1 The Cognitive Revolution

Thomas M. Ward and Ozanan Meireles

Introduction

In the last decade, we have been witnessing the great potential of a Cognitive Revolution in Medicine that promises to completely transform surgery. Before exploring its history and potential, we must establish some basic definitions. First, cognition is "the action or faculty of knowing" [1]. Cognition is fundamental, but formal study only started in the 1950s [2]. Recent progress in one of its major subfields, artificial intelligence (AI), has created the promise of revolution.

A revolution is a "dramatic or wide-reaching change in conditions" [1]. Examples include the Industrial Revolution, which jump-started modern society with its transition of manual labor into machine-assisted processes. Surgery has also undergone many revolutions. In the nineteenth century, the development of general anesthesia and asepsis allowed for surgeons to humanely, and safely, foray into invasive surgical procedures. The twentieth-century innovations of surgical staplers, endoscopy, and laparoscopy have created modern surgery as we know it [3]. Despite this progress, surgery is still fraught with dangers: almost 30% of surgical patients will suffer a complication [4]. Surgery needs to improve, and AI offers a potential solution: the Cognitive Revolution.

Artificial Intelligence

AI is "the study of computations that make it possible to perceive, reason, and act" [5]. The breadth and extent of these computational abilities lead to different types of AI. Movies and popular science portray AI as computers and automatons with cognition equivalent to humans'. This all-encompassing AI is termed *generalized* AI [6]. Taken to the extreme, some people even believe that AI will obtain superhuman intelligence and end humanity's reign, with an event known as the *singularity* [7]. Despite Hollywood's hyperboles, a more realistic and obtainable intelligence is a *narrow* one, where computer algorithms focus on specific tasks and excel. Narrow AI is pervasive throughout today's society, from movie recommendation systems to autonomous vehicles.

Narrow AI aligns closely to Warren McCulloch and Walter Pitts' original conception of AI in 1943. Based upon a knowledge of basic neurophysiology, propositional logic, and Turing's theory of computation, they proposed that any function could be computed with a network of neurons that were either on or off [8]. Minsky and Edmonds made this "neural network computer" a

T. M. Ward (✉)
Department of Surgery, Massachusetts General Hospital, Boston, MA, USA
e-mail: tmward@mgh.harvard.edu

O. Meireles
Surgical and Innovation Laboratory, Department of Surgery, Massachusetts General Hospital, Boston, MA, USA

S. Atallah (ed.), *Digital Surgery*, https://doi.org/10.1007/978-3-030-49100-0_1

reality in 1950 with their SNARC, a computer that simulated 40 neurons [9]. These neural networks led to initial great success in the 1950s, with an early prototype capable of winning games of Checkers [10]. Unfortunately, the early gains and lofty promises failed to deliver in the ensuing decades. Public and private sector opinion soured on AI, as epitomized by the Lighthill Report in 1973 which led to almost total cessation of governmental funding for AI in the United Kingdom [11]. The ensuing "AI Winter" brought with it a near total collapse of the AI industry in the 1980s and early 1990s [12]. As the 1990s progressed though, AI began to succeed, particularly with narrow tasks. Four decades after Samuel's machine learning success with Checkers, IBM created "Deep Blue," a chess computer capable of beating the World Chess Champion Garry Kasparov [13]. This landmark achievement was one of many to follow in the years ahead.

Artificial Intelligence Revolution

AI in the twentieth century had been a plodding field, filled with many promises but few results. AI's fortunes, though, have changed in the past two decades. AI, particularly *narrow* AI, is undergoing a resurgence and revolution – but why now? Its success results from the alignment of four major factors: (1) big data, (2) adequate compute power, (3) deep learning algorithms, and (4) increased investment.

The first key to the AI revolution is big data. Data inputs form the foundation for AI and its subfield – machine learning (ML). Without data, the algorithms cannot learn. Early successes in ML came with application to problems that have a small, finite data space. For example, Tic-Tac-Toe only has nine squares to fill with two possible markers (X or O), with only a thousand legal possible different positions. Checkers has over 10^{20} possibilities, and the game Go has 10^{170} possibilities [14, 15]. Mapping these possibilities is relatively easy for a computer but imagine the amount of data needed for an algorithm to not just play a board game, but classify objects, understand human speech, or even operate a motor vehicle.

