Artificial intelligence in medicine

What oncologists need to know about its potential — and its limitations

Omar F. Khan, BComm, MD, FRCPC; Gwyn Bebb, BMBCh, PhD, FRCPC; and Nimira A. Alimohamed, MD, FRCPC

ABSTRACT

Artificial intelligence is defined as the ability of a machine to independently replicate intellectual processes typical of humans. Machine learning — the ability of computers to autonomously learn from data — is a subset of artificial intelligence. Research into machine learning in healthcare is growing at an exponential rate. This article explores various applications of machine learning, such as risk modelling (including risk of developing cancer, risk of disease recurrence, or risk of treatment complications), improving diagnostic and staging investigation accuracy, predicting prognosis, and even predicting response to therapy. Machine learning has also been applied to radiation treatment planning, and has potential applications in psychosocial oncology. Large-scale artificial intelligence platforms (such as IBM Watson for Oncology) integrate many neural networks to process natural language, generate hypotheses, and integrate this information with medical databases to create recommendations. However, artificial intelligence has several vulnerabilities, including the black-box effect, potentially prohibitive costs, an overdependence on the quality of the data used for learning, and difficulty in gaining trust and acceptance by medical professionals. Despite these limitations, the role of artificial intelligence in oncology continues to expand. We must take steps now to prepare our healthcare systems for the arrival of machine learning, which will help increase research productivity while ultimately improving our ability to diagnose, prognosticate and make treatment decisions.

Keywords: artificial intelligence; machine learning; neural networks (computer)

INTRODUCTION

In 1936, Alan Turing published a theory on computability and the "universal machine," paving the road for the advent of modern computing.[1] The field of artificial intelligence was born at the Dartmouth Conference in 1956, founded on the premise that any facet of human intellect could be described with sufficient precision to be reproducible.[2] Artificial intelligence is defined as the ability of a computer to independently replicate intellectual processes typical of humans.[3,4] A subset of artificial intelligence — known as machine learning — focuses on the ability of computers to autonomously learn from data.[5]

Artificial intelligence (AI) has evolved rapidly in scope and capacity over the last 6 decades, due to a number of technologic advances. Firstly, the evolution of the microprocessor has dramatically advanced computing speed. The first commercially available microprocessor was released in 1971, executing 60,000 instructions per second;[6,7] in contrast, Intel recently announced a processor for home computers with computing power over one teraflop (a unit representing one trillion operations per second).[8] Secondly, increasingly efficient data collection, storage and retrieval capacity led to the advent of "Big Data" (enormous data sets accumulated from a variety of sources).[9] Finally, AI techniques have evolved, including the discovery of artificial neural networks — computer modelling algorithms loosely mimicking the human brain, utilizing interconnected nodes arranged into layers[10] (**Figure 1**) — that help improve the

Corresponding author:

Omar F. Khan, BComm, MD, FRCPC, is Resident Physician in the Division of Medical Oncology, Department of Oncology, Cumming School of Medicine, University of Calgary and Tom Baker Cancer Centre
Omar.Khan@ahs.ca

Gwyn Bebb, BMBCh, PhD, FRCPC, is a Medical Oncologist and Associate Professor in the Department of Oncology at the Cumming School of Medicine, University of Calgary

Nimira A. Alimohamed, MD, FRCPC, is a Medical Oncologist and Clinical Assistant Professor in the Department of Oncology at the Cumming School of Medicine, University of Calgary

Sources of funding: None
Acknowledgements: We thank Dr. Don Morris (Division Head, Medical Oncology, Cumming School of Medicine, University of Calgary) for discussions and comments that ultimately contributed to the manuscript.

FIGURE 1. Graphic representation of an artificial neural network with 3 hidden "layers"

