Artificial intelligence in radiation oncology

Elizabeth Huynh, Ahmed Hosny, Christian Guthier, Danielle S. Bitterman, Steven F. Petit, Daphne A. Haas-Kogan, Benjamin Kann, Hugo J. W. L. Aerts and Raymond H. Mak

Abstract | Artificial intelligence (AI) has the potential to fundamentally alter the way medicine is practised. AI platforms excel in recognizing complex patterns in medical data and provide a quantitative, rather than purely qualitative, assessment of clinical conditions. Accordingly, AI could have particularly transformative applications in radiation oncology given the multifaceted and highly technical nature of this field of medicine with a heavy reliance on digital data processing and computer software. Indeed, AI has the potential to improve the accuracy, precision, efficiency and overall quality of radiation therapy for patients with cancer. In this Perspective, we first provide a general description of AI methods, followed by a high-level overview of the radiation therapy workflow with discussion of the implications that AI is likely to have on each step of this process. Finally, we describe the challenges associated with the clinical development and implementation of AI platforms in radiation oncology and provide our perspective on how these platforms might change the roles of radiotherapy medical professionals.

Radiation therapy is a crucial pillar of cancer treatment and is indicated for ~50% of patients[1]. Estimates indicate, however, that millions of patients currently lack access to this vital treatment modality[2–6] owing to barriers such as a scarcity of infrastructure, technology and human resources (including treatment facilities, machines and planning systems as well as trained staff)[7]. Furthermore, radiation therapy has become increasingly complex over the past few decades owing to technological advances, resulting in a near-complete reliance on human–machine interactions including both software and hardware.

Despite technological advances, much of the radiation therapy workflow still requires time-consuming, manual input by a diverse team of health-care professionals, including radiation oncologists, medical physicists, medical dosimetrists and radiation therapists. The growing complexity of these human–machine interactions in conjunction with the increasing incidence of cancer has led to radiation oncology workforce shortages throughout the world and increasing variability in the quality of care[8]. Notably, variations in the radiotherapy treatment-planning process have been shown to negatively affect overall survival, even in clinical trials (a setting in which extra care is given to standardizing approaches)[9,10]. Furthermore, the radiation therapy knowledge and experience gap between adequately resourced and under-resourced health-care systems is one of the greatest global inequalities in cancer care and poses an enormous public health challenge.

Artificial intelligence (AI) involves the development and use of complex computer algorithms to perform tasks that normally require human intelligence, such as visual perception, pattern recognition, decision-making and problem solving, at a similar or improved level of performance. AI is transforming many fields of medicine and has the potential to address many of the challenges faced in radiation therapy and thereby improve the availability and quality of cancer care worldwide. Herein, we discuss the promise of AI to transform the field of radiation oncology by outlining each step of the clinical workflow and highlighting examples of how AI might increase the efficiency, accuracy and quality of radiation therapy, thus enhancing value-based cancer care delivery in today's resource-limited health-care environment. The possible applications of AI in radiation oncology are wide ranging and we have not covered them all in this article. Instead, we aim to provide an overview of the transformative potential of AI in radiation therapy and our perspective on the future of the radiation oncology workforce.

Artificial intelligence methods

Early AI platforms were predicated on rule-based reasoning performed by a computer system according to a set of steps and procedures defined by human experts[11,12]. However, the generalizability of these methods on variation of the input data and task scope is often limited by the lack of 'intelligent' components capable of processing 'edge cases' not explicitly described in the knowledge base[13]. These rule-based AI systems have achieved varying degrees of clinical utility[14]. Over the past decade, however, a fundamental shift has occurred in the algorithms powering the automation of image-based tasks. This shift has been marked by the revival of neural networks, a class of machine learning algorithms loosely based on our presumed understanding of how the human brain functions.

Research on neural networks has evolved from the mathematical development in the 1960s of the backpropagation algorithm, which is the main method of training neural networks and involves using the known output for each input value to fine tune the weights of a neural network, towards simple networks in the 1980s[15–17]. The increasingly large amounts of data available, along with increases in computational power and advances in algorithm development, have all revived interest in this field of research, leading to the development of 'deeper' neural networks with multiple intermediate hidden layers between the input and output layers (that is, the data fed into the network and the results generated, respectively). The function of hidden layers is to perform non-linear transformations of the input data to extract feature information in order to inform the output layer. The use of such algorithms has obviated the need to predefine reasoning rules because the map of 'hidden neurons' between input and output nodes can be learnt automatically from the training data. This approach provides deep learning algorithms with a greater learning capacity than that of preceding AI algorithms and consequently an ability to discern very complex, non-linear relationships in data. Deep learning can therefore begin to approximate or even surpass human capabilities for highly complex tasks and has been applied in several medical scenarios[18].

The radiation therapy workflow involves a multitude of complex tasks, including tumour and organ segmentation, dose optimization, outcome prediction and quality assurance (QA), which have seen varying degrees of digitization and consequent automation over the years. This heterogeneity is also reflected in the types of data used, ranging from radiographic images and radiation dose maps to hardware calibration log files and maintenance records. The multimodal nature of deep learning architectures[19] enables aggregation of these different data streams, cross-modality learning and algorithm generalizability, which might ultimately result in improved clinical decision-making and thus better quality care for all patients[20]. Indeed, various AI algorithms have been applied to each task of the radiation therapy workflow (TABLE 1).

Application in radiation oncology

The radiation therapy workflow can be divided into several steps including initial treatment decision-making, treatment planning and preparation, QA, delivery of radiation therapy and follow-up care (FIG. 1).

Table 1 | **A non-exhaustive list of the modern AI methods and their applications in radiation oncology**

AI method	Description	Selected applications in radiation oncology	Selected examples
XGBoost	A prediction modelling technique consisting of an ensemble of weaker prediction models, usually decision trees	Outcome prediction using structured data, such as tabular data on comorbidities, dosimetric indices, age, and so on, as well as radiomic features extracted from radiographic images[148]	Prediction of radiation-related fibrosis of neck muscles based on MRI data from patients with nasopharyngeal carcinoma[149]
		Artefact suppression in images; for example, in the context of motion management	Prediction of tumour motion ranges using 4D CT images in patients receiving radiotherapy for lung cancer[150]
Neural networks	Algorithms — loosely modelled on the neural networks of the human brain — comprising different layers, which are in turn composed of nodes that are activated based on input	Radiation dose quality assurance (QA)	Pretreatment dose verification in patients receiving radiotherapy for prostate cancer or nasopharyngeal carcinoma[151]
Convolutional neural networks (CNN)	Neural networks that are composed of convolutional layers (for perception), followed by fully connected layers (for cognition)	Outcome prediction from unstructured data; for example, derived from radiographic images	Prediction of rectal toxicities of radiotherapy for cervical cancer[152]
		Patient-specific QA measurements	QA of dose distribution in patients receiving radiotherapy for prostate cancer[153]
Fully convolutional neural networks (FCN)	Neural networks that are composed entirely of convolutional layers; images are encoded then decoded, thus producing a probability map per voxel indicating the probability of a specific prediction	Image segmentation using unstructured imaging data	Organ-at-risk segmentation in CT images of patients receiving radiotherapy for head and neck cancer[71]
		Prediction of radiation dose distribution	Prediction of the 3D dose distribution of stereotactic body radiation therapy (SBRT) in patients with prostate cancer[154]; prediction of dose distribution in patients receiving radiotherapy for nasopharyngeal carcinoma[155]
Variational auto-encoders (VAE)	Neural networks that perform dimensionality reduction on input data converting it into low-dimensional latent vectors	Outcome prediction from unstructured data; for example, radiographic images	Prediction of radiation pneumonitis in patients with non-small-cell lung cancer (NSCLC)[156]; prediction of intrahepatic failure of disease control and overall survival in patients who received SBRT for hepatocellular carcinoma[157]
Generative adversarial networks (GAN)	Neural networks comprised of 'generator' and 'discriminator' components that participate together in a zero-sum game; the generator attempts to generate synthetic samples that match the input data distribution, while the discriminator attempts to discern synthetic from real data	Generation of synthetic CT images	Generation of synthetic CT images using only MRI data to enable accurate calculation of radiation dose in the pelvis for patients with prostate, rectal or cervical cancer[35]
		Prediction of radiation dose distribution	Predicting optimal 3D radiation dose distributions for patients with oropharyngeal cancer[158]
Reinforcement learning (RL) with deep Q networks	RL involves training an agent to interact with its environment by performing 'actions' and arriving at 'states'; certain actions lead to 'rewards', which can be positive and negative	Radiation dose adaptation	Automated radiation adaptation protocols for patients with NSCLC[83]

AI, artificial intelligence.