Prior to 2003, humanity had generated 5 exabytes ($5 * 10^{18}$) of data. As of 2012, that much data was generated every 2 days [16]. Medicine has seen a similar explosion in data availability, with 1 minute of high-definition surgical video containing 25 times the amount of information in one CT scan image [17]. This glut of data has given AI and ML algorithms the information they need to learn and perform at human levels.

The second key to the AI revolution is adequate compute power. AI and ML algorithms, particularly those of the "deep learning" variety, are extraordinarily resource intensive. Much of AI's failure to launch in the mid-to-late twentieth century stemmed from lack of compute power. As Moore famously postulated in 1975, computer circuits have doubled in circuit complexity every 2 years, which roughly translates to a doubling in compute power [18]. Even this doubling of compute power failed to reach adequate levels for some of the newer ML algorithms found with "deep learning" which require millions of complex linear algebra calculations. The relatively recent employment of graphical processing units (GPUs) made these algorithms' utilization feasible. GPUs are special computer chips initially developed for computer graphical tasks, such as video games. For ML tasks, they perform calculations orders of magnitude faster than traditional computer chips [19]. Companies, such as Google, have expanded upon this idea with their creation of ML-specific chips like tensor processing units which run with improved energy costs and speed [20]. This additional computer "horsepower" has allowed for actual implementations of all the algorithms that AI's inventors could heretofore only imagine.

The third key to the AI revolution is deep learning algorithms. AI's inception started with the theory that computer networks could mirror a human's own neural networks to create intelligence. Relatively simple tasks with totally knowable data (such as a Checkers game) were quickly implemented [10]. However, complex tasks that many would define as a marker of true intelligence, such as image and speech recognition, escaped AI designers. Krizhevsky et al. created the breakthrough with their application of deep convolutional neural networks. They realized that the natural resources for AI, big data, and compute power finally were bountiful. To tackle complex tasks like image recognition, however, required development of a neural architecture – like the human brain, complex enough to use a limited set of training data to extrapolate recognition to all permutations. Their use of deep convolutional neural networks halved the error rate for image recognition compared to all other competitors [21]. These deep learning algorithms have become the primary approach to create intelligence and cognition at levels that meet or exceed a human's capabilities, from image recognition (with computer vision [CV]) to language (natural language processing [NLP]) [22].

The fourth key to the AI revolution is the increased investment that has accompanied the previous foundation. No longer is AI trapped in the unfundable "AI Winter" where private and public sector funding disappeared [12]. The US government invested $1.1 billion towards AI in 2015 alone [23]. The private sector has seen a similar increase, with a doubling in AI private equity investment from 2016 to 2017. In fact, 12% of worldwide private equity investment went to the AI industry in 2018 alone [24]. Healthcare, in particular, is seeing an order-of-magnitude increase in funding, from $600 million in 2014 to a projected $6.6 billion by 2021 [25]. Big data, adequate compute power, deep learning algorithms, and increased investment have generated the rich AI landscape of today.

AI in Healthcare

The AI revolution has led to an explosion of healthcare-related applications. The foundations of AI in healthcare rest upon the deep learning algorithm's ability, through CV and NLP, to emulate humans' cognitive capabilities. In the surgical arena, it mainly has served to augment, rather than supplant, the human element. Successful AI utilization is found in all phases of the surgery, from preoperative diagnosis and risk assessment to intraoperative assistance and postoperative complication prediction.

The preoperative phase has seen the largest application of AI technology. AI algorithms can go head-to-head with physicians – particularly in the image-predominant fields of radiology and pathology for preoperative diagnosis. Some examples of applications include diagnosis of intracranial hemorrhage from CT images, breast cancer from mammography, and lung cancer from tissue slides [26]. One exceptional example comes from work in dermatology. Esteva et al. developed a diagnostic system based on convolutional neural networks capable of classifying dermatologic lesions as malignant or benign with superior sensitivity and specificity when compared to board-certified dermatologists [27]. AI has also helped with preoperative patient risk stratification. One example includes the POTTER score, an algorithm based on the ML technology of optimal classification trees that outperformed traditional multivariable logistic regression model surgical risk calculators, such as the ACS-NSQIP calculator [28].