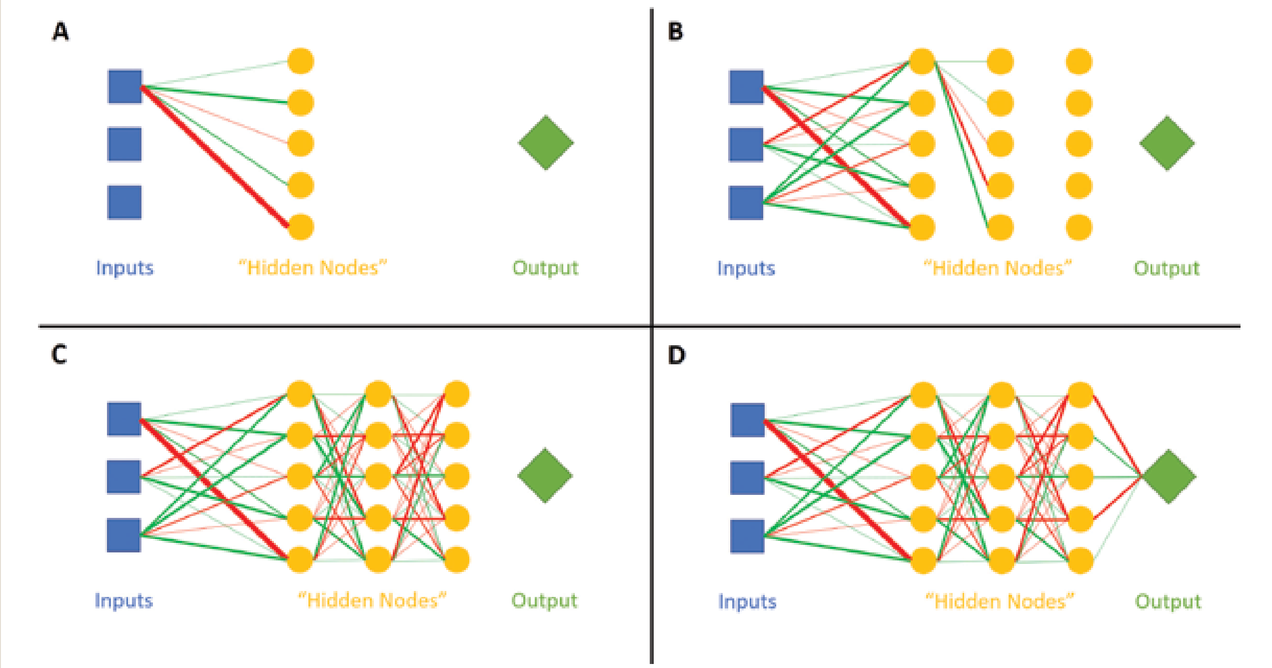

Each input variable (blue square) sends either a positive (green) or a negative (red) signal to the first layer of hidden nodes, with the weight of the line representing the strength of the signal (panel A). Information is processed in the first layer of hidden nodes, which then become the inputs for the second layer of hidden nodes (panel B), which in turn become the inputs for the third layer of hidden nodes (panel C). Data from the final layer of hidden nodes is processed to produce a final output (panel D).

accuracy of conventional modelling techniques. Unlike logistic regression techniques, neural networks and other machine learning models can model non-linear relationships between dependent and independent variables, improving the adaptability and flexibility of these models to handle complex problems.[11]

Current AI can rival (or even exceed) human ability in board games (such as chess[12] or Go[13]), television game shows like Jeopardy[14,15] and image recognition.[16] Machine learning is deployed in many sectors, including engineering, finance, art and medicine.[17-19] Other well-established applications for machine learning include facial recognition,[20] spam e-mail filtering[21] and autonomous vehicles.[22] AI and machine learning within healthcare has grown exponentially over the last several years. The growth trajectory in this field is best illustrated by the search results on the National Institute of Health's PubMed database for the term "machine learning"[23] (**Figure 2**).

This article explores the evolving role of AI and machine learning in medicine. While some examples are not from within oncology, the nature of the questions they answer are easily translatable to oncology; these examples thus serve to illustrate the breadth of potential applications of AI.

APPLYING MACHINE LEARNING TO MEDICINE

Modelling risk

Identifying patients at risk for developing a condition (as well as identifying those not at risk) can reduce costs while sparing patients from the morbidity and mortality associated with unnecessary treatment.

Using retrospective data collected from records of 378,256 patients in the United Kingdom, Weng et al. created 4 different machine-learning techniques to model cardiovascular risk.[24] These models were compared to the well-established American Heart Association/American College of Cardiology (AHA/ACC) algorithm. Overall, all 4 machine learning algorithms outperformed the AHA/ACC formula by between 3% and 5%. While some risk factors in the AHA/ACC algorithm (such as age, gender and smoking history) were also included in the machine learning algorithms, others (such as diagnosis of diabetes, and blood pressure) were not included. Interestingly, patient ethnicity and socioeconomic status were included in all 4 machine learning algorithms, but not in the AHA/ACC algorithm. In addition to improved model accuracy, this study demonstrates the ability of AI to identify associations that may not have previously been described. However, it also illustrates one of the limitations of machine learning — the "black box effect."[5] Since programmers teach neural networks how to recognize patterns (as opposed to what patterns to recognize), and as the complexity of neural networks grows, it becomes increasingly difficult to understand why they make certain associations. While some associations have a plausible biological basis, others (such as the neural network in this study identifying a lack of documented body mass index as a protective factor) are less clear.

