springer Science & Business Media; 2013.
- 14 Jolliffe IT. Principal component analysis and factor analysis. *Principal component analysis*. Springer; 1986. pp. 115–28.
- 15 Belkin M, Niyogi P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput*. 2003;15(6):1373–96.
- 16 van der Maaten L, Hinton G. Visualizing data using t-SNE. *J Mach Learn Res*. 2008 Nov;9: 2579–605.
- 17 Tseng HH, Naqa IE, Chien JT. Power-law stochastic neighbor embedding. 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP); March 2017. pp. 2347–51.
- 18 Baldi P. Autoencoders, unsupervised learning, and deep architectures. Proceedings of ICML Workshop on Unsupervised and Transfer Learning; 2012. pp. 37–49.
- 19 Dawson LA, Biersack M, Lockwood G, Eisbruch A, Lawrence TS, Ten Haken RK. Use of principal component analysis to evaluate the partial organ tolerance of normal tissues to radiation. *Int J Radiat Oncol Biol Phys*. 2005 Jul; 62(3):829–37.
- 20 Ng AY, Jordan MI. On discriminative vs. generative classifiers: A comparison of logistic regression and naive Bayes. In: Dietterich TG, Becker S, Ghahramani Z, editors. *Advances in Neural Information Processing Systems*, vol 14. MIT Press; 2002. pp. 841–8. [Online]. Available from: <http://papers.nips.cc/paper/2020-on-discriminative-vs-generative-classifiers-a-comparison-of-logistic-regression-and-naive-bayes.pdf>.
- 21 Andel J, Perez MG, Negrao AI. Estimating the dimension of a linear model. *Kybernetika*. 1981;17(6):514–25 [Online]. Available from: <http://eudml.org/doc/28796>.
- 22 Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput*. 1997 Nov; 9(8):1735–80.
- 23 Sarkka S, Solin A, Hartikainen J. Spatiotemporal learning via infinite-dimensional bayesian filtering and smoothing: A look at gaussian process regression through kalman filtering. *IEEE Signal Process Mag*. 2013 Jul;30(4): 51–61.
- 24 Madhavan P. *Systems Analytics: Adaptive Machine Learning Workbook*. 1st ed. USA: CreateSpace Independent Publishing Platform; 2016.
- 25 Silver D, Huang A, Maddison CJ, Guez A, Sifre L, van den Driessche G, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*. 2016 Jan;529(7587): 484–9.
- 26 Bellman R. A markovian decision process. *Indiana Univ Math J*. 1957;6(4):679–84.
- 27 LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015 May;521(7553):436–44.
- 28 Hornik K. Approximation capabilities of multilayer feedforward networks. *Neural Netw*. 1991;4(2):251–7.
- 29 Zeiler MD, Fergus R. Visualizing and understanding convolutional networks. In: Fleet D, Pajdla T, Schiele B, Tuytelaars T, editors. *Computer Vision – ECCV 2014*. Cham: Springer International Publishing; 2014. p. 818–33.
- 30 Hubel DH, Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. *J Physiol*. 1968 Mar;195(1):215–43.
- 31 Lecun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proc IEEE*. 1998 Nov;86(11): 2278–324.
- 32 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*. Curran Associates, Inc.; 2012. pp. 1097–105. [Online]. Available from: <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>.
- 33 Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, et al. Going deeper with convolutions. 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); June 2015. pp. 1–9.
- 34 Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. CoRR. arXiv:1409.1556. 2014. [Online]. Available from: <http://arxiv.org/abs/1409.1556>.
- 35 Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: Navab N, Hornegger J, Wells WM, Frangi AF, editors. *MICCAI 2015. Medical Image Computing and Computer-Assisted Intervention*. Cham: Springer International Publishing; 2015. pp. 234–41.
- 36 El Naqa I. *A Guide to Outcome Modeling in Radiotherapy and Oncology: Listening to the Data*. Series in Medical Physics and Biomedical Engineering. CRC Press, Taylor & Francis Group; 2018.
- 37 Gillies RJ, Kinahan PE, Hricak H. Radiomics: images are more than pictures, they are data. *Radiology*. 2016 Feb;278(2):563–77.
- 38 Segal E, Sirlin CB, Ooi C, Adler AS, Gollub J, Chen X, et al. Decoding global gene expression programs in liver cancer by noninvasive imaging. *Nat Biotechnol*. 2007 Jun;25(6): 675–80.
- 39 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 Jun;5(1):4006.
- 40 Vallières M, Kay-Rivest E, Perrin LJ, Liem X, Furstoss C, Aerts HJ, et al. Radiomics strategies for risk assessment of tumour failure in head-and-neck cancer. *Sci Rep*. 2017 Aug; 7(1):10117.
- 41 Naqa IE. The role of quantitative pet in predicting cancer treatment outcomes. *Clin Transl Imaging*. 2014 Aug;2(4):305–20.
- 42 Haralick RM, Shanmugam K, Dinstein I. Textural features for image classification. *IEEE Trans Syst Man Cybern*. 1973 Nov;SMC-3(6): 610–21.
- 43 Galloway MM. Texture analysis using gray level run lengths. *Comput Graph Image Process*. 1975;4(2):172–9.
- 44 Thibault G, Fertil B, Navarro C, Pereira S, Cau P, Levy J, et al. Texture indexes and gray level size zone matrix application to cell nuclei classification. 10th International Conference on Pattern Recognition and Information Processing. 2009.