The postoperative phase has also begun to see the introduction of AI technology. The majority of efforts have focused on complication prediction, as previous works have identified the concept of "failure to rescue," where overall complication rates between high- and low-performing hospitals are identical, but the lower-performing hospital have twice the mortality rates. These efforts hope to, through integration of myriad variables, detect complications early and therefore prevent a snowball effect that ultimately leads to higher mortality rates – which, in the case of pancreatic cancer, is over an order of magnitude higher [29, 30]. For example, one model considers over 175,000 data points per patient to predict mortality and morbidity [31]. Similar efforts aimed to predict postoperative surgical site infections from pre- and postoperative laboratory values [32].

Despite the development of AI technology for the pre- and postoperative phases, there has been relatively fewer applications to the intraoperative phase. A few computer vision groups focus on *temporal segmentation* of laparoscopic intraoperative videos with analysis of cholecystectomies, sleeve gastrectomies, and colectomies

[33–35]. One other group has worked to link intraoperative performance metrics from robotic surgery to predict postoperative events. For example, investigators could predict, based on intraoperative metrics alone, if a patient's length of stay postoperatively would exceed 2 days [36]. The past 7 years have generated significant progress for AI in healthcare, but there continues to remain numerous unexplored avenues for further applications.

Future Applications for AI in Surgery

The previously described AI innovations in healthcare are technologically revolutionary. The seemingly impossible tasks of image, speech, and language classification are now obtainable, at least at a rudimentary level. However, with respect to patient care, the advancements hardly seem to warrant the label of a "Cognitive Revolution." Fortunately, with the ever-increasing amount of generated data, more powerful computers, improved algorithms, and influx of funds, AI in healthcare is primed for a revolution.

This revolution will progress in incremental steps. Decision-support systems will become more pervasive at every stage of a patient's care. Consider a patient referred with a diagnosis of colon cancer. In the next few years, the patient's initial visit will seem relatively similar to the one from today, but it will incorporate ML algorithms throughout to augment their care. For example, an algorithm will classify their tumor at a granularity far superior to our currently crude TNM staging system to create an individualized treatment plan. Additionally, their metrics, including history, vitals, lab values, and imaging, will combine to form a comprehensive risk assessment. Initially, the risk assessment will help determine surgical readiness. However, in the coming years, it will evolve so that it is capable of providing recommendations for appropriate "pre-habilitation" to optimize the patient for surgery.

Intraoperative decision support will also start to slowly pervade the operating room. It will likely start with simple guidance, for example, to optimize laparoscopic port placement or help correlate preoperative imaging (such as tumor and major vasculature locations) with intraoperative displays. Work with *temporal-phase segmentation* will continue to build and begin to provide true operative guidance. Early implementations may offer a simple traffic light system, with a "green light" when dissection is going well, a "yellow light" when the surgeon is off course from a typical operation, and a "red light" when they are about to injure a vital structure. It will also offer a "phone-a-friend" functionality to connect to consultants for assistance. With continued development, this technology will ultimately develop into an intraoperative "GPS," guiding a surgeon through an operation step-by-step.

Postoperative decision support will include *early warning systems* to flag surgeons that a patient may have a certain complication. In the near future, integration of postoperative patient metrics with intraoperative video findings may lead to enhanced prediction that will predict not only that a complication may occur, but exactly the complication that will occur. Since these technologies will incorporate data from across hospitals and even countries, *their fund of knowledge will far exceed that of any surgeon and create a unified "collective surgical consciousness" that will provide the optimal care*.

Outside of decision-support systems, *automation* with underlying AI technology will also start to be incorporated into the operating room. It will start with automation of small tasks. For example, after recommending laparoscopic port placements, the machine may be able to dock a robotic-assisted surgery platform independently. Other small tasks may include fascial closure or anastomoses. Already, the Smart Tissue Autonomous Robot (STAR) can perform linear suturing and even autonomous sutured bowel anastomoses. In fact, its anastomoses can resist double the leak pressure compared to a human-sutured bowel anastomosis [37, 38]. In the coming years, the surgeon will perform the majority of the dissection, prepare the bowel, and then just press a "bowel anastomosis" button for a picture-perfect anastomosis constructed with optimal tension and precise bite-size throughout. In the more distant future, these incremental autonomous steps will combine until perhaps fully autonomous surgery is realized.

Challenges

AI promises a Cognitive Revolution, but with this revolution will come numerous hurdles and challenges. Without careful treatment, AI's progress may again derail, as it did in the 1980s, for a second coming of the AI Winter.

