



# What Are Machine and Deep Learning?

1

Issam El Naqa and Martin J. Murphy

## 1.1 Overview

A machine or a deep learning algorithm is a computational process that uses input data to achieve a desired task without being literally programmed (i.e., “hard coded”) to produce a particular outcome. These algorithms are in a sense “soft coded” in that they automatically alter or adapt their architecture through repetition (i.e., experience) so that they become better and better at achieving the desired task. The process of adaptation is called *training*, in which samples of input data are provided along with desired outcomes. The algorithm then optimally configures itself so that it cannot only produce the desired outcome when presented with the training inputs, but can generalize to produce the desired outcome from new, previously unseen data. This training is the “learning” part of machine and deep learning processes. The training does not have to be limited to an initial adaptation during a finite interval. As with humans, a good algorithm can practice “lifelong” learning as it processes new data and learns from its mistakes.

There are many ways that a computational algorithm can adapt itself in response to training. The input data can be selected and weighted to provide the most decisive outcomes. The algorithm can have variable numerical parameters that are adjusted through iterative optimization. It can have a network of possible computational pathways that it arranges for optimal results. It can determine probability distributions from the input data and use them to predict outcomes.

---

I. El Naqa (✉)

Department of Machine Learning, Moffitt Cancer Center, Tampa, FL, USA

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

e-mail: [ielnaqa@med.umich.edu](mailto:ielnaqa@med.umich.edu); [Issam.elnaqa@moffitt.org](mailto:Issam.elnaqa@moffitt.org)

M. J. Murphy

Department of Radiation Oncology, Virginia Commonwealth University,  
Richmond, VA, USA

© Springer Nature Switzerland AG 2022

I. El Naqa, M. J. Murphy (eds.), *Machine and Deep Learning in Oncology, Medical Physics and Radiology*, [https://doi.org/10.1007/978-3-030-83047-2\\_1](https://doi.org/10.1007/978-3-030-83047-2_1)

3

The ideal of machine learning is to emulate the way that human beings (and other sentient creatures) learn to process sensory (input) signals in order to accomplish a goal. Traditionally, a machine learning algorithm would feed computer-extracted human-engineered patterns (features) derived from the raw data by, e.g., computer vision methods, to an algorithm to perform a designated learning task; a process colloquially referred to now as *shallow learning*. This is in contrast to a special subcategory of machine learning that allows for combined data representation (e.g., feature extraction) and task learning (e.g., classification or detection) known as deep learning. Conceptually, deep learning comprises learning methods that are provided raw data and which then automatically discover the features needed for detection or classification using the designated machine learning approach. In either learning process, the goal could be, e.g., a task in pattern recognition, in which the learner wants to distinguish apples from oranges. Every apple and orange is unique, but we are still able (usually) to tell one from the other. Rather than hard code a computer with many, many exact representations of apples and oranges, or with an exhaustive set of defining characteristics, it can be programmed to learn to distinguish them through repeated experience with actual apples and oranges. This is a good example of *supervised learning*, in which each training example of input data with features (color, shape, texture, etc.) is paired with its known classification label (apple or orange). It allows the learner to deal with similarities and differences when the objects to be classified have many variable properties within their own classes but still have fundamental qualities that identify them. Most importantly, the successful learner should be able to recognize an apple or an orange that it has never seen before.

A second type of machine learning is the so-called *unsupervised algorithm*. This might have the objective of trying to throw a dart at a bull's-eye. The device (or human) has a variety of degrees of freedom in the mechanism that controls the path of the dart. Rather than try to exactly program the kinematics *a priori*, the learner practices throwing the dart. For each trial, the kinematic degrees of freedom are adjusted so that the dart gets closer and closer to the bull's-eye. This is unsupervised in the sense that the training doesn't associate a particular kinematic input configuration with a particular outcome. The algorithm finds its own way from the training input data. Ideally, the trained dart thrower will be able to adjust the learned kinematics to accommodate, for instance, a change in the position of the target.

A third type of machine learning is *semi-supervised learning*, where part of the data is labeled, and other parts are unlabeled. In such a scenario, the labeled part can be used to aid the learning of the unlabeled part. This kind of scenario lends itself to most processes in nature and more closely emulates how humans develop their skills.