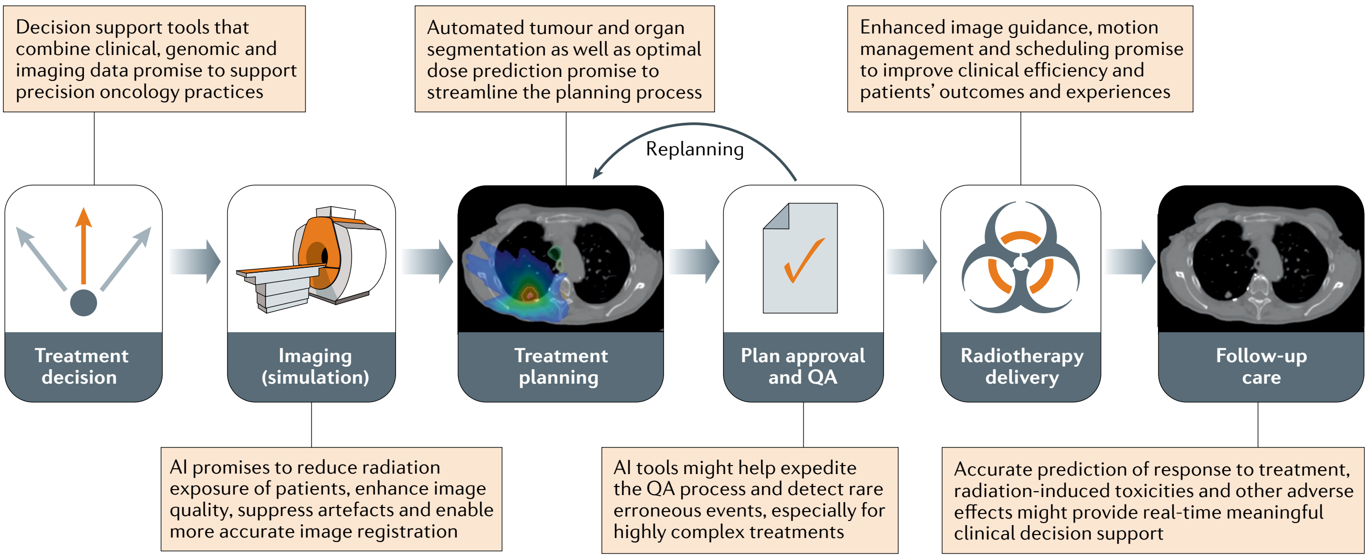

Fig. 1 | **Applications of AI in the radiation therapy workflow.** The image provides a general overview of the radiation therapy workflow with brief descriptions of expected applications of artificial intelligence (AI) at each step. The workflow begins with the decision to treat the patient with radiation therapy, followed by a simulation appointment during which medical images are acquired for treatment planning. Subsequently, the patient-specific treatment plan is created, and then the plan is subjected to approval, review and quality assurance (QA) measures prior to delivery of radiation to the patient. The patient then receives follow-up care. AI has the potential to improve radiation therapy for patients with cancer by increasing efficiency for the staff involved, improving the quality of treatments, and providing additional clinical information and predictions of treatment response to assist and improve clinical decision-making.

In the following sections, we described the key tasks at each step, the staff members involved and notable examples of the potential facilitatory roles of AI. For steps in the workflow where we do not anticipate an important role for AI (for example, the actual delivery of radiation) we have not provided examples.

Initial treatment decision-making

Patient evaluation. The clinical radiation therapy workflow starts with patient intake and evaluation. This step typically involves a consultation by the radiation oncologist that includes a review of the patient's symptoms, medical history, physical examination, pathological and genomic data, diagnostic studies, prognostication, comorbidities and risk of toxicities from radiotherapy; the radiation oncologist subsequently recommends a treatment plan based on a synthesis of these data (FIG. 2). An emerging challenge for clinicians involved in this process relates to the continuing accumulation of data to orders of magnitude beyond that which humans can rapidly absorb and interpret. AI-based methods that can automatically extract key clinically actionable features will be crucial to building decision support tools for clinicians at the initial point of care. AI approaches for medical imaging assessments[21] and natural language processing for electronic medical records[22,23] have shown initial promise in guiding treatment selection and/or the clinical management of patients with cancer. For example, prediction of the pathological response of involved lymph nodes in patients with non-small-cell lung cancer treated with chemoradiotherapy might inform the clinical decision to continue such therapy or proceed to surgery[21]. Moreover, such AI-based models have been reported to improve prognostication[24,25] and predict treatment outcomes[23,26–28], but have not yet been implemented in routine clinical practice.

Dose prescription. The prescribed dose of radiation to the tumour and dose constraints to the surrounding organs are determined by the radiation oncologist prior to treatment planning (FIG. 2), according to nationally accepted standards and evidence from clinical trials. However, variations in tumour biology can result in substantial differences in radiation sensitivity, even for a given cancer type. Furthermore, depending on the geometrical arrangement of the tumour and surrounding organs, the desired dose might not be achievable, which is often not realized until the planning process is near completion. AI platforms might enable the personalization of radiotherapy by predicting the radiation sensitivity of the tumour[29] and the optimal dose prescription that is achievable with a specific treatment plan, based on the contours of the tumour and organs[30].

Treatment planning and preparation

Treatment simulation — image acquisition, processing and registration. In preparation for treatment planning, simulation appointments take place during which the patient is immobilized to prevent substantial motion and, in most cases, medical images are acquired for use in formulating the treatment plan. Depending on the disease site, this process can be very complex, and optimal patient immobilization is subjective, and thus this process often requires radiation oncologist and medical physicist involvement (FIG. 2). For example, special consideration must be taken to evaluate potential interference between areas of the immobilization device and treatment beam angles or patient-specific issues that might result in collision with the treatment machine. Similar to how AI has been used to expedite treatment planning based on a patient's anatomy[30–32], we speculate that AI could have a role in identifying challenges that might be encountered at treatment simulation based on prior knowledge of the patient's anatomy (obtained, for example, through diagnostic imaging) and could offer solutions derived from the algorithm training data, thus expediting and optimizing the planning process.

For many patients scheduled to receive radiation therapy, multiple types of medical images are required for treatment planning, including CT images for calculating the

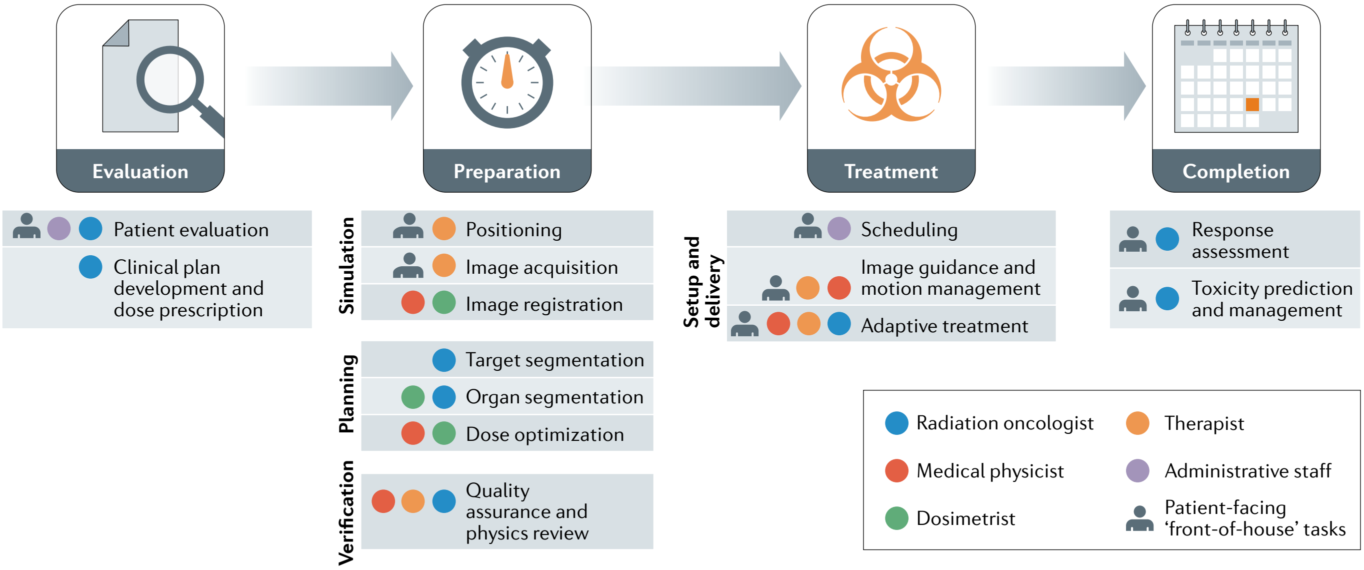

Fig. 2 | **Staff involvement and patient-facing steps in the radiation therapy workflow.** The radiation therapy workflow can be broken down into four main stages: patient evaluation and development of the clinical plan, preparation of the treatment (including quality assurance procedures), treatment setup and delivery, and completion. Each of these stages involves various steps and members of the radiation therapy team, such as radiation oncologists, medical physicists, dosimetrists, therapists and administrative staff, as indicated in the image. The staff members involved in patient-facing 'front-of-house' tasks are also indicated.

radiation dose and MRI scans for tumour segmentation. Typically, these images are acquired with the patient in different positions (in the treatment position during CT, but in various other positions during diagnostic imaging using other modalities), which introduces uncertainty when aligning the images. Eliminating the need for CT by acquiring MRI data that can also provide electron density information (that is, a synthetic CT) is one method of minimizing this uncertainty. AI has been used to generate synthetic CT images from MRI images of the brain[33,34] and pelvis[35], with minimal dose differences observed between the treatment plans formulated using synthetic CT versus true CT[33,35]. Additionally, this approach has the potential to improve clinical efficiency and costs by reducing the number of imaging appointments that patients need to attend, while also limiting their exposure to radiation from CT scans.

Advances in technology have led to the emerging roles of MRI in guiding radiation therapy and indeed the integration of MRI scanners with linear accelerators in a single treatment technology (MR Linac)[36–38]. High-resolution and low-noise MRI images require long acquisition times; thus, a compromise has to be made with regard to the resolution and signal-to-noise ratios that are achievable in the time available for image acquisition and other clinical tasks. AI has the potential to reduce MRI scan times by enabling reconstruction of fine details from undersampled MRI data, as demonstrated by the use of deep learning algorithms to generate high-resolution, high-contrast and low-noise brain[39–41] and cardiac MRI images[42] from undersampled data. Owing to the complexities of integrating MRI scanners with radiotherapy linear accelerators (that is, the distorting effects of the magnetic field on the radiation beams and the artefacts that the components of the linear accelerator can have on a magnetic field), current MR Linac systems are built with low-strength magnets, typically 0.35–1.5 T[43–45], which reduces image quality compared with the high-resolution images obtain using conventional high-field-strength MRI scanners. AI could enable the reconstruction of high-signal, high-resolution images from low-field-strength MRI scans (for example, 7-T MRI-like images of the brain from 3-T MRI data)[46] to improve the visualization of tumours throughout the course of treatment.