Ethics

AI and ML technologies present multiple ethical dilemmas. First is the issue of the "moral machine." The original "moral machine" problem asked a variety of questions to people across the world about autonomous driving scenarios, such as whether an autonomous vehicle should hit pedestrians to save the vehicle's passengers or swerve to avoid the pedestrian, thereby striking a barrier and killing the car's occupants. Answers depended on the scenario. For example, participants were more likely to favor saving the vehicle's occupants if they were younger than the pedestrian, or if the pedestrian was illegally crossing the street. *Interestingly, answers varied greatly across different world regions* [39]. Similar scenarios could arise as AI becomes pervasive in medicine. For example, will decision-support algorithms recommend against surgery for certain patients based upon their potential future societal contributions – or favor more aggressive treatment to wealthier patients? *AI model designers will need to provide algorithmic customization based on the locale's cultural norms and regularly work with communities to provide ethically acceptable decisions.*

A second ethical dilemma arises in the bias inherent to many AI algorithms. A recent analysis found one commercial prediction algorithm significantly under-triaged black patients compared to white patients due to use of healthcare costs as a surrogate for a patient's medical complexity. Since black patients had less access to more expensive treatment, their less-expensive care triaged them to an incorrect healthier risk strata [40]. *Training datasets need meticulous curation for fair representation of all patients; otherwise, algorithms unfairly trained may augment already present disparities* [41].

A third ethical dilemma comes from the training process for these models. ML model training is incredibly energy intensive, requiring powerful computers with hours to days of training time over multiple iterations before adequate model performance achievement. Since 2012, the amount of compute power used to train models has increased by 300,000-fold [42]. Training one model generates almost 80,000 pounds of carbon dioxide, which surpasses double the amount an average American produces annually [43]. *Development of these models must be done in an ordered and thoughtful fashion to minimize environmental repercussions.*

Privacy

AI and ML require big data, but big data raises numerous privacy issues. "De-identified" data should be anonymous; however, true "de-identification" is difficult, if not impossible. One researcher could link over 40% of newspaper stories regarding hospital admissions to "anonymized" public databases for hospital stays in the state of Washington [44]. In fact, based off gender, postal zip code, and date of birth (common information in "de-identified" datasets), 87% of United States' citizens are uniquely identified

[45]. Beyond issues with anonymity, many ML algorithms can make inferences about a patient to fill in missing data. As an example, some can infer smoking status (even if it is unknown) to help predict lung cancer risk. Future algorithms may be able to infer more sensitive information, such as HIV status, that the patient may not want known [46]. Other issues include ownership of data and patients' rights to withdraw their data and consent. Laws such as the European General Data Protection Regulation (GDPR) aim to protect the data rights of subjects. Concerted effort must be made to protect patients' privacy. One such solution may lie in split learning, where neural networks train across multiple data sources at different locations to prevent information leak from a central source [47].

Policy

To safely go forth and tackle the above issues, governmental and societal organizations must create sound policy. AI and ML algorithms continually "learn" and update, so regulatory agency approval and guarantee of safety may no longer apply after future training iterations. From a US perspective, the Food and Drug Administration (FDA) made a push in the early 2000s to classify software (smart phone application, stand-alone software, cloud-based solution) as a *medical device.* Congress then passed the 21st Century Cures Act as a response after software lobbying to remove many instances of software from the medical device list. Unfortunately, the Cures Act left a significant loophole for clinical decision support software, allowing it to be unregulated as long as it *intends* to explain to physicians its reasoning, even if this explanation is unsuccessful [48]. On the wings of this relative deregulation, the FDA has been approving increasing numbers of AI-related technologies and devices, from smart watches that can detect atrial fibrillation to algorithms that diagnose diabetic retinopathy [26]. Thus far, however, only "locked" algorithms (ones that will always return the same answer for a certain input) have been approved. The FDA recognized two main issues: first, the loophole and, second, the need for a framework that addressed evolving algorithms. As a response, they are working on a new regulatory framework [49]. As we march towards our AI future, concerted efforts, at the corporate, governmental, and societal level, must occur to ensure we proceed safely while still maximally benefiting from the new technology.