A fourth type of machine learning is *reinforcement learning*, where the algorithm learns to map inputs into optimized actions, i.e., goal-oriented tasks.

These algorithms currently represent the main categories of machine/deep learning, with supervised learning being the most common type in oncology, medical physics, and radiology with applications ranging from detection to diagnosis, drug discovery, and therapeutic interventions. However, several techniques are emerging to relieve the burden and cost of data labeling in supervised learning, including: the semi-supervised approach mentioned above, *transfer learning* (using knowledge

from other domains, such as natural images when learning medical ones), *active learning* (an interactive approach with human beings involved), and more recently *weak supervised learning*, where the labels are assumed to be imprecise or noisy.

There are two particularly important advantages to a successful algorithm. First, it can substitute for laborious and repetitive human effort. Second, and more significantly, it can potentially learn more complicated and subtle patterns in the input data than the average human observer is able to do. Both of these advantages are important to medical physics, oncology, and radiology applications. For example, the daily contouring of tumors and organs at risk during treatment planning is a time-consuming process of pattern recognition that is based on the observer's familiarity and experience with the appearance of anatomy in diagnostic images. That familiarity, though, has its limits, and consequently, there are uncertainties and inter-observer variability in the resulting contours. It is possible that an algorithm for contouring can pick up subtleties of texture or shape in one image or simultaneously incorporate data from multiple sources or blend the experience of numerous observers and thus reduce the uncertainty in the contour.

The complexity of medical physics, oncology, and radiology processes can vary and may involve several stages of sophisticated human-machine interactions and decision-making, which would naturally invite the use of machine/deep learning algorithms to optimize and automate these processes, including but not limited to computer-aided detection, diagnosis, triaging, radiation physics quality assurance, contouring and treatment planning, image-guidance, respiratory motion management, treatment response modeling, and treatment outcomes prediction.

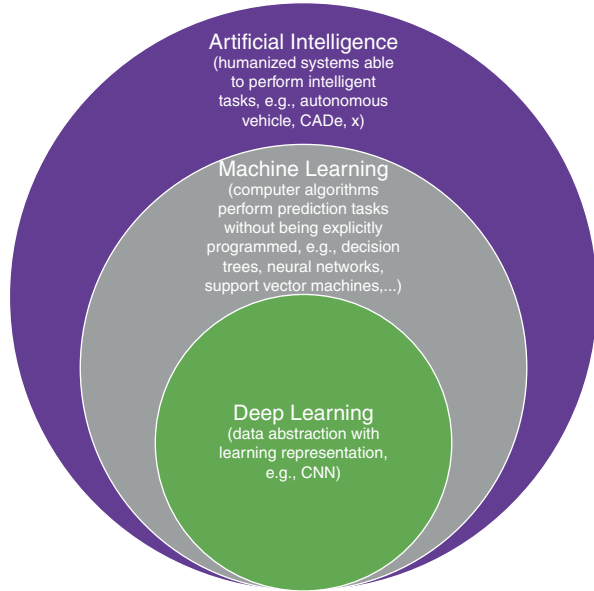
---

## 1.2 Background

Machine learning is a category of computer algorithms that are able to emulate some aspects of human intelligence. It draws on ideas from different disciplines such as artificial intelligence, probability and statistics, computer science, information theory, psychology, control theory, and philosophy [1–3]. The relationship between artificial intelligence, machine learning, and deep learning is depicted in Fig. 1.1 [4]. This technology has been applied in such diverse fields as pattern recognition [3], computer vision [5], spacecraft engineering [6], finance [7], entertainment [8, 9], ecology [10], computational biology [11, 12], and biomedical and medical applications [13, 14]. The most important property of these algorithms is their distinctive ability to learn the surrounding environment from input data with or without a teacher [1, 2].