Image registration is another integral part of the radiation therapy workflow in which data from multimodality and longitudinal imaging are used not only during treatment planning, but also immediately prior to delivery of each treatment fraction, as well as for real-time monitoring of radiation delivery. Commercially available automatic image-registration algorithms are typically designed to perform well only for modality-specific registration problems and are sensitive to image artefacts, which compromises accuracy and often requires additional manual edits to achieve a clinically acceptable registration. AI tools have been trained to determine the sequence of motion actions that result in optimal image alignment; these algorithms can achieve better accuracy and robustness than several state-of-the-art registration methods[47] and are generalizable across multiple imaging modalities[47,48]. Furthermore, AI approaches have been shown to mitigate the effects of image artefacts (for example, with X-ray images of the spine that contained artefacts resulting from metal screws and guide wires)[49] and motion artefacts (such as those commonly encountered with fetal MRI)[50] on registration accuracy. AI tools have been developed for initial applications in MRI[51], X-ray[49,52,53], CT–MRI[54] and MRI–PET[55] image registration. Although many of these algorithms have not been developed specifically in the context of radiation therapy, the challenges they address are also faced in this context; therefore, these algorithms could potentially be applied to improve the radiation therapy workflow.

Image segmentation and dosimetric treatment planning. Currently, manual segmentation of the primary tumour and affected lymph nodes is one of the most time-consuming but crucial tasks performed by the radiation oncologist (FIG. 2). The accuracy of tumour segmentation can directly affect outcomes: an incorrectly delineated tumour can lead to underdosing or overdosing, resulting in a decrease in the likelihood of tumour control or an increased risk of toxicities, respectively. Tumour segmentation is subject to inter-observer variation, even among expert radiation oncologists[56,57], which can lead to differences in the quality of treatment plans, with consequent effects on survival outcomes[9,10,58,59]. Current semi-automated segmentation tools that incorporate prior knowledge from reference images, such as segmentation atlases, are unreliable or inaccessible to many radiation oncologists owing to high costs and still require substantial manual input[60,61]. AI has the potential to dramatically increase the efficiency, reproducibility and quality of radiation treatment planning by enabling almost completely or possibly even fully automated segmentation approaches, such as those developed for contouring of nasopharyngeal carcinomas[62], primary lung tumours[63] and oropharyngeal carcinomas[64]. Importantly, the performance of these segmentation algorithms is similar to that achieved by human experts[62–64]. Nevertheless, further studies, particularly prospective studies, are required to directly compare the efficiency, accuracy and reproducibility of such AI tools against the current gold-standard approaches within the radiation therapy clinical workflow.

During radiation treatment planning, organs adjacent to the tumour are also segmented in order to calculate the radiation dose delivered to those crucial organs and ensure that it falls within safe limits. Early AI tools have demonstrated promise in delineating a variety of organs throughout the body, including the complex anatomy of the head and neck region[65], thoracic organs[66], kidneys[67], liver[68,69] and cardiac substructures[70]; however, these findings are limited by small training sets and thus potential overfitting of the AI algorithms. The largest scale example of this approach reported to date involved an academic–industry partnership between the University College London Hospitals Department of Radiotherapy and Google DeepMind, in which a training dataset of CT images from 663 patients was used to develop an algorithm capable of segmenting organs in the head and neck region with performance comparable to that of human experts[71]. With commercially available AI-based auto-segmentation tools now starting to feature in treatment planning systems, additional tools are required for QA to identify errors. QA of auto-segmentations is a labour-intensive and time-consuming task, and is in turn another area in which AI-based QA tools could potentially reduce the required time and resources[72].

Once provided with medical images, tumour and organ segmentations and the dose prescription, the medical dosimetrist aims to generate the optimal treatment plan for the patient, with the goal of maximizing the dose delivered to the tumour while sparing surrounding organs (FIG. 2). Treatment planning is a time-intensive, iterative process whereby the dosimetrist designs the dose distribution, making necessary changes on a trial-and-error basis in order to achieve the goals outlined in the dose prescription. The treatment plan is then evaluated by the radiation oncologist before approval for implementation. The quality of radiation treatment plans is dependent on several different human factors, such as the choice of radiation beam angles and optimization parameters for the plan, resulting in large variations both intra-institutionally and inter-institutionally[73].

Current strategies to standardize and improve the efficiency of dosimetric treatment planning are not AI-based and involve the automation of repetitive tasks using hard-coded rules and/or the optimization of plan parameters according to predefined objectives using statistical methods[74–78]. The methods used are typically designed for specific anatomical sites and have a limited capacity to account for variations in plan complexity and patient-specific trade-offs.

AI tools for automating treatment planning have two main steps: 1) predicting the optimal dose distribution; and 2) identifying the treatment machine parameters required to achieve that distribution. The results of several studies demonstrate the ability of deep learning algorithms to predict the optimal dose distributions for individual patients based on their anatomy[30–32] and to accelerate dose calculations[79]. In order for AI-based treatment-planning algorithms to generate a high-quality plan, information regarding the complex decision-making process needs to be included in the underlying model, similar to the approach used in the development of AI algorithms that are able to play Atari games[80] or the board game Go[81]. In retrospective studies, researchers have applied these gamification concepts to automatically generate treatment plans for high-dose-rate brachytherapy in patients with cervical cancer[82] or for radiation dose adaptation in patients with non-small-cell lung cancer[83], with comparable or superior performance to that of human planners. Overall, AI techniques have the potential to substantially improve this crucial step in the radiation workflow, first by predicting what radiation dose distributions can be safely achieved in order that radiation oncologists can select the optimal treatment approach and, second, by subsequently generating the treatment plan for delivery of the optimal radiation dose. Thus, AI might enable full automation of the treatment-planning process in the near future.

Pretreatment review and verification

After the radiation oncologist approves the treatment plan, the medical physicist performs plan checks and other QA checks to ensure that all the technical components involved in treatment delivery are functioning and set correctly to deliver the intended dose to the patient (FIG. 2). AI tools have been developed to minimize the need for repetitive, time-consuming manual measurements and improve the efficiency of some QA activities, such as patient-specific and machine QA assessments.

Patient-specific QA involves assessment of treatment plans to detect human errors and potential anomalies in the performance of the treatment machine software and hardware as a whole (as opposed to machine QA, in which isolated parts of the device are tested) that might affect delivery of the specific treatment plan. These assessments include checking the suitability and accuracy of the plan and treatment parameters, and verifying the planned dose against the delivered dose. AI tools have been shown to expedite this process and to detect rare errors. For example, for highly complex treatment plans, a physical measurement of the delivered dose is obtained using a dosimeter-containing phantom and compared to the planned dose. The majority of plans pass this QA step, but in the rare case that a plan fails, many potential contributing factors require investigation, which might delay treatment. An AI algorithm has been designed to predict QA passing rates based on the treatment plan itself and to identify the possible sources of errors, potentially eliminating the need for physical dose measurements[84,85].

Machine QA involves various assessments of treatment machine function, accuracy and precision that are conducted on a daily, weekly, monthly or annual basis. The plethora of data acquired during these evaluations has provided the means to develop AI algorithms that are capable of predicting trends and errors, such as multileaf collimator positional errors[86] and beam symmetry trends[87], and of automatically detecting imaging artefacts (such as scatter artefacts)[88]. These tools might improve the efficiency of the QA process and thus provide medical physicists with more time for other tasks.

Treatment setup and delivery

Scheduling. Patients receiving radiation therapy are required to attend the radiation oncology department for several appointments, including consultation, radiation dose planning, treatment and follow-up assessments, all of which can have varying durations and waiting times. Long waiting times negatively affect not only the efficiency of the clinic but also patient anxiety and satisfaction[89]. AI has the potential to identify the most important factors contributing to waiting time durations (such as the time of day, number of radiotherapy dose fractions, median past duration of treatments, number of treatment fields and previous treatment duration) and predict waiting times[90], thus enabling optimization of clinic flow and efficiency. Using AI models, appointment scheduling could potentially be further optimized by organizing the sequence of patients according to anatomical treatment site and the immobilization and treatment techniques used in order to decrease the room turnover time between patients and accommodate a higher number of patients.

Image guidance and motion management. Setting up the patient in the same position that was used to create the treatment plan is a key part of radiotherapy delivery. Currently, the integrated cone beam CT (CBCT) device of the treatment machine is most commonly used for 'on-treatment' imaging to position the patient; however, CBCT provides images of much lower quality than the planning CT images. AI has been applied to improve the image quality of CBCT in order to enable more accurate positioning of patients for treatment[91]. Increasingly complex and multimodality imaging techniques are being incorporated into image-guided radiation therapy, including on-board MRI, ultrasonography and optical surface imaging, which presents a unique opportunity for imaging-based AI methods to enhance and/or synthesize complex data at the point of care.

Patient or organ motion throughout treatment can necessitate increases in the radiation dose delivered to non-malignant tissues in order to ensure that the tumour volume is adequately irradiated. Motion-management methods aim to reduce, capture and/or monitor the extent of motion from respiration and/or digestion[92]. However, considerable variability in motion exists between and within individuals in terms of magnitude, amplitude and frequency as well as the movement of organs relative to each other, which complicates predictive modelling of tumour motion. AI can be used to account for these diverse variables by generating patient-specific dynamic motion-management models that adapt to changes in patterns of motion in order to improve tumour tracking. To date, research in this area has largely focused on the prediction of respiratory motion using data collected from external surrogate positional markers as inputs for the models[93–95]. These algorithms could automatically adjust for complex breathing patterns in real time to accurately track tumour motion and predict the position of the tumour up to 800 ms in advance[93].