AI can also help model risks of treatment complications. Neural networks were utilized to predict the risk of 30-day morbidity after surgery to treat head and neck squamous cell carcinoma.[25] Auditing 1075 patient charts, risk models created using logistic regression were compared with AI techniques, including a neural network. The artificial neural network outperformed all other models, with an area under the curve c-statistic (AUC) of 0.85 (the closer an AUC is to 1.0, the more accurate the model). Another study using neural networks to predict pneumonectomy complications after surgical resection for non-small cell lung cancer reported an AUC of 0.98.[26]

For resected malignancies, accurate prediction of recurrence risk is critical and informs the risk-benefit ratio for adjuvant therapies. Kim et al. applied machine learning to predict risk of recurrence after breast cancer surgical resection, using longitudinal observational data from 733 Korean patients.[11] They found their AI models (AUCs between 0.80 and 0.85) outperformed Cox regression analysis (AUC 0.73), the Adjuvant! Online tool (AUC 0.71) and the Nottingham Prognostic Index (AUC 0.70).

FIGURE 2. Number of publications (sorted by year of publication) retrieved during a PubMed search for "machine learning."[23]

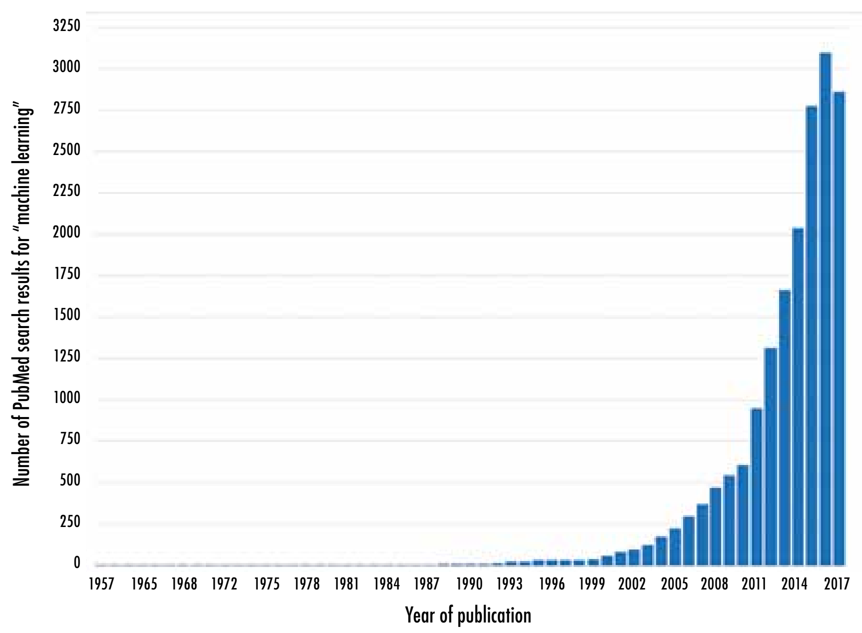

Improving diagnostic and staging accuracy

Machine learning techniques are being applied to diagnostic investigations. In a recent study, 2 artificial neural networks (AlexNet and GoogLeNet) classified chest x-ray images of 1007 patients (492 with tuberculosis, and 515 healthy controls) from Belarus, China and the United States.[27] In the first part of the study, AlexNet and GoogLeNet were immediately used to differentiate "tuberculosis" chest x-rays from "healthy" ones. The neural networks were trained and internally validated using 857 patients, and their accuracy was then tested on the remaining 150 patients. Using this model, AlexNet had an AUC of 0.90, while GoogLeNet had an AUC of 0.88 to predict tuberculosis. In the second part of the study, after resetting AlexNet and GoogLeNet to erase their memory, the researchers first asked the neural networks to classify 1.2 million "everyday" images (e.g. images of computers, people, plants, etc.), including "augmented" images that were rotated and/or cropped. The "pretrained augmented" neural networks were then taught and tested on chest x-ray classification. This time, both AlexNet and GoogLeNet had an AUC of 0.98, and the combination of the two neural networks had an AUC of 0.99. The value of pretraining the neural networks illustrates the concept of machine learning — both AlexNet and GoogLeNet learned something from classifying random images, which they applied to improve their ability to diagnose tuberculosis.

Machine learning is also being applied to pathologic diagnosis and grading of cervical cancer,[28] breast cancer[29] and gliomas.[30] Additionally, in a study of 85 bladder wash samples, artificial neural networks had an AUC of 0.71 to predict a positive cystoscopy for urothelial cell carcinoma, compared to an AUC of 0.58 for light microscopy by an experienced cytopathologist.[31] Similarly, researchers at Google trained a neural network to identify breast cancer lymph node metastases with a sensitivity of 92.4%, as compared to 73.2% sensitivity for a human pathologist.[32]