Annotations

The majority of the aforementioned algorithms represent *supervised learning,* where machines effectively learn from human-labeled examples. To teach a model surgical intraoperative phases, a surgeon will watch a video and label each phase, and then the algorithm will be given both the labels and the video and learn what constitutes each phase. *Learning requires immense amounts of data and a commensurate amount of labelling time.* Other areas solved this labelling problem through outsourcing, such as the "reCAPTCHA" tests seen online to ascertain whether a user is a human or computer. Ahn et al. used the reCAPTCHA tests to have regular Internet users transcribe over 440 million words from ancient texts with 99% accuracy [50]. Unfortunately for healthcare, the data is too complex for labelling by untrained annotators, so our annotation capabilities are severely limited by the relatively few expert annotators.

To solve the labor requirement, recent efforts have looked at streamlining the process. One group looked at pretraining models with unlabeled data to hopefully reduce the amount of required labelled data for accurate model training [51]. Others used a clever trick: they trained the model on a small number of videos and then used the model to auto-annotate further videos, achieving similar accuracy to models trained with four times the amount of data [52]. Future efforts will hopefully continue this "auto-annotation" process. The ML model's strength and ability to truly revolutionize surgery will require it to obtain superhuman knowledge and capabilities. Training with thousands to ultimately millions of videos from across the world will give it the collective experience of countless surgeons. A rare example seen in rural Canada could then prevent a complication the next day on the opposite side of the globe. *The sum of surgical experience will create a collective surgical consciousness greater than its individual parts.*

Surgical Training

The Cognitive Revolution that AI promises will require a different workforce in the future. Fewer physicians will be needed, particularly in fields that decision-support systems are well-suited to replacing (such as radiology and pathology). Automation will remove the need for physicians to do more quotidian tasks. Instead, the physician of the future will need more training in probability and statistical learning to accurately interpret algorithms that will assist their care. They will also need increased exposure to ethics to help morally apply these recommendations and computer science to understand the machinations providing them with daily assistance.

AI also promises to revolutionize surgical credentialing. The current process of regular written examinations fails to test actual surgical skills. With intraoperative ML models, surgeons in the future will be able to submit videos for recertification. If the video falls within a level of acceptable practice, they will then successfully recertify (of course, provided they also demonstrate aptitude in the management and care of the surgical patient). Similarly, when new procedures and technologies are introduced into surgical practice, the certification process will start first with intraoperative GPS guidance to train the surgeon, followed by automated video assessment to certify practice-ready performance.

Conclusion

We are at the start of the Cognitive Revolution driven by advancements in AI. The combination of big data, improved compute power, deep learning algorithms, and increased investment has led to an explosion in AI innovation and applications. Despite less than 10 years of innovation, AI models are matching, and often exceeding, human performance across all phases of patient care. This explosion brings with it numerous challenges, ranging from ethical dilemmas to privacy issues, which will require thoughtful and measured policies. The power of a collective surgical consciousness promises an exciting future and, more importantly, a safer future for patients.

References

1. Stevenson A, editor. Oxford dictionary of English. 3rd ed. New York: Oxford University Press; 2010.
2. Paul Thagard. Cognitive science. The Stanford Encyclopedia of Philosophy https://plato.stanford.edu/archives/spr2019/entries/cognitive-science/. Accessed 17 Dec 2019.
3. Tilney NL. Invasion of the body: revolutions in surgery. Cambridge: Harvard University Press; 2011.
4. Healey MA, Shackford SR, Osler TM, Rogers FB, Burns E. Complications in surgical patients. Arch Surg. 2002;137(5):611–8. https://doi.org/10.1001/archsurg.137.5.611.
5. Winston PH. Artificial intelligence. 3rd ed. Reading: Addison-Wesley Pub. Co; 1992.
6. Goertzel B, Pennachin C, editors. Artificial general intelligence. Berlin/New York: Springer; 2007.
7. Vinge V. The coming technological singularity: how to survive in the post-human era. 1993. http://hdl.handle.net/2060/19940022856.
8. McCulloch W, Pitts W. A logical calculus of the ideas immanent in nervous activity. Bull Math Biophys. 1943;5(4):115–33.
9. Russell SJ, Norvig P, Davis E. Artificial intelligence: a modern approach. 3rd ed. Upper Saddle River: Prentice Hall; 2010.
10. Samuel AL. Some studies in machine learning using the game of checkers. IBM J Res Dev. 1959;3(3):210–29.
11. Mccarthy J. Artificial intelligence: a paper symposium: professor sir James Lighthill, FRS. Artificial intelligence: a general survey. In: science research council, 1973. Artif Intell. 1974;5(3):317–22.
12. Crevier D. AI: the tumultuous history of the search for artificial intelligence. New York: Basic Books; 1993.
13. Simon HA, Munakata T. AI lessons. (Artificial intelligence; IBM's Deep Blue chess computer). Communicat ACM. 1997;40(8):23.
14. Schaeffer J, Burch N, Björnsson Y, et al. Checkers is solved. Science. 2007;317(5844):1518–22. https://www.jstor.org/stable/20037797. Accessed 29 Dec 2019.
15. Tromp J, Farnebäck G. Combinatorics of Go. In: van den Herik HJ, Ciancarini P, HHLM D, editors.