Adaptive treatment. Substantial changes in a patient's anatomy between the planning appointment and delivery of treatment (typically days or weeks later) or throughout treatment (often over several weeks) can warrant re-planning. These changes often reflect tumour shrinkage or growth, or anatomical variations (such as movement of internal organs or differences in gas or liquid filling of the bowels and stomach) that could potentially result in altered doses to the tumour and organs. Adaptive treatments involve creating a new treatment plan based on up-to-date images of the patient's anatomy. Currently, the radiation oncologist must decide when anatomical changes are large enough to be clinically relevant based on their own qualitative assessment of the patient's clinical parameters and images. AI might provide tools to predict which patients require adaptation of treatment and the ideal time point at which it should occur. In retrospective studies, AI models predicted geometric changes occurring in patients with head and neck cancer throughout treatment and identified the fourth week as the ideal time point for treatment adaptation[96,97]. Similar approaches have been applied in patients with lung cancer to identify the need to adapt treatment plans in order to maximize local tumour control[98] and reduce radiation-induced pneumonitis[83].

Completion of treatment

Response assessment and follow-up care. The Response Evaluation Criteria in Solid Tumors[99] (RECIST) is the most widely adopted system for evaluating treatment response in patients with solid tumours based on their presence or absence and changes in their size. AI algorithms have the potential to provide more detailed information on tumour responses throughout the course of radiation therapy — for example, changes in tumour phenotype that are captured in imaging features and might provide better assessments of response and predictions of outcome than changes in size alone. Initial studies of the use of AI with pretreatment and post-treatment imaging for early assessment of response to various therapies have been conducted in patients with lung cancer and enabled prediction of cancer-specific outcomes, such as disease progression, development of distant metastasis and locoregional recurrence, as well as overall survival[29,100,101]. Similarly, AI has been used to predict treatment responses in patients with bladder[102] or pancreatic cancer[103,104].

The presence of radiation-induced tissue damage can not only reduce the reliability of RECIST definitions of response, but also obfuscate the detection of disease recurrence. Studies have shown that AI algorithms have the potential to detect early changes in the lung that are associated with local recurrence and might be overlooked by physicians as radiation-induced fibrosis[105]. This additional information would enable earlier, personalized treatment interventions to improve outcomes.

Toxicity prediction and management. The proactive, rather than reactive, management of acute and late toxicities in patients is complicated by the largely unpredictable occurrence and/or severity of such adverse effects. Nevertheless, predictive models of radiation toxicities can be generated based on imaging data and risk factors, including certain clinical characteristics, germline genomic variations and the radiation dose distribution, and can be used to guide treatment planning. To date, such approaches have focused mostly on subsets of these data sources and/or the extrapolation of radiobiological modelling from preclinical and observational studies[106]. AI is poised to enable these data streams to be analysed more comprehensively and thereby build more robust predictive models incorporating comorbidities, radiation dose and pretreatment imaging data[107], which could provide clinical decision support for both the anticipatory management and secondary prevention of toxicities. For example, AI-based probability models of non-malignant tissue complications have been developed to predict the severity of acute dysphagia[108], xerostomia[109] and oral mucositis[110] in patients with head and neck cancer. AI tools for predicting radiation-induced pneumonitis[111–113], oesophagitis[114], rectal toxicities[115] and epilepsy[116] in patients with other cancer types have also been developed.

Pretreatment clinical data can also be used to provide guidance on potentially severe toxicities. In a retrospective study, several AI algorithms trained on clinical data from electronic medical records accurately predicted the risk of acute toxicities leading to emergency room visits and hospital admission in patients receiving radiotherapy or chemoradiotherapy (sensitivity of 81.0% and specificity of 67.3% with the best-performing gradient tree-boosting method)[23]. The integration of multiple clinical data streams for advanced forecasting of adverse events during radiation therapy is a representative example of the power of AI to provide real-time, meaningful clinical decision support at the point of care.

Development challenges

Multiple challenges lie ahead on the path to developing clinical AI tools, with the availability of high-quality datasets for algorithm training and validation of these algorithms arguably being the most crucial factor. The amount of data needed to construct high-accuracy AI models is strongly dependent on the application and the nature of the outcome data. The wealth of data generated for every patient often requires laborious curation before it can be utilized in developing AI models, especially given the lack of consistent standards in the generation of these data. Areas that suffer from limited standardization of definitions include organ and non-malignant tissue annotation and contouring[117], treatment techniques, the nature and timing of tumour recurrence, severity grading of toxicities, and the concepts and metrics used to evaluate treatment plans[118,119]. This deficiency hinders the sharing and aggregation of data across institutions, which is a prerequisite for developing AI models that accurately capture the full breadth of clinical variations while avoiding biases towards local standards. Although the creation of medical data repositories, such as The Cancer Imaging Archive[120], has helped to promote data sharing, and professional organizations have attempted to standardize the radiation oncology ontology[121,122], more work is needed in this area.

The proprietary nature of the treatment-planning software packages, and thus the limited knowledge of their optimization algorithms, is another hurdle facing the development of AI methods for radiation therapy. This challenge is being alleviated as some vendors start to release application programming interfaces that enable research efforts to communicate with and integrate AI algorithms into clinical software, albeit with restricted scopes.

Early AI research in radiation oncology has typically been focused on easily measured outcomes, such as overall survival, which might not be the outcome of greatest interest for all patients treated with radiotherapy. Instead, AI solutions will begin to move towards the prediction of outcomes that are more directly pertinent to radiation therapy, such as local tumour control and radiation-induced toxicities; however, the collection of robust outcome data continues to be a challenge.

Challenges to clinical implementation

Clinical adoption is a key barrier to realizing the potential of AI in radiation oncology; the introduction of AI tools will require upfront investment of time and resources as well as efforts to understand the utility and limitations of these tools and to redesign the current clinical workflows. Many AI tools remain at the proof-of-concept stage and lack external validation[123], resulting in a slow translation into routine practice such that demonstration of generalizability and effectiveness becomes unattainable[124]. Establishing trust in AI systems is also crucial, given the ‘black box’ nature of many machine learning algorithms and specifically deep learning. Despite active research into the ‘interpretability’ and ‘explainability’ of AI[125] (that is, understanding what an algorithm is doing and the underlying mechanics, respectively), the lack of transparency of AI hinders our ability to understand the outputs, predict failures and troubleshoot generalizability issues. Without actively monitoring the performance of deployed AI tools as well as continuous assessment of training data fit to the problem at hand, errors might increase as systematic biases are introduced into these systems.

Current AI tools are not perfectly accurate, and three criteria can be used to evaluate their potential for clinical implementation: 1) the time available for and the ability of the user to judge the accuracy of the result; 2) whether erroneous results can be corrected; and 3) the consequence of errors for a patient. Even in the case of potentially severe consequences, clinical implementation can be fairly straightforward as long as errors by the model are detected and corrected before moving on to the next step in the radiotherapy workflow. The potential for clinical implementation will be lower, however, if the time and ability required for the user to judge the accuracy of the result outweighs the efficiency or accuracy gains of using the AI tool. Furthermore, the risk-to-benefit ratio of using the AI-based tool is much more challenging to determine for applications in which the user cannot judge the correctness of the result (for example, when a tumour is not visible on an image and an AI tool is used for auto-segmentation). Tasks assisted or completed by AI that could have a substantial effect on a patient’s treatment will present a particular challenge to clinical implementation owing to the potential consequence for the patient.

From a legal standpoint, a means of governing algorithm-based decision-making has yet to be fully developed, including the right of patients to be given an explanation for algorithm outputs as well as the implications of data protection laws[126,127]. AI has the potential to reduce medical errors, but is also expected to alter the legal landscape surrounding clinical liabilities and responsibilities[128]. Indeed, the increased utilization of AI will change the dynamics of the patient–doctor relationship, likely with a shift towards a patient–health-care system relationship, thus potentially eroding the notion of personal responsibility of the doctor for the patient. In terms of ethics, algorithms used for facial detection[129] or for predicting an offender’s risk of recidivism[130] have already demonstrated inherent racial biases, and applications of AI in health care are already starting to present similar problems[131]. Moreover, unethical AI approaches could potentially be developed by parties with ulterior motives to skew results towards financial gain[132]. All of these challenges must be addressed to enable the effective widespread clinical adoption of AI-based tools.

Regulation and clinical evaluation

Currently, AI technologies are classified as ‘software as a medical device’ by the FDA and international regulatory bodies[133]. Many of the applications of these technologies in radiation therapy will fall under these regulatory standards; for example, treatment-planning decision support software has been explicitly identified as software as a medical device[134,135]. Much discussion has focused on schedules for re-evaluation of new devices and on the regulatory requirements for locked versus continuously learning AI algorithms[136],

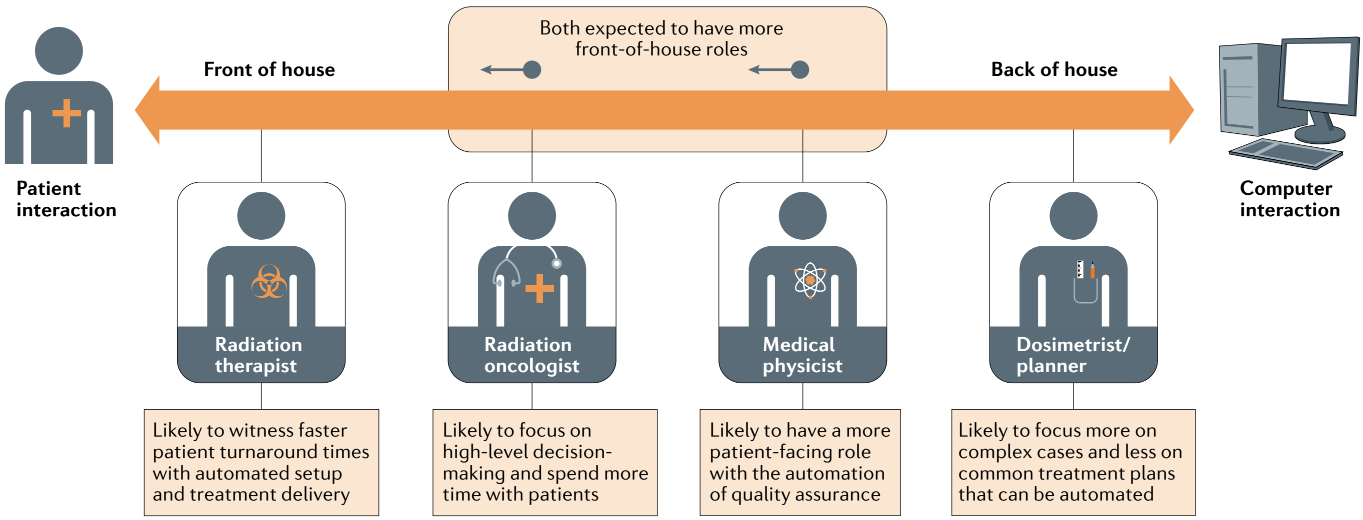

Fig. 3 | **Potential implications of applying AI in radiation oncology for members of the radiation therapy workforce.** Radiation therapists, radiation oncologists, medical physicists and dosimetrists are shown along a spectrum according to their overall level of involvement in patient-facing 'front-of-house' tasks versus predominantly computational 'back-of-house' tasks. Our projections regarding how each profession is expected to evolve with the integration of artificial intelligence (AI) tools into the radiation therapy workflow are summarized.