AI also can accurately stage cancers, significantly impacting patient selection for appropriate therapy. Toney and Vesselle utilized an artificial neural network to predict surgical-pathologic nodal staging of non-small cell lung cancer in 133 patients using fluorodeoxyglucose positron emission tomography (FDG-PET) scan data.[33] An expert radiologist interpreted FDG-PET scans to identify clinical nodal stage; they then provided data to the neural network, including primary tumour size and intensity, as well as the largest lymph node size and intensity at the N1, N2 and N3 stations (regardless of whether the lymph node was pathologically enlarged or not). Based on pathologic nodal staging after mediastinoscopy, the expert FDG-PET reader had a staging accuracy of 72.4 ± 3.7%, while the artificial neural network had a staging accuracy of 99.2 ± 0.8% ($p<0.001$). When differentiating resectable disease (i.e. stage N0 or N1) from unresectable disease (i.e. stage N2 or N3), the expert FDG-PET reader performed better (with an accuracy of 92.2 ± 2.3%), but were still outperformed by the neural network, which had an accuracy of 99.2 ± 0.8% ($p<0.001$).

Predicting prognosis and response to therapy

Artificial neural networks are used to predict patient prognosis. One study successfully utilized gene expression data available from 440 non-small cell lung cancer patients treated in various American cancer centres.[34] Constructing an artificial neural network, the authors achieved an 83.0% accuracy in classifying patients into low-risk (with a median overall survival of 55.7 months) and high-risk (with a median overall survival of 16 months) groups ($p<0.00001$). Another study utilized machine learning algorithms to fully automate pathology slide analysis, and used this data to predict survival among stage I lung adenocarcinoma patients more accurately than could be predicted by tumour grade.[35]

Some studies also demonstrate an ability to predict treatment response with machine learning. One small study assessed the ability of dynamic contrast-enhanced magnetic resonance imaging (MRI) features to predict responsiveness to neoadjuvant chemotherapy in 68 breast cancer patients.[36] Patients had dynamic contrast-enhanced MRI scans performed at diagnosis and repeated after completion of neoadjuvant therapy. Clinical tumour response was determined using RECIST 1.1 (Response Evaluation Criteria In Solid Tumours) criteria. While regression analysis techniques were effective (AUC 0.85 ± 0.05) at predicting a complete response or non-response to chemotherapy, artificial neural networks (AUC 0.96 ± 0.03) were significantly better ($p<0.01$), with a sensitivity of 88% and a specificity of 98%.

Radiation treatment planning

Ibragimov and Xing designed artificial neural networks to identify organs-at-risk on computed tomography (CT) simulation scans done for head and neck radiation treatment planning.[37] While their neural network performed better in identifying the spinal cord, mandible, pharynx, optic nerves and orbits, it was inferior at segmenting objects with poorly recognizable boundaries (such as parotid glands, submandibular glands and the optic chiasm). This illustrates some of the limitations of neural networks — much like the humans they are modelled after, they are limited by the quality of information provided to them (in this case, by the ability of CT scan imaging to differentiate structures with similar densities).

Drug discovery and development

The drug discovery and development process is costly and time-consuming; AI can help screen for molecules more likely to be effective, significantly cutting costs and speeding up the research process. For example, the DeepVS neural network was used to predict molecular docking of a test compound into a specified biologic target.[38] Using a test set of 40 receptors, DeepVS outperformed other, more labour-intensive virtual docking programs (which rely on manually-derived scoring functions), while requiring relatively minimal human intervention.

Psychosocial oncology

Within the psychosocial sphere, artificial neural networks have analyzed online breast cancer forums to identify trends discussed by patients. Zhang et al. analyzed more than 3.2 million posts authored by over 58,000 members on the breastcancer.org forum, using artificial neural networks to identify common themes discussed by members.[39] They found participants were likely to join discussion forums seeking informational support, but that over time their posts were more likely to relate to emotional support and exchange of personal stories. While this study focused on the ability of neural networks to accurately classify posts, the authors suggest machine learning algorithms like theirs can analyze patient replies and comments to assess the efficacy of online psychosocial interventions.

Another study assessed voice recorded interactions between 91 breast cancer patients and their healthcare providers.[40] The authors created machine learning algorithms that used implicit emotional "cues" and explicit emotional "concerns" identified by patients to predict healthcare provider responses (either as "providing space" or "reducing space" for patient emotional disclosure), with an AUC of up to 0.81. Machine learning in this study helped identify which cues were more likely to be missed by healthcare providers, presenting opportunities to train healthcare providers more effectively.