Historically, the inception of machine learning can be traced to the seventeenth century and the development of machines that can emulate human ability to add and subtract by Pascal and Leibniz [15]. In modern history, Arthur Samuel from IBM coined the term “machine learning” and demonstrated that computers could be programmed to learn to play checkers [16]. This was followed by the development of the perceptron by Rosenblatt as one of the early neural network architectures in 1958 [17]. However, early enthusiasm about the perceptron was dampened by the

**Fig. 1.1** Venn diagram of the relationship between artificial intelligence, machine learning, and deep learning from [4]



observation made by Minsky that the perceptron classification ability is limited to linearly separable problems and not common nonlinear problems such as a simple XOR logic [18]. A breakthrough was achieved in 1975 by the development of the multilayer nonlinear perceptron (MLP) by Werbos [19]. This was followed by the development of decision trees by Quinlan in 1986 [20] and support vector machines by Cortes and Vapnik [21]. Ensemble machine learning algorithms, which combine multiple learners using boosting of weak learners or bagging (model averaging), were subsequently proposed, including Adaboost [22] and random forests [23]. More recently, distributed multilayered learning algorithms such as convolutional neural networks (CNN) have emerged under the notion of deep learning [24]. These algorithms are able to learn good representations of the data that make it easier to automatically extract useful information when building classifiers or other predictors, compared to conventional machine learning algorithms [25] as discussed further below.

### 1.3 Machine Learning Definition

The field of machine learning has received several formal definitions in the literature. Arthur Samuel in his seminal work defined machine learning as “a field of study that gives computers the ability to learn without being explicitly programmed” [16]. Using a computer science lexicon, Tom Mitchell presented it as “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  $E$ ” [1]. Ethem Alpaydin in his textbook defined machine

learning as the field of “Programming computers to optimize a performance criterion using example data or past experience” [2]. These various definitions share the notion of coaching computers to intelligently perform tasks beyond traditional number crunching by learning the surrounding environment through repeated examples. The various conventional machine learning algorithms will be reviewed in Chap. 3.


## 1.4 Deep Learning Definition

Deep learning (DL), as noted earlier, comprises a subcategory of machine learning that deals with representation learning, where raw information or data are fed directly into the algorithm, which can then automatically discover the underlying patterns (features) needed for the detection or classification task [26]. Conceptually, it can be applied to any machine learning technology as depicted in Fig. 1.2, but has been practically shown to be most effective currently with deep neural networks methods [27, 28], which will be thoroughly discussed in Chap. 4.

## 1.5 Learning from Data

The ability to learn through input from the surrounding environment, whether it is playing checkers or chess games, or recognizing written patterns, or solving the daunting problems in medical physics, oncology, or radiology, is the key to a successful machine learning application. Learning is defined in this context as estimating dependencies from data [29].

The fields of data mining and machine learning are intertwined. Data mining utilizes machine learning algorithms to interrogate large databases and discover hidden knowledge in the data, while many machine learning algorithms employ data



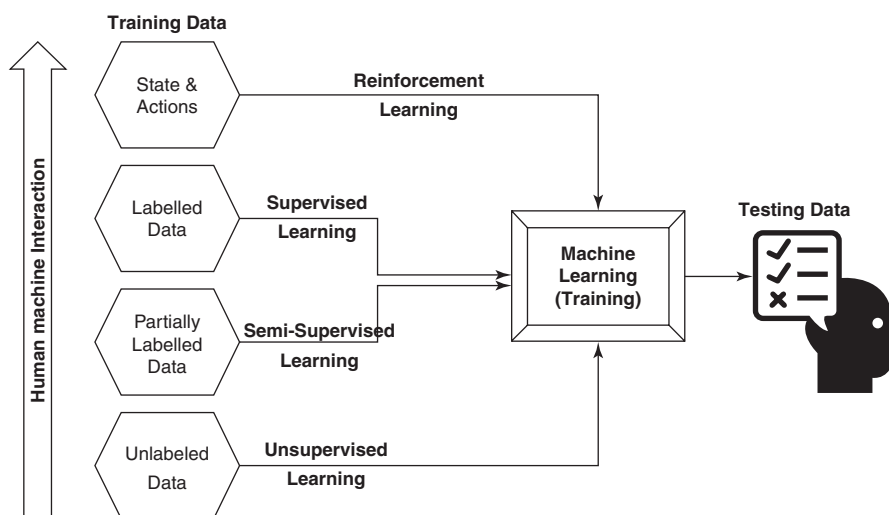
**Fig. 1.2** Conventional “shallow” machine learning (top) versus deep learning algorithms, where image data representation and classification are handled within the same framework

mining methods to preprocess the data before learning the desired tasks [30]. However, it should be noted that machine learning is not limited to solving database-like problems but also extends into solving complex artificial intelligence challenges by learning and adapting to a dynamically changing situation, as is encountered in a busy radiation oncology practice, for instance.