although clearer standards for clinical evaluation of the utility of these devices are needed. Notably, AI tools can have implications for patient outcomes that can only be identified through robust retrospective or prospective studies carried out in representative populations.

Whereas randomized clinical trials are the gold standard in the evaluation of anticancer therapies, such studies are neither feasible nor necessary for all AI tools. Nevertheless, tools that have the potential to affect patient outcomes, costs and the efficiency of the clinical workflow (as opposed to AI tools for machine QA, for example) should be considered for prospective clinical evaluation[137]. Given the rapid proliferation of AI technologies, master protocols evaluating multiple technologies of a single class across a range of malignancies might increase the feasibility and efficiency of prospective evaluation[137]. Phase I/II studies might be adequate for low-risk devices that will remain under ongoing surveillance by health-care providers, although phase III studies will be needed for high-risk tools that are used without standard clinical oversight. Postmarketing surveillance will be crucial to assess the value of AI-based radiation therapy technologies, considering that the function of these tools might be affected by interactions with other hardware and software. High-quality, risk-stratified clinical validation can establish the value of, and engender trust in, AI technologies, which will be particularly important for these black-box systems that can have considerable effects on cancer care.

AI and the radiotherapy workforce

As the shift towards the integration of AI into radiation oncology clinics unfolds over the next few decades, the roles of staff members will be redefined, especially those that currently spend substantial amounts of time on repetitive tasks requiring manual input. AI will predominantly affect staff members that perform 'back-of-house' activities, including the technical aspects of radiation therapy (such as tumour and organ segmentation, plan design and QA), with less of an effect on 'front-of-house' activities involving direct interaction with patients, which are typically carried out by physicians, radiation therapists and nurses (FIGS 2,3). In particular, nursing is a predominantly patient-facing profession, and thus the roles of nurses are unlikely to change substantially with the integration of AI into the clinic.

Implications for radiation oncologists

As AI-based segmentation algorithms begin to replace the manual segmentation performed by radiation oncologists, the focus of these physicians will shift to quality control of AI output and high-value, front-of-house activities of human interaction, such as patient counselling, education, support and clinical management (FIG. 3). Moreover, implementation of AI solutions will probably result in increased standardization of tumour segmentation and reduce unwarranted variation, particularly in under-resourced health-care environments, which might translate into improved clinical outcomes and quality of care.

Training of radiation oncologists will need to evolve from the current residency training models that focus on memorizing clinical facts and lengthy apprenticeships to gain expertise in performing manual segmentation and evaluating treatment plans. Instead, we predict that future training programmes will have an increased focus on instilling a deeper understanding of how to integrate and interpret information from large datasets in order to support clinical decision-making.

Implications for medical physicists

By analysing patterns and trends to predict when a technology needs to be serviced, AI tools have the potential to reduce the frequency and/or breadth of routine QA tasks performed by medical physicists. This change would cause a shift in the focus of medical physicists towards proactive prevention of non-routine, high-risk problems as well as the development and implementation of new technologies that require human creativity and intuition. As the field of radiation therapy moves towards more complex treatments, the role of the medical physicist will continue to be key to ensuring the accuracy, precision and clinical release of the technologies involved, including AI-based systems.

Additionally, thought leaders have called for the transitioning of medical physicists from back-of-house work to a more clinical, patient-facing role as a means of improving the quality of information provided to patients, as well as enhancing patients' experiences and satisfaction with their care[138–140]. If this transition is realized, and with appropriate training[139], our perspective is that the medical physicist's role will be further strengthened, despite automation of their technical tasks (FIG. 3).

Implications for medical dosimetrists

Medical dosimetrists currently perform many of the manual treatment-planning tasks that are most likely to be superseded by AI approaches. Studies have revealed that variation in the quality of treatment plans is generally attributable to the overall 'planner skill'[141], as opposed to other parameters such as experience, certification and education. This finding underlines the potential benefits of automating dosimetrists' tasks, especially the possibility of reducing the variability of delivered care. The potential for automation of treatment planning to reduce the workload of medical dosimetrists has been suggested to be dependent on the clinical accuracy of the plans generated[142]. Further evidence is required to provide sufficient confidence for a shift towards complete automation, yet data from early studies have demonstrated promising potential[82,83]. In the short term, we expect that the remit of dosimetrists will be focused on more high-risk and complex situations that present a challenge for current AI approaches (FIG. 3). We predict that automation with AI will probably disrupt this profession substantially in the long term. According to the 2017 American Association of Medical Dosimetry salary survey[143], 45% of respondents felt they were affected by understaffing. Automation could potentially reduce the dosimetrists' workload to reach appropriate staffing levels, although it might lead to substantial reductions in the number of dosimetrists.

Implications for radiation therapists

Radiation therapists serve as the final gatekeeper of treatment delivery to ensure patient safety and avoid misadministration of radiotherapy. As we have outlined, AI could provide software tools to help radiation therapists ensure accurate and safe treatment, as well as increase efficiency and patient access; however, we believe that the radiation therapists will continue to have an important role in being present to monitor the performance of these automated systems and the patient (FIG. 3).

Conclusions

Beyond gains in accuracy, reproducibility and consistency, partnering human intuition and the capacity of AI to leverage diverse information from large datasets has the potential to drastically improve efficiency and throughput in radiation therapy. These benefits have become of prime importance in the current era of cost reduction together with the shift from fee-for-service to value-based care[144].

The global health landscape also stands to benefit from AI-based interventions[145]. Over half of all patients with cancer live in low-income or middle-income countries[146]. Workforce and equipment shortages in these resource-constrained settings have left >50% of patients who are expected to benefit from radiotherapy without access to this treatment, with this value being up to 90% in some low-income countries[147]. Software AI applications promise to alleviate some of these shortages by providing specialized expert knowledge across disease sites and treatment modalities. Whether hardware equipment shortages can be addressed with AI remains unclear, although AI might help to support the upkeep of existing equipment by facilitating the analysis of machine QA reports[89].

Ultimately, the availability of AI tools will undoubtedly change the composition and skillset of the radiation oncology workforce; however, these changes will largely be positive and will enable the field to continue to bend the cost curve through greater efficiency while improving the quality of care.

Elizabeth Huynh[1,2], Ahmed Hosny[1,2], Christian Guthier[2], Danielle S. Bitterman[1,2,3], Steven F. Petit[4], Daphne A. Haas-Kogan[1,2], Benjamin Kann[1,2], Hugo J. W. L. Aerts[1,2,5,6]✉ and Raymond H. Mak[1,2]

[1]Artificial Intelligence in Medicine (AIM) Program, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA.

[2]Department of Radiation Oncology, Dana-Farber Cancer Institute, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA.

[3]Computational Health Informatics Program, Boston Children's Hospital, Harvard Medical School, Boston, MA, USA.

[4]Department of Radiation Oncology, Erasmus MC Cancer Institute, Rotterdam, The Netherlands.

[5]Department of Radiology, Dana-Farber Cancer Institute, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA.

[6]Department of Radiology and Nuclear Medicine, CARIM & GROW, Maastricht University, Maastricht, The Netherlands.