IBM Watson for Oncology

While the above examples of AI utilized single algorithms to answer their research questions, the process of performing an oncologic consult is far more complex. IBM Watson can process 500 gigabytes per second — approximately the amount of data contained in one million books.[41] It consists of layer upon layer of neural networks, allowing it to process natural language, generate hypotheses, and integrate this information with available databases. After being trained by oncologists from Memorial Sloan Kettering Cancer Center, Watson for Oncology is now deployed in several cancer centres around the world. Compared with multidisciplinary tumour board recommendations for 362 patients treated in hospitals in India, IBM Watson's recommendations were concordant in 96.4% of lung, 81.0% of colon and 92.7% of rectal cancer cases (for colon cancer, Watson recommended treatments not available in India up to 11% of the time, likely contributing to the lower concordance rate for that tumour site).[42]

Criticisms of Watson include its reliance on Memorial Sloan Kettering Cancer Center (potentially establishing their treatment recommendations as a de facto gold standard, even if evidence supporting those recommendations is weak). Additionally, in many cases, the best treatment option is clear, decreasing the value of Watson's recommendations. A partnership between MD Anderson Cancer Center and IBM Watson fell through in February 2017 (after spending $62 million USD), due to cost overruns and delays; contributing to this failure was the decision by MD Anderson to move to a new electronic medical record, which ultimately was incompatible with Watson.[43,44] Finally, Watson has not undergone validation in prospective clinical trials to prove its efficacy in altering patient outcomes or treatment decisions.

CURRENT AND FUTURE DIRECTIONS IN CANADA

Within Canada, numerous machine learning projects are currently underway. In British Columbia, the Shah Lab for Computational Cancer Biology used machine learning algorithms to predict cell migration patterns in ovarian cancer.[45] The Alberta Machine Intelligence Institute created a patient-specific survival prediction tool, available online at http://pssp.srv.ualberta.ca.[46] The Princess Margaret Bioinformatics and Computational Genomics Laboratory in Toronto is using machine learning to assess pharmacogenomic drug classification, cancer subtyping, and identification of predictive and prognostic markers.[47-49]

AI technologies have several strengths — they can process extensive amounts of data at incredible speed, model complex nonlinear relationships between variables, and update themselves as new data emerges. However, several limitations exist. As a machine learning algorithm gets progressively more accurate at analyzing the data used to train it, its results become less generalizable to the larger population (known as the overfitting problem[50]). Interestingly, AI algorithms may also acquire the inherent biases of their human trainers — a recent report found machine learning algorithms used to predict reoffending risk in American prisoners were more likely to misclassify African-American individuals as high risk, while more often misclassifying Caucasian individuals as low risk.[51]

Several barriers to implementing machine learning programs in Canadian cancer centres exist. The first of these barriers is data quality — while most provincial health systems are moving towards electronic health records, the degree to which these have been implemented is highly variable. Most data are thus locked within the confines of paper charts and dictated notes, rendering them inaccessible to electronic algorithms. Additionally, privacy and confidentiality issues regarding sharing of health information are a significant concern. Solutions to the data problem will need to be multifactorial — e.g. utilizing natural-language and handwriting processing algorithms to convert dictated and paper notes into usable data, codifying electronic health record data into analyzable forms, and collaborating to standardize data collection and reporting procedures locally, provincially and nationally.

Cost is another barrier. However, single-payer Canadian healthcare systems contain large patient datasets, and can leverage this resource to drive costs down. As machine learning algorithms depend on the amount and quality of data they learn from, partnerships between the public and private sector can be mutually beneficial — with the healthcare system providing anonymized data, and the private sector providing the technology at a discounted cost. Finally, costs may be offset somewhat (at least theoretically) by improved patient selection for appropriate therapies, though further research will be required to confirm this.

Another major consideration is the acceptance of AI by medical professionals and patients. The degree to which people trust automated systems depends on human factors (including personality, self-confidence and fatigue), machine factors (such as system reliability) and environmental factors (such as workload).[52] In particular, the transparency of logic behind an automated system's decision correlates with its perceived trustworthiness.[53] As a result, the black box effect (due to its lack of transparency) is a major hindrance to widespread adoption of neural networks, though dependably producing accurate results can overcome this hesitation. Machine learning is most trusted when using data with clear, consistent differences between normal and abnormal parameters (such as lab tests). Fostering trust in AI is more difficult in fields like diagnostic imaging, where "normal" can take on many different forms; in pathology, a near-infinite range of images can occur from the same representative tissue sample, due to non-uniform slide staining and other technical factors. Thus, adoption of AI is likely to be gradual, and its function will likely be to augment human expertise, rather than to replace it.

To successfully move AI forward within oncology, several steps are critical. Firstly, we must increase our understanding of benefits and drawbacks associated with AI. Secondly, we must forge partnerships with experts in information technology, both public and private. Thirdly, electronic health records must be planned years in advance, in tandem with AI systems, to store data in easily usable and retrievable formats. Fourthly, machine learning capabilities must be validated prospectively. Finally, machine learning algorithms need to be available for use in other centres, allowing for external validation of data, and ultimately promoting improved health outcomes for patients.