Machine/deep learning has both engineering science aspects such as data structures, algorithms, probability and statistics, and information and control theory and social science aspects that draw on ideas from psychology and philosophy.

## 1.6 Overview of Machine and Deep Learning Approaches

Machine or deep learning can be divided according to the nature of the data labeling into supervised, unsupervised, semi-supervised, and reinforcement learning as shown in Fig. 1.3. Supervised learning is used to estimate an unknown input-output mapping from known input-output samples, where the output is labeled (e.g., classification and regression). In unsupervised learning, only input samples are given to the learning system (e.g., clustering and estimation of probability density function). Semi-supervised learning is a combination of both supervised and unsupervised where part of the data is partially labeled and the labeled part is used to infer the unlabeled portion (e.g., text/image retrieval systems). In reinforcement learning, the machine learning algorithm aims to control learning by accommodating a feedback system, in which an agent attempts to take a sequence of actions that may maximize a cumulative reward such as winning a game of checkers, for instance [31]. This kind of approach is particularly useful for adaptive or sequential decision-making applications as will be discussed in Chap. 19.



**Fig. 1.3** Categories of machine learning algorithms according to training data nature

From a concept learning perspective, machine learning can be categorized into transductive and inductive learning [32]. Transductive learning involves the inference from specific training cases to specific testing cases using discrete labels as in clustering or using continuous labels as in manifold learning. On the other hand, inductive learning aims to predict outputs from inputs that the learner has not encountered before. Along these lines, Mitchell argues for the necessity of an inductive bias in the training process to allow for a machine learning algorithm to generalize beyond unseen observation [33].

From a probabilistic perspective, machine learning algorithms can be divided into discriminant or generative models. A discriminant model measures the conditional probability of an output given typically deterministic inputs, such as neural networks or a support vector machine. A generative model is fully probabilistic whether it is using a graph modeling technique such as Bayesian networks, or not, as in the case of naïve Bayes.

---

## 1.7 Quantifying the Data and Learning Objectives

The first step in the execution of a machine learning algorithm is the identification of the salient characteristics of the process to be emulated or the entity to be recognized or classified. These characteristics must necessarily be quantitative because this is, after all, a computational problem. The characteristics are extracted from the raw input data and then assembled into a “feature vector” that is presented to the algorithm. The extraction almost invariably involves data compression to avoid completely overwhelming the subsequent computational steps. For example, when we look at an image, we don’t see individual pixels, we see recognizable structures. The art of feature extraction is to make the algorithm “see” structures and traits in the input data. The smaller the feature vector, the better, but it is critical that it be adequate to accurately represent the data and learning objectives. The identification and quantification of the most useful features is a fundamental part of the art of designing a machine learning algorithm, which has recently been automated in the context of deep learning.

In object classification (e.g., apples and oranges), the features could be empirical attributes that are directly quantifiable, such as dimensions, weight, density, etc., or indirectly quantifiable, such as color, texture, or smell. The indirect features need to be preprocessed further to convert them to numerical measures.

Formal features can be extracted via data transformation or reduction techniques. If the raw input data have many, many discrete elements, such as pixel values in an image, then using the entire image as the feature vector would have prohibitive computational overhead. However, if those elements are not random, then the size of the input feature vector can be dramatically reduced with minimal loss by methods of dimensionality reduction and compression such as principal component analysis (PCA) or Fourier analysis. PCA transforms a complex set of correlated data elements into a set of maximally uncorrelated principal component basis vectors and their associated coefficients. A linear combination of the basis vectors and

coefficients reproduces the original data set with an accuracy that is determined by the number of vectors that are retained from the analysis. In highly correlated data, a very small number of PCA vectors and coefficients can be sufficient to characterize its structure. The most significant coefficients are then collected into the feature vector. Fourier decomposition of the input data into a set of Fourier basis vectors and coefficients achieves the same goal, but the difference is that the PCA method requires an initial set of representative training examples to determine the principal components, while Fourier decomposition can be done case by case using fixed basis vectors. The Fourier transform method lends itself naturally to image compression, as is well known from the JPEG algorithm, but it can require many more coefficients to capture salient image content than the PCA method. Both of these methods lend themselves naturally to pattern recognition and classification algorithms such as neural networks and support vector machines. Formal feature extraction or representation also lends itself naturally to deep learning applications, which automates the process by functioning as the interface between the raw input data and the learning algorithm.