Computers and games. Lecture notes in computer science. Berlin, Heidelberg: Springer; 2007. p. 84–99. https://doi.org/10.1007/978-3-540-75538-8_8.

16. Sagiroglu S, Sinanc D. Big data: a review. In: 2013 International Conference on Collaboration Technologies and Systems (CTS); 2013. p. 42–7. https://doi.org/10.1109/CTS.2013.6567202.
17. Natarajan P, Frenzel JC, Smaltz DH. Demystifying big data and machine learning for healthcare. Boca Raton: CRC Press, Taylor & Francis Group; 2017.
18. Moore GE. Progress in Digital Integrated Electronics. In: International Electron Devices Meeting. Vol 21. Washington, DC; 1975:35–40.
19. Mittal S, Vaishay S. A survey of techniques for optimizing deep learning on GPUs. J Syst Archit. 2019;99.
20. Jouppi N, Young C, Patil N, Patterson D. Motivation for and evaluation of the first tensor processing unit. IEEE Micro. 2018;38(3):10–9. https://doi.org/10.1109/MM.2018.032271057.
21. Krizhevsky A, Sutskever I, Hinton GE. ImageNet Classification with Deep Convolutional Neural Networks. In: Pereira F, Burges CJC, Bottou L, Weinberger KQ, eds. Advances in Neural Information Processing Systems 25. Curran Associates, Inc.; 2012:1097–1105.
22. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521(7553):436–44. https://doi.org/10.1038/nature14539.
23. Preparing for the Future of Artificial Intelligence. Executive Office of the President; 2016.
24. OECD. Private equity investment in artificial intelligence. 2018. www.oecd.org/going-digital/ai/private-equity-investment-in-artificial-intelligence.pdf.
25. Collier M, Fu R, Yin L, Christiansen P. Artificial intelligence: healthcare's new nervous system. Viewable at https://www.accenture.com/_acnmedia/pdf-49/accenture-health-artificial-intelligence.pdf.
26. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019;25(1):44. https://doi.org/10.1038/s41591-018-0300-7.
27. Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature, London. 2017;542(7639):115–8G. http://dx.doi.org.ezp-prod1.hul.harvard.edu/10.1038/nature21056
28. Bertsimas D, Dunn J, Velmahos GC, Kaafarani HMA. Surgical risk is not linear: derivation and validation of a novel, user-friendly, and machine-learning-based Predictive OpTimal Trees in Emergency Surgery Risk (POTTER) calculator. Ann Surg. 2018;268(4):574–83. https://doi.org/10.1097/SLA.0000000000002956.
29. Silber JH, Williams SV, Krakauer H, Schwartz JS. Hospital and patient characteristics associated with death after surgery. A study of adverse occurrence and failure to rescue. Med Care. 1992;30(7):615–29.
30. Ghaferi AA, Birkmeyer JD, Dimick JB. Complications, failure to rescue, and mortality with major inpatient surgery in medicare patients. Ann Surg. 2009;250(6):1029–34. https://doi.org/10.1097/SLA.0b013e3181bef697.
31. Rajkomar A, Oren E, Chen K, et al. Scalable and accurate deep learning with electronic health records. NPJ Digit Med. 2018;1(1):18. https://doi.org/10.1038/s41746-018-0029-1.
32. Soguero-Ruiz C, Fei WME, Jenssen R, et al. Data-driven temporal prediction of surgical site infection. AMIA Ann Symp Proc 2015;2015:1164–1173. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4765613/. Accessed 18 Feb 2019.
33. Twinanda AP, Shehata S, Mutter D, Marescaux J, de Mathelin M, Padoy N. EndoNet: a deep architecture for recognition tasks on laparoscopic videos. IEEE Trans Med Imaging. 2017;36(1):86–97. https://doi.org/10.1109/TMI.2016.2593957.
34. Hashimoto DA, Rosman G, Witkowski ER, et al. Computer vision analysis of intraoperative video: automated recognition of operative steps in laparoscopic sleeve gastrectomy. Ann Surg. 2019;270(3):414–21. https://doi.org/10.1097/SLA.0000000000003460.
35. Kitaguchi D, Takeshita N, Matsuzaki H, et al. Real-time automatic surgical phase recognition in laparoscopic sigmoidectomy using the convolutional neural network-based deep learning approach. Surg Endosc. 2019; https://doi.org/10.1007/s00464-019-07281-0.
36. Hung AJ, Chen J, Gill IS. Automated performance metrics and machine learning algorithms to measure surgeon performance and anticipate clinical outcomes in robotic surgery. JAMA Surg. 2018;153(8):770. https://doi.org/10.1001/jamasurg.2018.1512.
37. Leonard S, Wu KL, Kim Y, Krieger A, Kim PCW. Smart Tissue Anastomosis Robot (STAR): a vision-guided robotics system for laparoscopic suturing. IEEE Trans Biomed Eng. 2014;61(4):1305–17. https://doi.org/10.1109/TBME.2014.2302385.
38. Shademan A, Decker RS, Opfermann JD, Leonard S, Krieger A, Kim PCW. Supervised autonomous robotic soft tissue surgery. Science Translational Medicine. 2016;8(337):337ra64. https://doi.org/10.1126/scitranslmed.aad9398.
39. Awad E, Dsouza S, Kim R, et al. The moral machine experiment. Nature. 2018;563(7729):59–64. https://doi.org/10.1038/s41586-018-0637-6.
40. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447–53. https://doi.org/10.1126/science.aax2342.
41. O'neil C. Weapons of math destruction: how big data increases inequality and threatens democracy. Broadway Books; 2016.
42. AI and Compute. OpenAI. 2018. https://openai.com/blog/ai-and-compute/. Accessed 30 Dec 2019.
43. Strubell E, Ganesh A, McCallum A. Energy and policy considerations for deep learning in NLP. arXiv:190602243 [cs]. 2019. http://arxiv.org/abs/1906.02243. Accessed 4 Sept 2019.