✉*e-mail: hugo_aerts@dfci.harvard.edu*

https://doi.org/10.1038/s41571-020-0417-8

Published online: 25 August 2020

1. Delaney, G., Jacob, S., Featherstone, C. & Barton, M. The role of radiotherapy in cancer treatment: estimating optimal utilization from a review of evidence-based clinical guidelines. *Cancer* **104**, 1129–1137 (2005).
2. Pan, H. Y. et al. Supply and demand for radiation oncology in the United States: updated projections for 2015 to 2025. *Int. J. Radiat. Oncol. Biol. Phys.* **96**, 493–500 (2016).
3. Ferlay, J. et al. Cancer incidence and mortality worldwide: sources, methods and major patterns in GLOBOCAN 2012. *Int. J. Cancer* **136**, E359–E386 (2015).
4. Miller, K. D. et al. Cancer treatment and survivorship statistics, 2016. *CA Cancer J. Clin.* **66**, 271–289 (2016).
5. Atun, R. et al. Expanding global access to radiotherapy. *Lancet Oncol.* **16**, 1153–1186 (2015).
6. Grover, S. et al. A systematic review of radiotherapy capacity in low- and middle-income countries. *Front. Oncol.* **4**, 380 (2015).
7. Elmore, S. N. C., Ben Prajogi, G., Rubio, J. A. P. & Zubizarreta, E. The global radiation oncology workforce in 2030: estimating physician training needs and proposing solutions to scale up capacity in low- and middle-income countries. *Adv. Radiat. Oncol.* 1–8 (2019).
8. Kresl, J. J. & Drummond, R. L. A historical perspective of the radiation oncology workforce and ongoing initiatives to impact recruitment and retention. *J. Am. Coll. Radiol.* **1**, 641–648 (2004).
9. Peters, L. J. et al. Critical impact of radiotherapy protocol compliance and quality in the treatment of advanced head and neck cancer: results from TROG 02.02. *J. Clin. Oncol.* **28**, 2996–3001 (2010).
10. Brade, A. M. et al. Radiation therapy quality assurance (RTQA) of concurrent chemoradiation therapy for locally advanced non-small cell lung cancer in the PROCLAIM phase 3 trial. *Int. J. Radiat. Oncol. Biol. Phys.* **101**, 927–934 (2018).
11. Kalet, I. J. & Paluszynski, W. Knowledge-based computer systems for radiotherapy planning. *Am. J. Clin. Oncol.* **13**, 344–351 (1990).
12. Laramore, G. E. et al. Applications of data bases and AI/expert systems in radiation therapy. *Am. J. Clin. Oncol.* **11**, 387–393 (1988).
13. Sanders, G. D. & Lyons, E. A. The potential use of expert systems to enable physicians to order more cost-effective diagnostic imaging examinations. *J. Digit. Imaging* **4**, 112–122 (1991).
14. Hosny, A., Parmar, C., Quackenbush, J., Schwartz, L. H. & Aerts, H. J. W. L. Artificial intelligence in radiology. *Nat. Rev. Cancer.* **18**, 500–510 (2018).
15. Dreyfus, S. The numerical solution of variational problems. *J. Math. Anal. Appl.* **5**, 30–45 (1962).
16. Fukushima, K. Neocognitron: a self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* **36**, 193–202 (1980).
17. LeCun, Y., Haffner, P., Bottou, L. & Bengio, Y. in *Shape, Contour and Grouping in Computer Vision* (eds Forsyth, D. A., Mundy, J. L., di Gesú, V. & Cipolla, R.) 319–345 (Springer, 1999).
18. Litjens, G. et al. A survey on deep learning in medical image analysis. *Med. Image Anal.* **42**, 60–88 (2017).
19. Ngiam, J. et al. Multimodal deep learning. *in Proceedings of the 28th international conference on machine learning (ICML-11)* 689–696 (2011).
20. Feng, M., Valdes, G., Dixit, N. & Solberg, T. D. Machine learning in radiation oncology: opportunities, requirements, and needs. *Front. Oncol.* **8**, 110 (2018).
21. Kann, B. H. et al. Pretreatment identification of head and neck cancer nodal metastasis and extranodal extension using deep learning neural networks. *Sci. Rep.* **8**, 14036 (2018).
22. Savova, G. K. et al. DeepPhe: a natural language processing system for extracting cancer phenotypes from clinical records. *Cancer Res.* **77**, e115–e118 (2017).
23. Hong, J. C., Niedzwiecki, D., Palta, M. & Tenenbaum, J. D. Predicting emergency visits and hospital admissions during radiation and chemoradiation: an internally validated pretreatment machine learning algorithm. *JCO Clin. Cancer Inform.* **2**, 1–11 (2018).
24. Oberije, C. et al. A validated prediction model for overall survival from stage III non-small cell lung cancer: toward survival prediction for individual patients. *Int. J. Radiat. Oncol. Biol. Phys.* **92**, 935–944 (2015).
25. Jochems, A. et al. Developing and validating a survival prediction model for NSCLC patients through distributed learning across 3 countries. *Int. J. Radiat. Oncol. Biol. Phys.* **99**, 344–352 (2017).
26. Deist, T. M. et al. Expert knowledge and data-driven Bayesian networks to predict post-RT dyspnea and 2-year survival. *Radiother. Oncol.* **118**, S29–S30 (2016).
27. Deist, T. M. et al. Machine learning algorithms for outcome prediction in (chemo)radiotherapy:

an empirical comparison of classifiers. *Med. Phys.* **45**, 3449–3459 (2018).
28. Gilmer, V., Timothy, D. S., Marina, H., Lyle, U. & Charles, B. S. II. Using machine learning to predict radiation pneumonitis in patients with stage I non-small cell lung cancer treated with stereotactic body radiation therapy. *Phys. Med. Biol.* **61**, 6105 (2016).
29. Lou, B. et al. An image-based deep learning framework for individualising radiotherapy dose: a retrospective analysis of outcome prediction. *Lancet Digital Health* **1**, e136–e147 (2019).
30. Nguyen, D. et al. A feasibility study for predicting optimal radiation therapy dose distributions of prostate cancer patients from patient anatomy using deep learning. *Sci. Rep.* **9**, 1076 (2019).
31. Campbell, W. G. et al. Neural network dose models for knowledge-based planning in pancreatic SBRT. *Med. Phys.* **44**, 6148–6158 (2017).
32. Häring, M., Groβhans, J., Wolf, F. & Eule, S. Automated segmentation of epithelial tissue using cycle-consistent generative adversarial networks. Preprint at *bioRxiv* (2018).
33. Dinkla, A. M. et al. MR-only brain radiation therapy: dosimetric evaluation of synthetic CTs generated by a dilated convolutional neural network. *Int. J. Radiat. Oncol. Biol. Phys.* **102**, 801–812 (2018).
34. Han, X. MR-based synthetic CT generation using a deep convolutional neural network method. *Med. Phys.* **44**, 1408–1419 (2017).
35. Maspero, M. et al. Dose evaluation of fast synthetic-CT generation using a generative adversarial network for general pelvis MR-only radiotherapy. *Phys. Med. Biol.* **63**, 185001 (2018).
36. Chandarana, H., Wang, H., Tijssen, R. H. N. & Das, I. J. Emerging role of MRI in radiation therapy. *J. Magn. Reson. Imaging* **48**, 1468–1478 (2018).
37. Rai, R. et al. The integration of MRI in radiation therapy: collaboration of radiographers and radiation therapists. *J. Med. Radiat. Sci.* **64**, 61–68 (2017).
38. Kerkmeijer, L. G. W. et al. The MRI-linear accelerator consortium: evidence-based clinical introduction of an innovation in radiation oncology connecting researchers, methodology, data collection, quality assurance, and technical development. *Front. Oncol.* **6**, 215 (2016).
39. Hyun, C. M., Kim, H. P., Lee, S. M., Lee, S. & Seo, J. K. Deep learning for undersampled MRI reconstruction. *Phys. Med. Biol.* **63**, 135007 (2018).
40. Wang, S. et al. Accelerating magnetic resonance imaging via deep learning. *Proc. IEEE Int. Symp. Biomed. Imaging* 514–517 (2016).
41. Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R. & Rosen, M. S. Image reconstruction by domain-transform manifold learning. *Nature* **555**, 487–492 (2018).
42. Schlemper, J., Caballero, J., Hajnal, J. V., Price, A. N. & Rueckert, D. A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Trans. Med. Imaging* **37**, 491–503 (2018).
43. Fallone, B. G. The rotating biplanar linac–magnetic resonance imaging system. *Semin. Radiat. Oncol.* **24**, 200–202 (2014).
44. Mutic, S. & Dempsey, J. F. The ViewRay system: magnetic resonance-guided and controlled radiotherapy. *Semin. Radiat. Oncol.* **24**, 196–199 (2014).
45. Raaymakers, B. W. et al. Integrating a 1.5 T MRI scanner with a 6 MV accelerator: proof of concept. *Phys. Med. Biol.* **54**, N229–N237 (2009).
46. Bahrami, K., Shi, F., Rekik, I. & Shen, D. in *Deep Learning and Data Labeling for Medical Applications* 39–47 (Springer, 2016).
47. de Tournemire P. et al. An artificial agent for robust image registration. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* 4168–4175 (2017).
48. Wu, G., Kim, M., Wang, Q., Munsell, B. C. & Shen, D. Scalable high-performance image registration framework by unsupervised deep feature representations learning. *IEEE Trans. Biomed. Eng.* **63**, 1505–1516 (2016).
49. Miao, S., et al. Dilated FCN for multi-agent 2D/3D medical image registration. *Thirty-Second AAAI Conference on Artificial Intelligence* (2018).
50. Hou, B. et al. Predicting slice-to-volume transformation in presence of arbitrary subject motion. *International Conference on Medical Image Computing and Computer-Assisted Intervention* 296–304 (2017).
51. Yang, X., Kwitt, R., Styner, M. & Niethammer, M. Fast predictive multimodal image registration. *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI)* 858–862 (2017).
52. Miao, S., Jane Wang, Z., Zheng, Y. & Liao, R. Real-time 2D/3D registration via CNN regression. *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*1430–1434 (2016).
53. Kearney, V., Haaf, S., Sudhyadhom, A., Valdes, G. & Solberg, T. D. An unsupervised convolutional neural network-based algorithm for deformable image registration. *Phys. Med. Biol.* **63**, 185017 (2018).
54. Ma, K. et al. Multimodal image registration with deep context reinforcement learning. in *Lecture Notes in Computer Science* 240–248 (2017).
55. Suk, H.-I., Lee, S.-W. & Shen, D. Hierarchical feature representation and multimodal fusion with deep learning for AD/MCI diagnosis. *Neuroimage* **101**, 569–582 (2014).
56. Van de Steene, J. et al. Definition of gross tumor volume in lung cancer: inter-observer variability. *Radiother. Oncol.* **62**, 37–49 (2002).
57. Cui, Y. et al. Contouring variations and the role of atlas in non-small cell lung cancer radiation therapy: analysis of a multi-institutional preclinical trial planning study. *Pract. Radiat. Oncol.* **5**, e67–e75 (2015).
58. Wuthrick, E. J. et al. Institutional clinical trial accrual volume and survival of patients with head and neck cancer. *J. Clin. Oncol.* **33**, 156–164 (2015).
59. Ohri, N. et al. Radiotherapy protocol deviations and clinical outcomes: a meta-analysis of Cooperative Group clinical trials. *J. Clin. Oncol.* **30**, 181–181 (2012).
60. Delpon, G. et al. Comparison of automated atlas-based segmentation software for postoperative prostate cancer radiotherapy. *Front. Oncol.* **6**, 178 (2016).
61. Kim, Y. et al. Impact of contouring accuracy on expected tumor control probability for head and neck cancer: semiautomated segmentation versus manual contouring. *Int. J. Radiat. Oncol. Biol. Phys.* **96**, E545 (2016).
62. Men, K. et al. Deep deconvolutional neural network for target segmentation of nasopharyngeal cancer in planning computed tomography images. *Front. Oncol.* **7**, 315 (2017).
63. Mak, R. H. et al. Use of crowd innovation to develop an artificial intelligence-based solution for radiation therapy targeting. *JAMA Oncol.* **5**, 654 (2019).
64. Cardenas, C. E. et al. Deep learning algorithm for auto-delineation of high-risk oropharyngeal clinical target volumes with built-in dice similarity coefficient parameter optimization function. *Int. J. Radiat. Oncol. Biol. Phys.* **101**, 468–478 (2018).
65. Ibragimov, B. & Xing, L. Segmentation of organs-at-risks in head and neck CT images using convolutional neural networks. *Med. Phys.* **44**, 547–557 (2017).
66. Lustberg, T. et al. Clinical evaluation of atlas and deep learning based automatic contouring for lung cancer. *Radiother. Oncol.* **126**, 312–317 (2018).
67. Jackson, P. et al. Deep learning renal segmentation for fully automated radiation dose estimation in unsealed source therapy. *Front. Oncol.* **8**, 215 (2018).
68. Peijun, H., Fa, W., Jialin, P., Ping, L. & Dexing, K. Automatic 3D liver segmentation based on deep learning and globally optimized surface evolution. *Phys. Med. Biol.* **61**, 8676 (2016).
69. Ibragimov, B., Toesca, D., Chang, D., Koong, A. & Xing, L. Combining deep learning with anatomical analysis for segmentation of the portal vein for liver SBRT planning. *Phys. Med. Biol.* **62**, 8943–8958 (2017).
70. Morris, E. D. et al. Cardiac substructure segmentation with deep learning for improved cardiac sparing. *Med. Phys*. **47**, 576–586 (2020).
71. Nikolov, S. et al. Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy. Preprint at *arXiv* (2018).
72. Zhang, J., Ates, O. & Li, A. Implementation of a machine learning-based automatic contour quality assurance tool for online adaptive radiation therapy of prostate cancer. *Int. J. Radiat. Oncol. Biol. Phys.* **96**, E668 (2016).
73. Berry, S. L., Boczkowski, A., Ma, R., Mechalakos, J. & Hunt, M. Interobserver variability in radiation therapy plan output: results of a single-institution study. *Pract. Radiat. Oncol.* **6**, 442–449 (2016).
74. Appenzoller, L. M., Michalski, J. M., Thorstad, W. L., Mutic, S. & Moore, K. L. Predicting dose-volume histograms for organs-at-risk in IMRT planning. *Med. Phys.* **39**, 7446–7461 (2012).
75. Babier, A., Boutilier, J. J., McNiven, A. L. & Chan, T. C. Y. Knowledge-based automated planning for oropharyngeal cancer. *Med. Phys.* **45**, 2875–2883 (2018).
76. Boutilier, J. J., Lee, T., Craig, T., Sharpe, M. B. & Chan, T. C. Y. Models for predicting objective function weights in prostate cancer IMRT. *Med. Phys.* **42**, 1586–1595 (2015).
77. Voet, P. W. J. et al. Toward fully automated multicriterial plan generation: a prospective clinical study. *Int. J. Radiat. Oncol. Biol. Phys.* **85**, 866–872 (2013).
78. Hussein, M., Heijmen, B. J. M., Verellen, D. & Nisbet, A. Automation in intensity modulated radiotherapy treatment planning — a review of recent innovations. *Br. J. Radiol.* **91**, 20180270 (2018).
79. Xing, Y., Nguyen, D., Lu, W., Yang, M. & Jiang, S. Technical note: a feasibility study on deep learning-based radiotherapy dose calculation. *Med. Phys*. **47**, 753–758 (2019).
80. Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
81. Chen, J. X. The evolution of computing: AlphaGo. *Comput. Sci. Eng.* **18**, 4–7 (2016).
82. Shen, C. et al. Intelligent inverse treatment planning via deep reinforcement learning, a proof-of-principle study in high dose-rate brachytherapy for cervical cancer. *Phys. Med. Biol.* **64**, 115013 (2019).
83. Tseng, H.-H. et al. Deep reinforcement learning for automated radiation adaptation in lung cancer. *Med. Phys.* **44**, 6690–6705 (2017).
84. Valdes, G. et al. A mathematical framework for virtual IMRT QA using machine learning. *Med. Phys.* **43**, 4323 (2016).
85. Valdes, G. et al. IMRT QA using machine learning: a multi-institutional validation. *J. Appl. Clin. Med. Phys.* **18**, 279–284 (2017).
86. Carlson, J. N. K. et al. A machine learning approach to the accurate prediction of multi-leaf collimator positional errors. *Phys. Med. Biol.* **61**, 2514–2531 (2016).
87. Li, Q. & Chan, M. F. Predictive time-series modeling using artificial neural networks for Linac beam symmetry: an empirical study. *Ann. N. Y. Acad. Sci.* **1387**, 84–94 (2017).
88. Valdes, G. et al. Use of TrueBeam developer mode for imaging QA. *J. Appl. Clin. Med. Phys.* **16**, 322–333 (2015).
89. Paul, C. et al. Cancer patients' concerns regarding access to cancer care: perceived impact of waiting times along the diagnosis and treatment journey. *Eur. J. Cancer Care* **21**, 321–329 (2012).
90. Joseph, A., Hijal, T., Kildea, J., Hendren, L. & Herrera, D. in *Machine Learning and Applications (ICMLA), 2017 16th IEEE International Conference* 1024–1029 (McGill Univ. Health Centre, 2018).
91. Kida, S. et al. Cone beam computed tomography image quality improvement using a deep convolutional neural network. *Cureus* **10**, e2548 (2018).
92. Langen, K. M. & Jones, D. T. Organ motion and its management. *Int. J. Radiat. Oncol. Biol. Phys.* **50**, 265–278 (2001).
93. Isaksson, M., Jalden, J. & Murphy, M. J. On using an adaptive neural network to predict lung tumor motion during respiration for radiotherapy applications. *Med. Phys.* **32**, 3801–3809 (2005).
94. Kakar, M., Nyström, H., Aarup, L. R., Nøttrup, T. J. & Olsen, D. R. Respiratory motion prediction by using the adaptive neuro fuzzy inference system (ANFIS). *Phys. Med. Biol.* **50**, 4721–4728 (2005).
95. Murphy, M. J. & Pokhrel, D. Optimization of an adaptive neural network to predict breathing. *Med. Phys.* **36**, 40–47 (2009).
96. Guidi, G. et al. A support vector machine tool for adaptive tomotherapy treatments: prediction of head and neck patients criticalities. *Phys. Med.* **31**, 442–451 (2015).
97. Guidi, G. et al. A machine learning tool for re-planning and adaptive RT: a multicenter cohort investigation. *Phys. Med.* **32**, 1659–1666 (2016).
98. Varfalvy, N., Piron, O., Cyr, M. F., Dagnault, A. & Archambault, L. Classification of changes occurring in lung patient during radiotherapy using relative γ analysis and hidden Markov models. *Med. Phys.* **44**, 5043–5050 (2017).
99. Eisenhauer, E. A. et al. New response evaluation criteria in solid tumours: revised RECIST guideline (version 1.1). *Eur. J. Cancer* **45**, 228–247 (2009).
100. Xu, Y. et al. Deep learning predicts lung cancer treatment response from serial medical imaging. *Clin. Cancer Res.* **25**, 3266–3275 (2019).
101. Hosny, A. et al. Deep learning for lung cancer prognostication: a retrospective multi-cohort radiomics study. *PLoS Med.* **15**, e1002711 (2018).
102. Cha, K. H. et al. Bladder cancer treatment response assessment in CT using radiomics with deep-learning. *Sci. Rep.* **7**, 8738 (2017).