---

## 1.8 Application in Biomedicine

Machine learning algorithms have witnessed increased use in biomedicine, starting naturally in neuroscience and cognitive psychology through the seminal work of Donald Hebb in his 1949 book [34] developing the principles of associative or Hebbian learning as a mechanism of neuron adaptation and the work of Frank Rosenblatt developing the perceptron in 1958 as an intelligent agent [17]. This was shortly followed by Ledley and Lusted in their 1959 paper, where they anticipated the role of a probabilistic logic-based approach to understand and support physicians' reasoning [35]. An early major machine learning initiative was the MYCIN project at Stanford in the 1970s, which was a rule-based system to identify bacteria types that may cause infectious diseases [36], achieving an acceptability rating of 65% from a panel of experts [37]. Recent reviews of the application of machine learning in biomedicine and medicine can be found in [12, 13, 38, 39].

---

## 1.9 Applications in Radiology and Oncology

Among the earliest adoptions of machine learning algorithms was in the field of radiological and medical image analysis. Winsberg et al. reported in 1967 on a computer detection algorithm for radiographic abnormalities in mammograms [40]. Lodwick et al. presented a roentgenograms concept for analyzing bone and lung cancer images [41, 42] and Meyers et al. developed an automated computer analysis of cardiothoracic ratios [43]. However, the major thrust happened in the 1980s, when tremendous developments occurred in computer-aided detection (CADe) and computer-aided diagnosis (CADx), providing radiologists with computer output as a "second opinion" to aid in making final decisions [44–49]. These CAD systems



utilized image feature-based analysis for the detection of microcalcifications in mammogram images [50–53] and lung nodules in digital chest radiographs [54]. This expanded into every area in radiology, in the form of decision support systems. In the field of oncology and specifically, radiation oncology, early applications of machine learning have focused on treatment planning and predicting normal tissue toxicity [55–57], but its application has since branched into almost every part of the field, including tumor response modeling, radiation physics quality assurance, contouring and treatment planning, image-guided radiotherapy, and respiratory motion management. Examples of the application of machine and deep learning will be the main subject of the second half of this book.

---

## 1.10 Ethical Challenges in the Application of Machine Learning

The application of machine learning in medicine has not been without challenges and even controversies. This is understandable given the data-driven nature of these algorithms and caveats related to data sharing, provenance, patient privacy, and the nature of medical data acquisition, which not only vary in technologies and parameters but also shift over time with new developments. Moreover, issues related to learning bias [58] and adversarial examples [60, 61] need to be accounted for. For instance, a machine learning algorithm developed for predicting the risk of pneumonia counter-intuitively suggested that patients with pneumonia and asthma would be at a lower risk of death than patients with pneumonia but without asthma [59]. Similar controversial examples were noted in the case of skin cancer risk prediction, where the presence of a ruler in the image may be a cue for the ML algorithm of high risk [62] or the appearance of a tube in a chest X-ray being indicative of severe lung disease [63]. These examples and others stress the importance of data quality and context when training and applying these powerful tools.


These challenges have led the Food and Drug Administration (FDA) in the United States, the European Union, and other international bodies to advocate for lawful, ethical and robust application from technological and societal perspectives. Towards this goal there have been shifts towards developing more explainable/interpretable machine learning algorithms [64], which would allow for better transparency, oversight, and accountability.

---

## 1.11 Steps to Machine Learning Heaven

For the successful application of machine learning in general and in medical physics, radiology and oncology in particular, one first needs to properly characterize the nature of the problem, in terms of the input data and the desired outputs. Secondly, despite the robustness of machine learning to noise, a good model cannot substitute for bad data, keeping in mind that models are primarily built on approximations, and it has been stated that “All models are wrong; some models are useful (George

**Fig. 1.4** GIGO paradigm. Learners cannot be better than the data



Box).” Additionally, this has been stated as the GIGO principle, garbage in garbage out as shown in Fig. 1.4 [65].