CONCLUSIONS

This article illustrates the incredible breadth of applications for AI and machine learning within oncology. While many projects currently underway are highlighted here, machine learning can potentially help answer many, possibly any future research questions. Over the coming years, AI will help increase research productivity and improve our ability to diagnose, prognosticate and make treatment decisions. By adopting these rapidly changing technologies now, oncologists can learn to leverage their capabilities to improve patient care.

In 1964, when asked what he thought about computers, Pablo Picasso replied, "but they are useless… they can only give you answers." His words still ring true today — for all its answers, AI still does not ask the questions; at least for now, curiosity remains inherently human.

References

1. Hodges A. Computer science. Beyond Turing's machines. *Science.* 2012;336(6078):163-164. doi:10.1126/science.1218417.
2. Cotter TS. RESEARCH AGENDA INTO HUMAN-INTELLIGENCE/MACHINE-INTELLIGENCE GOVERNANCE. https://search-proquest-com.ezproxy.lib.ucalgary.ca/docview/1751245873/fulltextPDF/BD61C485FF1743C1PQ/1?accountid=9838. Accessed October 7, 2017.
3. Hamet P, Tremblay J. Artificial intelligence in medicine. *Metabolism.* 2017;69:S36-S40. doi:10.1016/j.metabol.2017.01.011.
4. Copeland BJ. artificial intelligence | Definition, Examples, and Applications | Britannica.com. In: Encyclopedia Britannica. https://www.britannica.com/technology/artificial-intelligence. Accessed October 8, 2017.
5. Deo RC. Machine Learning in Medicine. *Circulation.* 2015;132(20):1920-1930. doi:10.1161/CIRCULATIONAHA.115.001593.
6. Aspray W. The Intel 4004 microprocessor: what constituted invention? *IEEE Ann Hist Comput.* 1997;19(3):4-15. doi:10.1109/85.601727.
7. Faggin F, Hoff ME, Mazor S, Shima M. The history of the 4004. *IEEE Micro.* 1996;16(6):10-20. doi:10.1109/40.546561.