- 45 Thibault G, Angulo J, Meyer F. Advanced statistical matrices for texture characterization: application to cell classification. *IEEE Trans Biomed Eng*. 2014 Mar;61(3):630–7.
- 46 Amadasun M, King R. Textural features corresponding to textural properties. *IEEE Trans Syst Man Cybern*. 1989 Sep;19(5):1264–74.

- 47 Sun C, Wee WG. Neighboring gray level dependence matrix for texture classification. *Comput Vis Graph Image Process*. 1983; 23(3):341–52.
- 48 Zwanenburg A, Leger S, Vallières M, Löck S. Image biomarker standardisation initiative - feature definitions. CoRR. arXiv:1612.07003. 2016. [Online]. Available from: <http://arxiv.org/abs/1612.07003>.
- 49 Lowe DG. Distinctive image features from scale-invariant keypoints. *Int J Comput Vis*. 2004 Nov;60(2):91–110.
- 50 Freeman WT, Roth M. Orientation histograms for hand gesture recognition. MERL – Mitsubishi Electric Research Laboratories, Cambridge, MA, 02139. Tech. Rep. TR94-03, Dec 1994 [online]. Available from: <http://www.merl.com/publications/TR94-03/>.
- 51 Guyon I, Elisseeff A. An introduction to variable and feature selection. *J Mach Learn Res*. 2003 Mar;3:1157–82.
- 52 Baum EB, Haussler D. What size net gives valid generalization? In: *Advances in neural information processing systems*. 1989. pp. 81–90.
- 53 Vapnik V. *The Nature of Statistical Learning Theory*. New York: Springer; 2013.
- 54 Hastie T, Tibshirani R, Friedman J. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. New York: Springer; 2013.
- 55 Efron B. Estimating the error rate of a prediction rule: improvement on cross-validation. *J Am Stat Assoc*. 1983;78(382):316–31.
- 56 Breiman L, Spector P. Submodel selection and evaluation in regression. the x-random case. *Int Stat Rev*. 1992;60(3):291–319.
- 57 Kohavi R. A study of cross-validation and bootstrap for accuracy estimation and model selection. Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 2. IJCAI'95. San Francisco, CA: Morgan Kaufmann Publishers Inc.; 1995. pp. 1137–43. [Online]. Available from: <http://dl.acm.org/citation.cfm?id=1643031.1643047>.
- 58 Efron B, Tibshirani R. Improvements on cross-validation: The .632+ bootstrap method. *J Am Stat Assoc*. 1997;92(438):548–60.
- 59 Collins GS, Reitsma JB, Altman DG, Moons KG. Transparent Reporting of a multivariable prediction model for Individual Prognosis or Diagnosis (TRIPOD): the TRIPOD statement. *Ann Intern Med*. 2015 Jan;162(1):55–63.
- 60 Wang D, Khosla A, Gargeya R, Irshad H, Beck AH. Deep learning for identifying metastatic breast cancer. arXiv preprint. arXiv:1606.05718. 2016.
- 61 Antropova N, Huynh BQ, Giger ML. A deep feature fusion methodology for breast cancer diagnosis demonstrated on three imaging modality datasets. *Med Phys*. 2017 Oct; 44(10):5162–71.
- 62 Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. Proceedings of the 25th International Conference on Neural Information Processing Systems; 2012 Dec 3–6; Lake Tahoe; 2012. p. 1097–105.
- 63 Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. arXiv preprint. arXiv:1409.1556. 2014.
- 64 Wang D, Otto C, Jain AK. Face search at scale: 80 million gallery. arXiv preprint. arXiv:1507.07242. 2015.
- 65 Nie D, Trullo R, Lian J, Petitjean C, Ruan S, Wang Q, et al. English medical image synthesis with context-aware generative adversarial networks. *English Medical Image Computing and Computer Assisted Intervention – MICCAI 2017 - 20th International Conference Proceedings. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 10435 LNCS. New York: Springer; 2017. p. 417–25.
- 66 Dwork C, Roth A, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*. 2014;9(3–4):211–407.
- 67 Vu D, Slavkovic A. Differential privacy for clinical trial data: Preliminary evaluations. 2009 *IEEE International Conference on Data Mining Workshops*, Dec 2009. pp. 138–143.
- 68 Wikipedia contributors. Differential privacy. Wikipedia, the free encyclopedia [Online]. 2018 [cited 2018 Apr 4. Available from: https://en.wikipedia.org/w/index.php?title=Differential_privacy&oldid=833357507.
- 69 Lambin P, Zindler J, Vanneste B, van de Voorde L, Jacobs M, Eekers D, et al. Modern clinical research: how rapid learning health care and cohort multiple randomised clinical trials complement traditional evidence based medicine. *Acta Oncol*. 2015;54(9):1289–300.
- 70 Lambin P, Zindler J, Vanneste BG, De Voorde LV, Eekers D, Compter I, et al. Decision support systems for personalized and participative radiation oncology. *Adv Drug Deliv Rev*. 2017 Jan;109:131–53.
- 71 El Naqa I, Ruan D, Valdes G, Dekker A, McNutt T, Ge Y, et al. Machine learning and modeling: Data, validation, communication challenges. *Med Phys*. 2018 Oct;45(10):e834–40.