Thirdly, the model needs to generalize beyond the observed data into unseen data, as indicated by the inductive bias mentioned earlier. To achieve this goal, the model needs to be kept as simple as possible but not simpler, a property known as parsimony, which follows from Occam’s razor that “Among competing hypotheses, the hypothesis with the fewest assumptions should be selected.” Analytically, the complexity of a model could be derived using different metrics such as Vapnik–Chervonenkis (VC) dimension discussed in Chap. 2 for instance [32]. However, deep learning algorithms with their large number of layers for learning data representation and performing model prediction in the same architecture, may present a future challenge to this classical notion, but the overall objective remains the same, that is, to achieve generalizability to out-of-sample data, which should be carefully evaluated as discussed in Chap. 6. Finally, a major limitation in the adoption of machine learning in general and deep learning in particular by the larger medical community is the “black box” stigma and the inability to provide an intuitive interpretation of the learned process that could help clinical practitioners better understand their data and trust the model predictions. This is an active and necessary area of research that requires special attention from the machine learning community working in biomedicine. Solutions such as deriving proxy models, developing attention maps, providing disentangled representation or learning with known operators have been emerging to create a more interpretable/explainable machine learning paradigm [66–70].

## 1.12 Conclusions

Machine and deep learning present computer algorithms that are able to learn from the surrounding environment to optimize the solution for the task at hand. It builds on expertise from diverse fields such as artificial intelligence, probability and statistics, computer science, information theory, and cognitive neuropsychology. Machine learning algorithms can be categorized into different classes according to the nature of the data, its representation, the learning process, and the model type. Machine

learning has a long history in biomedicine, particularly in radiology, but its application in medical physics and oncology is in its infancy, with high potential and promising future to improve the safety and efficacy of clinical care and advance cancer research discovery.

---

## References

1. Mitchell TM. Machine learning. New York: McGraw-Hill; 1997.
2. Alpaydin E. Introduction to machine learning. 3rd ed. Cambridge, MA: The MIT Press; 2014.
3. Bishop CM. Pattern recognition and machine learning. New York: Springer; 2006.
4. El Naqa I, Haider MA, Giger ML, Ten Haken RK. Artificial Intelligence: reshaping the practice of radiological sciences in the 21st century. *Br J Radiol.* 2020;93:20190855.
5. Apolloni B. Machine learning and robot perception. Berlin: Springer; 2005.
6. Ao S-I, Rieger BB, Amouzegar MA. Machine learning and systems engineering. Dordrecht/New York: Springer; 2010.
7. Györfi L, Ottucsák G, Walk H. Machine learning for financial engineering. Singapore/London: World Scientific; 2012.
8. Gong Y, Xu W. Machine learning for multimedia content analysis. New York/London: Springer; 2007.
9. Yu J, Tao D. Modern machine learning techniques and their applications in cartoon animation research. 1st ed. Hoboken: Wiley; 2013.
10. Fielding A. Machine learning methods for ecological applications. Boston: Kluwer Academic; 1999.
11. Mitra S. Introduction to machine learning and bioinformatics. Boca Raton: CRC; 2008.
12. Yang ZR. Machine learning approaches to bioinformatics. Hackensack: World Scientific; 2010.
13. Cleophas TJ. Machine learning in medicine. New York: Springer; 2013.
14. Malley JD, Malley KG, Pajevic S. Statistical learning for biomedical data. Cambridge: Cambridge University Press; 2011.
15. Ifrah G. The universal history of computing: from the abacus to the quantum computer. New York: John Wiley; 2001.
16. Samuel AL. Some studies in machine learning using the game of checkers. *IBM J Res Dev.* 1959;3:210–29.
17. Rosenblatt F. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev.* 1958;65:386–408.
18. Minsky ML, Papert S. Perceptrons; an introduction to computational geometry. Cambridge, MA: MIT Press; 1969.
19. Werbos PJ. Beyond regression: new tools for prediction and analysis in the behavioral sciences. PhD thesis, Harvard University; 1974.
20. Quinlan JR. Induction of decision trees. *Mach Learn.* 1986;1:81–106.
21. Cortes C, Vapnik V. Support-vector networks. *Mach Learn.* 1995;20:273–97.
22. Schapire RE. A brief introduction to boosting. In: Proceedings of the 16th international joint conference on artificial intelligence, vol. 2. Stockholm: Morgan Kaufmann; 1999. p. 1401–6.
23. Breiman L. Random forests. *Mach Learn.* 2001;45:5–32.
24. Hinton GE. Learning multiple layers of representation. *Trends Cogn Sci.* 2007;11:428–34.
25. Bengio Y, Courville A, Vincent P. Representation learning: a review and new perspectives. *IEEE Trans Pattern Anal Mach Intell.* 2013;35:1798–828.
26. Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. *Nature.* 1986;323(6088):533–6. <https://doi.org/10.1038/323533a0>.
27. Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. *Science.* 2006;313(5786):504–7. <https://doi.org/10.1126/science.1127647>.

28. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521(7553):436–44. <https://doi.org/10.1038/nature14539>.
29. Cherkassky VS, Mulier F. Learning from data: concepts, theory, and methods. 2nd ed. Hoboken: IEEE Press/Wiley-Interscience; 2007.
30. Kargupta H. Next generation of data mining. Boca Raton: CRC Press; 2009.
31. Sutton RS, Barto AG. Reinforcement learning: an introduction. Cambridge, MA: MIT Press; 1998.
32. Vapnik VN. Statistical learning theory. New York: Wiley; 1998.
33. Mitchell TM. The need for biases in learning generalizations. New Brunswick: Rutgers University; 1980.
34. Hebb DO. The organization of behavior; a neuropsychological theory. New York: Wiley; 1949.
35. Ledley RS, Lusted LB. Reasoning foundations of medical diagnosis; symbolic logic, probability, and value theory aid our understanding of how physicians reason. *Science*. 1959;130(3366):9–21.
36. Shortliffe EH, Buchanan BG. A model of inexact reasoning in medicine. *Math Biosci*. 1975;23(3):351–79. [https://doi.org/10.1016/0025-5564\(75\)90047-4](https://doi.org/10.1016/0025-5564(75)90047-4).
37. Yu VL, Fagan LM, Wraith SM, Clancey WJ, Scott AC, Hannigan J, Blum RL, Buchanan BG, Cohen SN. Antimicrobial selection by a computer: a blinded evaluation by infectious diseases experts. *JAMA*. 1979;242(12):1279–82. <https://doi.org/10.1001/jama.1979.03300120033020>.
38. Rajkomar A, Dean J, Kohane I. Machine learning in medicine. *N Engl J Med*. 2019;380(14):1347–58. <https://doi.org/10.1056/NEJMr1814259>.
39. Saria S, Butte A, Sheikh A. Better medicine through machine learning: what’s real, and what’s artificial? *PLoS Med*. 2018;15(12):e1002721. <https://doi.org/10.1371/journal.pmed.1002721>.
40. Winsberg F, et al. Detection of radiographic abnormalities in mammograms by means of optical scanning and computer analysis. *Radiology*. 1967;89(2):211–5.
41. Lodwick GS, Keats TE, Dorst JP. The coding of Roentgen images for computer analysis as applied to lung cancer. *Radiology*. 1963;81(2):185–200.
42. Lodwick GS, et al. Computer diagnosis of primary bone tumors. *Radiology*. 1963;80(2):273–5.
43. Meyers PH, et al. Automated computer analysis of radiographic images. *Radiology*. 1964;83(6):1029–34.
44. Doi K. Computer-aided diagnosis in medical imaging: historical review, current status and future potential. *Comput Med Imaging Graph*. 2007;31(4-5):198–211.
45. Doi K, et al. Artificial intelligence and neural networks in radiology: Application to computer-aided diagnostic schemes. In: Hendee W, Trueblood J, editors. Digital imaging. AAPM medical physics monograph; 1993. p. 301–22.
46. Giger M, et al. Computer-aided diagnosis in mammography. In: Sonka M, Fitzpatrick M, editors. Handbook of medical imaging. Philadelphia, PA: SPIE; 2000. p. 915–1004.
47. Giger ML. Future of breast imaging. Computer-aided diagnosis. In: Haus A, Yaffe M, editors. AAPM/RSNA categorical course on the technical aspects of breast imaging; 1992. p. 257–70.
48. Giger ML. Computer-aided diagnosis in radiology. *Acad Radiol*. 2002;9(1):1–3.
49. Swett H, Giger M, Doi K. Computer vision and decision support. In: Hendee W, Wells P, editors. Perception of visual information. Berlin: Springer-Verlag; 1993. p. 272–315.
50. Chan HP, et al. Image feature analysis and computer-aided diagnosis in digital radiography. I. Automated detection of microcalcifications in mammography. *Med Phys*. 1987;14(4):538–48.
51. El-Naqa I, Yang Y, Wernick MN, Galatsanos NP, Nishikawa RM. A support vector machine approach for detection of microcalcifications. *IEEE Trans Med Imaging*. 2002;21:1552–63.
52. Gurcan MN, Chan HP, Sahiner B, Hadjiiski L, Petrick N, Helvie MA. Optimal neural network architecture selection: improvement in computerized detection of microcalcifications. *Acad Radiol*. 2002;9:420–9.
53. El-Naqa I, Yang Y, Galatsanos NP, Nishikawa RM, Wernick MN. A similarity learning approach to content-based image retrieval: application to digital mammography. *IEEE Trans Med Imaging*. 2004;23:1233–44.