8. Tung L. Intel unveils monster 18-core Core i9: "First teraflop-speed" consumer PC chip | ZDNet. ZDNet. http://www.zdnet.com/article/intel-unveils-monster-18-core-core-i9-first-teraflop-speed-consumer-pc-chip/. Published 2017. Accessed October 14, 2017.
9. Barton AJ. Big Data. *J Nurs Educ*. 2016;55(3):123-124. doi:10.3928/01484834-20160216-01.
10. Walczak S, Velanovich V. An Evaluation of Artificial Neural Networks in Predicting Pancreatic Cancer Survival. *J Gastrointest Surg*. 2017;21(10):1606-1612. doi:10.1007/s11605-017-3518-7.
11. Kim W, Kim KS, Lee JE, et al. Development of Novel Breast Cancer Recurrence Prediction Model Using Support Vector Machine. *J Breast Cancer*. 2012;15(2):230. doi:10.4048/jbc.2012.15.2.230.
12. Bloomfield BP, Vurdubakis T. IBM's Chess Players: On AI and Its Supplements. *Inf Soc*. 2008;24(2):69-82. doi:10.1080/01972240701883922.
13. Silver D, Huang A, Maddison CJ, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*. 2016;529(7587):484-489. doi:10.1038/nature16961.
14. Ferrucci DA. Introduction to "This is Watson." *IBM J Res Dev*. 2012;56(3.4):1:1-1:15. doi:10.1147/JRD.2012.2184356.
15. Chandrasekar R. Elementary? Question answering, IBM's Watson, and the Jeopardy! challenge. *Resonance*. 2014;19(3):222-241. doi:10.1007/s12045-014-0029-7.
16. He K, Zhang X, Ren S, Sun J. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. 2015:1026-1034. https://www.cv-foundation.org/openaccess/content_iccv_2015/html/He_Delving_Deep_into_ICCV_2015_paper.html. Accessed October 7, 2017.
17. Hassabis D, Kumaran D, Summerfield C, Botvinick M. Neuroscience-Inspired Artificial Intelligence. *Neuron*. 2017;95(2):245-258. doi:10.1016/j.neuron.2017.06.011.
18. Cao MS, Pan LX, Gao YF, et al. Neural network ensemble-based parameter sensitivity analysis in civil engineering systems. *Neural Comput Appl*. 2017;28(7):1583-1590. doi:10.1007/s00521-015-2132-4.
19. Erdal HI, Ekinci A. A Comparison of Various Artificial Intelligence Methods in the Prediction of Bank Failures. *Comput Econ*. 2013;42(2):199-215. doi:10.1007/s10614-012-9332-0.
20. Amos B, Ludwiczuk B, Satyanarayanan M. OpenFace: A general-purpose face recognition library with mobile applications. 2016. http://cmusatyalab.github.io/openface/. Accessed October 8, 2017.
21. Barushka A, Hájek P. Spam Filtering Using Regularized Neural Networks with Rectified Linear Units. In: Springer, Cham; 2016:65-75. doi:10.1007/978-3-319-49130-1_6.
22. Fleetwood J. Public Health, Ethics, and Autonomous Vehicles. *Am J Public Health*. 2017;107(4):532-537. doi:10.2105/AJPH.2016.303628.
23. US National Library of Medicine National Institutes of Health. Machine Learning - PubMed. https://www.ncbi.nlm.nih.gov/pubmed/?term=Machine+Learning. Published 2017. Accessed October 8, 2017.
24. Weng SF, Reps J, Kai J, Garibaldi JM, Qureshi N. Can machine-learning improve cardiovascular risk prediction using routine clinical data? Liu B, ed. PLoS One. 2017;12(4):e0174944. doi:10.1371/journal.pone.0174944.
25. Tighe D, Thomas A, Sassoon I, Kinsman R. Developing a risk stratification tool for audit of outcome after surgery for head and neck squamous cell carcinoma. Head. 2017. http://onlinelibrary.wiley.com/doi/10.1002/hed.24769/full. Accessed September 15, 2017.
26. Santos-García G, Varela G, Novoa N, Jiménez MF. Prediction of postoperative morbidity after lung resection using an artificial neural network ensemble. *Artif Intell Med*. 2004;30(1):61-69. http://www.ncbi.nlm.nih.gov/pubmed/14684265. Accessed October 8, 2017.
27. Lakhani P, Sundaram B. Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks. *Radiology*. 2017;284(2):574-582. doi:10.1148/radiol.2017162326.
28. Chankong T, Theera-Umpon N, Auephanwiriyakul S. Automatic cervical cell segmentation and classification in Pap smears. Comput Methods Programs *Biomed*. 2014;113(2):539-556. doi:10.1016/j.cmpb.2013.12.012.
29. Dey P, Logasundaram R, Joshi K. Artificial neural network in diagnosis of lobular carcinoma of breast in fine-needle aspiration cytology. *Diagn Cytopathol*. 2013;41(2):102-106. doi:10.1002/dc.21773.
30. Ertosun MG, Rubin DL. Automated Grading of Gliomas using Deep Learning in Digital Pathology Images: A modular approach with ensemble of convolutional neural networks. AMIA . Annu Symp proceedings AMIA Symp. 2015;2015:1899-1908. http://www.ncbi.nlm.nih.gov/pubmed/26958289. Accessed October 8, 2017.
31. Vriesema JL, van der Poel HG, Debruyne FM, Schalken JA, Kok LP, Boon ME. Neural network-based digitized cell image diagnosis of bladder wash cytology. *Diagn Cytopathol*. 2000;23(3):171-179. http://www.ncbi.nlm.nih.gov/pubmed/10945904. Accessed October 8, 2017.
32. Liu Y, Gadepalli K, Norouzi M, et al. Detecting Cancer Metastases on Gigapixel Pathology Images. Mountain View, California; 2017. https://drive.google.com/file/d/0B1T58bZ5vYa-QlR0QlJTa2dPWVk/view. Accessed October 8, 2017.