44. Sweeney L. Only you, your doctor, and many others may know. Technol Sci. 2015. /a/2015092903/. Accessed 12 Sept 2019.
45. Sweeney L. Simple demographics often identify people uniquely. Health (San Francisco). 2000;671:1–34.
46. Wiens J, Saria S, Sendak M, et al. Do no harm: a roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337–40. https://doi.org/10.1038/s41591-019-0548-6.
47. Gupta O, Raskar R. Distributed learning of deep neural network over multiple agents. arXiv:181006060 [cs, stat]. 2018. http://arxiv.org/abs/1810.06060. Accessed 17 Sept 2019.
48. Evans B, Ossorio P. The challenge of regulating clinical decision support software after 21st century cures. Am J Law Med. 2018;44(2–3):237–51. https://doi.org/10.1177/0098858818789418.
49. Food US, Administration D, others. Proposed Regulatory Framework for Modifications to Artificial Intelligence. Machine learning (AI/ML)-based Software as a Medical Device (SaMD) discussion paper and request for feedback. Available at: https://www.fda.gov/media/122535/download. Accessed 12 June 2019.
50. von Ahn L, Maurer B, Mcmillen C, Abraham D, Blum M. reCAPTCHA: human-based character recognition via web security measures. Science. 2008;321(5895):1465–8. https://doi.org/10.1126/science.1160379.
51. Funke I, Jenke A, Mees ST, Weitz J, Speidel S, Bodenstedt S. Temporal coherence-based self-supervised learning for laparoscopic workflow analysis. In: Stoyanov D, Taylor Z, Sarikaya D, et al., eds. OR 2.0 context-aware operating theaters, computer assisted robotic endoscopy, clinical image-based procedures, and skin image analysis. Lecture notes in computer science. Springer International Publishing; 2018. p. 85–93.
52. Yu T, Mutter D, Marescaux J, Padoy N. Learning from a tiny dataset of manual annotations: a teacher/student approach for surgical phase recognition. arXiv:181200033 [cs, stat]. 2018. http://arxiv.org/abs/1812.00033. Accessed 11 Sept 2019.