54. Giger ML, Doi K, MacMahon H. Image feature analysis and computer-aided diagnosis in digital radiography. 3. Automated detection of nodules in peripheral lung fields. *Med Phys*. 1988;15(2):158–66.
55. Gulliford SL, Webb S, Rowbottom CG, Corne DW, Dearnaley DP. Use of artificial neural networks to predict biological outcomes for patients receiving radical radiotherapy of the prostate. *Radiother Oncol*. 2004;71:3–12.
56. Munley MT, Lo JY, Sibley GS, Bentel GC, Anscher MS, Marks LB. A neural network to predict symptomatic lung injury. *Phys Med Biol*. 1999;44:2241–9.
57. Su M, Miften M, Whiddon C, Sun X, Light K, Marks L. An artificial neural network for predicting the incidence of radiation pneumonitis. *Med Phys*. 2005;32:318–25.
58. Raji ID, Buolamwini J. Actionable auditing: investigating the impact of publicly naming biased performance results of commercial AI products. In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. Honolulu, HI: ACM; 2019. p. 429–35.
59. Caruana R, et al. Intelligible models for healthcare: predicting pneumonia risk and hospital 30-day readmission. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Sydney, NSW: ACM; 2015. p. 1721–30.
60. Biggio B, Roli F. Wild patterns: ten years after the rise of adversarial machine learning. *Pattern Recognit*. 2018;84:317–31.
61. Finlayson SG, et al. Adversarial attacks on medical machine learning. *Science*. 2019;363(6433):1287–9.
62. Esteva A, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017;542(7639):115–8.
63. Rajpurkar P, et al. CheXNet: radiologist-level pneumonia detection on chest x-rays with deep learning. In: *arXiv e-prints*; 2017.
64. Luo Y, et al. Balancing accuracy and interpretability of machine learning approaches for radiation treatment outcomes modeling. *BJR Open*. 2019;1(1):20190021.
65. Tweedie R, Mengersen K, Eccleston J. Garbage in, garbage out: can statisticians quantify the effects of poor data? *Chance*. 1994;7:20–7.
66. Philbrick KA, et al. What does deep learning see? Insights from a classifier trained to predict contrast enhancement phase from CT images. *Am J Roentgenol*. 2018;211(6):1184–93.
67. Seah JCY, et al. Chest radiographs in congestive heart failure: visualizing neural network learning. *Radiology*. 2019;290:514–22.
68. Luna JM, et al. Building more accurate decision trees with the additive tree. *Proc Natl Acad Sci U S A*. 2019;116(40):19887–93.
69. Nazmul Haque K, Latif S, Rana R. Disentangled representation learning with information maximizing autoencoder. In: *arXiv e-prints*; 2019.
70. Maier AK, et al. Learning with known operators reduces maximum error bounds. *Nat Mach Intell*. 2019;1(8):373–80.