33. Toney L, Vesselle H. Neural networks for nodal staging of non–small cell lung cancer with FDG PET and CT: importance of combining uptake values and sizes of nodes and primary tumor. *Radiology*. 2014. http://pubs.rsna.org/doi/abs/10.1148/radiol.13122427. Accessed September 15, 2017.
34. Chen Y-C, Ke W-C, Chiu H-W. Risk classification of cancer survival using ANN with gene expression data from multiple laboratories. *Comput Biol Med*. 2014;48:1-7. doi:10.1016/J.COMPBIOMED.2014.02.006.
35. Yu K-H, Zhang C, Berry GJ, et al. Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features. *Nat Commun*. 2016;7:12474. doi:10.1038/ncomms12474.
36. Aghaei F, Tan M, Hollingsworth AB, Qian W, Liu H, Zheng B. Computer-aided breast MR image feature analysis for prediction of tumor response to chemotherapy. *Med Phys*. 2015;42(11):6520-6528. doi:10.1118/1.4933198.
37. Ibragimov B, Xing L. Segmentation of organs at risks in head and neck CT images using convolutional neural networks. *Med Phys*. 2017. http://onlinelibrary.wiley.com/doi/10.1002/mp.12045/full. Accessed September 15, 2017.
38. Pereira JC, Caffarena ER, dos Santos CN. Boosting Docking-Based Virtual Screening with Deep Learning. *J Chem Inf Model*. 2016;56(12):2495-2506. doi:10.1021/acs.jcim.6b00355.
39. Zhang S, Grave E, Sklar E, Elhadad N. Longitudinal analysis of discussion topics in an online breast cancer community using convolutional neural networks. *J Biomed Inform*. 2017;69:1-9. doi:10.1016/J.JBI.2017.03.012.
40. Barracliffe L, Arandjelovic O, Humphris G. A Pilot Study of Breast Cancer Patients: Can Machine Learning Predict Healthcare Professionals' Responses to Patient Emotions? St Andrews, United Kingdom; 2016. https://www.researchgate.net/profile/Ognjen_Arandjelovic/publication/312900919_A_Pilot_Study_of_Breast_Cancer_Patients_Can_Machine_Learning_Predict_Healthcare_Professionals'_Responses_to_Patient_Emotions/links/58991dd9a6fdcc32dbdd100e/A-Pilot-Study-of-Bre. Accessed October 8, 2017.
41. Rennie J. How IBM's Watson Computer Excels at Jeopardy! | Retort. PLOS Blogs. http://blogs.plos.org/retort/2011/02/14/how-ibm's-watson-computer-will-excel-at-jeopardy/. Published 2011. Accessed October 8, 2017.
42. Somashekhar SP, Sepulveda M-J, Norden AD, et al. Early experience with IBM Watson for Oncology (WFO) cognitive computing system for lung and colorectal cancer treatment. In: ASCO. Chicago, IL; 2017. doi:10.1200/JCO.2017.35.15_SUPPL.8527.
43. Schmidt C. M. D. Anderson Breaks With IBM Watson, Raising Questions About Artificial Intelligence in Oncology. *JNCI J Natl Cancer Inst*. 2017;109(5). doi:10.1093/jnci/djx113.
44. Harper M. MD Anderson Benches IBM Watson In Setback For Artificial Intelligence In Medicine. Forbes. https://www.forbes.com/sites/matthewherper/2017/02/19/md-anderson-benches-ibm-watson-in-setback-for-artificial-intelligence-in-medicine/#76d01e813774. Published 2017. Accessed October 14, 2017.
45. Roth A, McPherson A, Laks E, et al. Clonal genotype and population structure inference from single-cell tumor sequencing. *Nat Methods*. 2016;13(7):573-576. doi:10.1038/nmeth.3867.
46. Yu C-N, Greiner R, Lin H-C, Baracos V. Learning Patient-Specific Cancer Survival Distributions as a Sequence of Dependent Regressors. 2011:1845-1853. http://papers.nips.cc/paper/4210-learning-patient-specific-cancer-survival-distributions-as-a-sequence-of-dependent-regressors. Accessed October 8, 2017.
47. Grossmann P, Stringfield O, El-Hachem N, et al. Defining the biological basis of radiomic phenotypes in lung cancer. Elife. 2017;6:e23421. doi:10.7554/eLife.23421.
48. Safikhani Z, Thu KL, Silvester J, et al. Gene isoforms as expression-based biomarkers predictive of drug response in vitro. doi.org. July 2017:160937. doi:10.1101/160937.
49. El-Hachem N, Gendoo DMA, Ghoraie LS, et al. Integrative Cancer Pharmacogenomics to Infer Large-Scale Drug Taxonomy. Cancer Res. 2017;77(11):3057-3069. doi:10.1158/0008-5472.CAN-17-0096.
50. Dietterich T. Overfitting and Undercomputing in Machine Learning. http://delivery.acm.org.ezproxy.lib.ucalgary.ca/10.1145/220000/212114/p326-dietterich.pdf?ip=136.159.235.223&id=212114&acc=ACTIVE SERVICE&key=FD0067F557510FFB.49D736CDF0172E98.4D4702B0C3E38B35.4D4702B0C3E38B35&CFID=819016510&CFTOKEN=33195942&__acm__=1507970993_a128664b6daa7eaf28b816cab003bba8. Accessed October 14, 2017.
51. Angwin J, Mattu S, Kirchner L. Machine Bias — ProPublica. ProPublica. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing. Published 2016. Accessed October 14, 2017.
52. Chen JYC, Barnes MJ. Human–Agent Teaming for Multirobot Control: A Review of Human Factors Issues. IEEE Trans Human-Machine Syst. 2014;44(1):13-29. doi:10.1109/THMS.2013.2293535.
53. Lyons JB, Koltai KS, Ho NT, Johnson WB, Smith DE, Shively RJ. Engineering Trust in Complex Automated Systems. 2016. http://journals.sagepub.com.ezproxy.lib.ucalgary.ca/doi/pdf/10.1177/1064804615611272. Accessed October 15, 2017.