103. Chen, X. et al. Assessment of treatment response during chemoradiation therapy for pancreatic cancer based on quantitative radiomic analysis of daily CTs: an exploratory study. *PLoS ONE* **12**, e0178961 (2017).
104. Horvat, N. et al. MR imaging of rectal cancer: radiomics analysis to assess treatment response after neoadjuvant therapy. *Radiology* **287**, 833–843 (2018).
105. Mattonen, S. A. et al. Detection of local cancer recurrence after stereotactic ablative radiation therapy (SABR) for lung cancer: physician performance versus radiomic assessment. *Int. J. Radiat. Oncol. Biol. Phys.* **96**, S48 (2016).
106. Lambin, P. et al. Predicting outcomes in radiation oncology — multifactorial decision support systems. *Nat. Rev. Clin. Oncol.* **10**, 27–40 (2013).
107. Lee, S. et al. Machine learning on a genome-wide association study to predict late genitourinary toxicity after prostate radiation therapy. *Int. J. Radiat. Oncol. Biol. Phys.* **101**, 128–135 (2018).
108. Dean, J. et al. Incorporating spatial dose metrics in machine learning-based normal tissue complication probability (NTCP) models of severe acute dysphagia resulting from head and neck radiotherapy. *Clin. Transl. Radiat. Oncol.* **8**, 27–39 (2018).
109. Gabryś, H. S., Buettner, F., Sterzing, F., Hauswald, H. & Bangert, M. Design and selection of machine learning methods using radiomics and dosiomics for normal tissue complication probability modeling of xerostomia. *Front. Oncol.* **8**, 35 (2018).
110. Dean, J. A. et al. Normal tissue complication probability (NTCP) modelling using spatial dose metrics and machine learning methods for severe acute oral mucositis resulting from head and neck radiotherapy. *Radiother. Oncol.* **120**, 21–27 (2016).
111. Cunliffe, A. et al. Lung texture in serial thoracic computed tomography scans: correlation of radiomics-based features with radiation therapy dose and radiation pneumonitis development. *Int. J. Radiat. Oncol. Biol. Phys.* **91**, 1048–1056 (2015).
112. Chen, S., Zhou, S., Yin, F.-F., Marks, L. B. & Das, S. K. Investigation of the support vector machine algorithm to predict lung radiation-induced pneumonitis. *Med. Phys.* **34**, 3808–3814 (2007).
113. Moran, A., Daly, M. E., Yip, S. S. F. & Yamamoto, T. Radiomics-based assessment of radiation-induced lung injury after stereotactic body radiotherapy. *Clin. Lung Cancer* **18**, e425–e431 (2017).
114. Luna, J. M. et al. Novel use of machine learning for predicting radiation esophagitis in locally advanced stage II–III non-small cell lung cancer. *Int. J. Radiat. Oncol. Biol. Phys.* **99**, E476–E477 (2017).
115. Zhen, X. et al. Deep convolutional neural networks with transfer learning for rectum toxicity prediction in combined brachytherapy and external beam radiation therapy for cervical cancer. *Int. J. Radiat. Oncol. Biol. Phys.* **99**, S168 (2017).
116. Liu, Z. et al. Radiomics analysis allows for precise prediction of epilepsy in patients with low-grade gliomas. *NeuroImage Clin.* **19**, 271–278 (2018).
117. Wright, J. L. et al. Standardizing normal tissue contouring for radiation therapy treatment planning: an ASTRO consensus paper. *Pract. Radiat. Oncol.* **9**, 65–72 (2019).
118. Covington, E. L. et al. Improving treatment plan evaluation with automation. *J. Appl. Clin. Med. Phys.* **17**, 16–31 (2016).
119. Evans, S. B. et al. Standardizing dose prescriptions: an ASTRO white paper. *Pract. Radiat. Oncol.* **6**, e369–e381 (2016).
120. Clark, K. et al. The cancer imaging archive (TCIA): maintaining and operating a public information repository. *J. Digit. Imaging* **26**, 1045–1057 (2013).
121. Mayo, C. S. et al. American Association of Physicists in Medicine task group 263: standardizing nomenclatures in radiation oncology. *Int. J. Radiat. Oncol. Biol. Phys.* **100**, 1057–1066 (2018).
122. Hayman, J. A. et al. Minimum data elements for radiation oncology: an ASTRO consensus paper. *Pract. Radiat. Oncol*. **9**, 395–401 (2019).
123. Kim, D. W., Jang, H. Y., Kim, K. W., Shin, Y. & Park, S. H. Design characteristics of studies reporting the performance of artificial intelligence algorithms for diagnostic analysis of medical images: results from recently published papers. *Korean J. Radiol.* **20**, 405 (2019).
124. Allen, B. Jr et al. A road map for translational research on artificial intelligence in medical imaging: from the 2018 National Institutes of Health/RSNA/ACR/ the Academy workshop. *J. Am. Coll. Radiol.* **16**, 1179–1189 (2019).
125. Gilpin, L. H. et al. in *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)* 80–89 (2018).
126. Goodman, B. & Flaxman, S. European Union regulations on algorithmic decision-making and a ‘right to explanation’. *AI Mag.* **38**, 50–57 (2017).
127. Kaminski, M. E. The right to explanation, explained. *Berkeley Technol. Law J.* **34**, 1 (2019).
128. Harned, Z., Lungren, M. P. & Rajpurkar, P. Machine Vision, Medical AI, and Malpractice (JOLT, 2019).
129. Buolamwini, J. & Gebru, T. in *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* Vol. 81 (eds Friedler, S. A. & Wilson, C.) 77–91 (PMLR, 2018).
130. Angwin, J., Larson, J., Mattu, S. & Kirchner, L. Machine bias. *ProPublica* **23**, 2016 (2016).
131. Obermeyer, Z., Powers, B., Vogeli, C. & Mullainathan, S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* **366**, 447–453 (2019).
132. Char, D. S., Shah, N. H. & Magnus, D. Implementing machine learning in health care — addressing ethical challenges. *N. Engl. J. Med.* **378**, 981–983 (2018).
133. IMDRF. “Software as a Medical Device”: Possible Framework for Risk Categorization and Corresponding Considerations. http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-140918-samd-framework-risk-categorization-141013.pdf (2014).
134. FDA. 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. https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf (2019).
135. FDA. Draft Guidance for Industry and Food and Drug Administration Staff. https://www.fda.gov/media/109618/download (2019).
136. Hwang, T. J., Kesselheim, A. S. & Vokinger, K. N. Lifecycle regulation of artificial intelligence- and machine learning-based software devices in medicine. *JAMA* **322**, 2285–2286 (2019).
137. Bitterman, D. S. et al. Master protocol trial design for efficient and rational evaluation of novel therapeutic oncology devices. *J. Natl Cancer Inst*. **112**, 229–237 (2020).
138. Schuller, B. W., Hendrickson, K. R. G. & Rong, Y. Medical physicists should meet with patients as part of the initial consult. *J. Appl. Clin. Med. Phys.* **19**, 6–9 (2018).
139. Brown, D. W. et al. A program to train medical physicists for direct patient care responsibilities. *J. Appl. Clin. Med. Phys.* **19**, 332–335 (2018).
140. Atwood, T. F. et al. Establishing a new clinical role for medical physicists: a prospective phase II trial. *Int. J. Radiat. Oncol. Biol. Phys.* **102**, 635–641 (2018).
141. Nelms, B. E. et al. Variation in external beam treatment plan quality: an inter-institutional study of planners and planning systems. *Practical Radiat. Oncol.* **2**, 296–305 (2012).
142. Adams, R. D. The future of medical dosimetry. *Med. Dosim.* **40**, 159–165 (2015).
143. American Association of Medical Dosimetrists. *2017 Salary Survey of Currently Active Medical Dosimetrists*. (American Association of Medical Dosimetrists, 2018).
144. Center for Medicare & Medicaid Services. Radiation Oncology Model. https://innovation.cms.gov/initiatives/radiation-oncology-model (2019).
145. Hosny, A. & Hugo, J. W. Artificial intelligence for global health. *Science* **366**, 955–956 (2019).
146. Barton, M. B., Frommer, M. & Shafiq, J. Role of radiotherapy in cancer control in low-income and middle-income countries. *Lancet Oncol.* **7**, 584–595 (2006).
147. Zubizarreta, E. H., Fidarova, E., Healy, B. & Rosenblatt, E. Need for radiotherapy in low and middle income countries — the silent crisis continues. *Clin. Oncol.* **27**, 107–114 (2015).
148. Aerts, H. J. W. L. et al. Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nat. Commun.* **5**, 4006 (2014).
149. Wang, J. et al. A predictive model of radiation-related fibrosis based on radiomic features of magnetic resonance imaging. *Int. J. Radiat. Oncol. Biol. Phys.* **105**, E599 (2019).
150. Lin, H. et al. A super-learner model for tumor motion prediction and management in radiation therapy: development and feasibility evaluation. *Sci. Rep.* **9**, 14868 (2019).
151. Mahdavi, S. R. et al. Use of artificial neural network for pretreatment verification of intensity modulation radiation therapy fields. *Br. J. Radiol.* **92**, 20190355 (2019).
152. Zhen, X. et al. Deep convolutional neural network with transfer learning for rectum toxicity prediction in cervical cancer radiotherapy: a feasibility study. *Phys. Med. Biol.* **62**, 8246–8263 (2017).
153. Tomori, S. et al. A deep learning-based prediction model for gamma evaluation in patient-specific quality assurance. *Med. Phys.* **45**, 4055–4065 (2018).
154. Kearney, V., Chan, J. W., Haaf, S., Descovich, M. & Solberg, T. D. DoseNet: a volumetric dose prediction algorithm using 3D fully-convolutional neural networks. *Phys. Med. Biol.* **63**, 235022 (2018).
155. Chen, X., Men, K., Li, Y., Yi, J. & Dai, J. A feasibility study on an automated method to generate patient-specific dose distributions for radiotherapy using deep learning. *Med. Phys.* **46**, 56–64 (2019).
156. Cui, S., Luo, Y., Tseng, H., Ten Haken, R. K. & El Naqa, I. Combining handcrafted features with latent variables in machine learning for prediction of radiation-induced lung damage. *Med. Phys.* **46**, 2497–2511 (2019).
157. Wei, L. et al. Variational autoencoder graph-based radiomics outcome modeling of intrahepatic progression risk and overall survival for HCC post-SBRT patients. *Int. J. Radiat. Oncol. Biol. Phys.* **105**, S83–S84 (2019).
158. Mahmood, R., Babier, A., McNiven, A., Diamant, A. & Chan, T. C. Y. Automated treatment planning in radiation therapy using generative adversarial networks. Preprint at *arXiv* (2018).

Acknowledgements
The authors acknowledge financial support from the US NIH (grants U24CA194354, U01CA190234, U01CA209414 and R35CA220523).

Author contributions
E.H., A.H. and R.H.M. researched data for the article; E.H., A.H., H.J.W.L.A. and R.H.M. contributed substantially to the discussion of content; and E.H., A.H., C.G., D.S.B., H.J.W.L.A. and R.H.M. wrote the article. All authors reviewed/edited the manuscript before submission.

Competing Interests
A.H. is a shareholder of and receives consulting fees from Altis Labs. The other authors declare no competing interests.

Peer review information
Nature Reviews Clinical Oncology thanks I. El Naqa, G. Valdes and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Publisher’s note
Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.