


Cancerworld

AI in Managing Clinical Data and Trials

Adriana Albini / 6 June 2023




AI is reshaping the way we live and our horizons, including medicine and clinical practice. Currently, artificial intelligence is regarded as a significant novel instrument for completing the personalised medicine revolution and creating fresh opportunities to enhance patient management. Held on 17 April 2023, as part of the SPCC Educational Project on Artificial Intelligence in Cancer Care, [this webinar](#) focussed on important issues and topics related to this field. The session was moderated by **Claudio Luchini**, Surgical Pathologist, Diagnostics and Public Health Department, and ARC-Net Research Centre, University and Hospital Trust of Verona, Verona, IT

The Value of AI in Oncology and related fields: from Research to Clinical

Trials

Antonio Pea is Consultant Pancreatic Surgeon at G.B. Rossi Hospital University of Verona, IT, and Honorary Fellow at the Institute of Cancer Science of the University of Glasgow, UK. He is a clinician scientist with a specific focus on translational research and integrative multi-omics analysis in pancreatic cancer. In the context of oncological research and possible AI applications, various research avenues in oncology come to mind. For instance, there are opportunities to **accelerate drug discovery and development** by identifying potential drug candidates and forecasting drug efficacy. Additionally, AI can **advance our understanding** of cancer biology by predicting protein structures and unravelling complex molecular and cellular interactions. It can also **enable personalised medicine through multi-omics integration**. These objectives can be achieved by combining genomics, transcriptomics, proteomics, and metabolomics data, to identify molecular subtypes, forecast patient survival, and tailor treatment strategies based on individual patient profiles.

2) Identifying Relevant Features for AI Models in Oncology Research



In the context of oncology research, artificial intelligence can be applied to digital pathology data, and contribute to many research objectives. For instance, Vivek Nimgaonkar and colleagues developed an AI derived histologic signature associated with response to gemcitabine in pancreatic cancer (Cell Rep. Med. 2023). In another study, deep learning algorithms were used to analyse the sub-tumour microenvironment, particularly the cell-to-cell interaction between cancer cells and immune cells (Barbara T. Grünwald et al., Cell 2021). And in a study published in the Lancet in 2020, a deep learning model was used to predict survival in colorectal cancer patients using histological slides. Currently, over 75% of AI applications in oncology are concentrated in cancer radiology and pathology. This raises the question: why are these two areas being prioritised? Firstly, the nature of image-based data is inherently complex and extensive. Each digital image comprises pixels that can be represented by numbers, and the complete image is a metric of these numbers. Even if we are not pathologists or radiologists, we can see that a vast amount of data can be extrapolated from each single image. Moreover, many tasks in pathology and radiology are objective, making them suitable for visual training and evaluation of AI models. There is a high demand for automation in these fields, and the achievements of AI have garnered strong research and industry support for these

specialties.

What makes AI so effective at analysing complex biological data? One reason is that biological data is often high-dimensional, with numerous features, such as genes, proteins, or imaging characteristics, that can be measured simultaneously. AI techniques, especially deep learning models, are designed to handle high-dimensional data effectively, identifying complex patterns and relationships among the features. Additionally, AI can capture non-linear relationships among data. Biological systems are known for their non-linear relationships and interactions, as seen in gene regulatory networks, where the over-expression of one gene due to a mutation can change the expression of a large number of other genes through the gene interaction network. AI can model these non-linearities and uncover the underlying structure of the data. AI models also possess the characteristic of being robust to noise and capable of efficiently processing large-scale data. Considering the example of large-scale genomic data, such as whole genome sequencing or whole exome sequencing, which can generate millions of genetic variants, AI can manage the noise in sequencing data, such as sequencing errors and alignment issues, while identifying the genetic variants associated with disease or traits relevant to oncological outcomes.


In oncology research, the AI workflow involves several steps, including data collection and pre-processing, feature selection and extraction, model development, validation, and evaluation, as well as interpretation and application. If we want to **prepare our data for our AI analysis in oncology**, we must consider the various **data sources** available in the field of oncology: genomic data, gene expression data, protein expression data, clinical information and digital pathology or radiology images. Clinical information is among the less curated data, as it is less objective compared to sequencing and imaging data, and often contains errors. **Data quality** is essential for AI projects, requiring accurate labelling and annotation of features. For example, if we are doing a digital pathology project, we need expert pathology on that disease to do all the labelling and annotation. Furthermore, it is crucial to identify and remove outliers, correct for batch effects and technical variations, and address all possible sources of bias in the experimental design and data collection. We need **data standardisation**. We need to normalise data for consistent representation, and we need to harmonise data from different sources and platforms. And finally, we must convert our data into a suitable format for analysis. Lastly, it is crucial to determine the type of missing data and select the suitable imputation methods to fill the incomplete data points.

After preparing our data, it is important to identify relevant features for our AI model. Collaboration with domain experts is necessary for this task to rank features and select those that are **biologically meaningful**. Not all extracted features from slides may be meaningful for the outcome.

Dimensionality reduction techniques can be utilised to manage high-dimensional data, prevent overfitting, and eliminate unnecessary variables. Additionally, deep learning techniques can be used for **feature learning**. For instance, in the case of an H&E slide of pancreatic cancer, we can perform a cell segmentation to identify cells and extract various features from each cell, such as stain intensity, morphology, spatial distribution, and relationship with surrounding cells. We can then use these features to classify cells as either tumour or stroma. Although the number of features that can be extracted is extensive, not all of them may be relevant to our outcome of interest. For instance, we can employ deep learning techniques to extract features from the surrounding tissue at different radii, but we must carefully select which features to include in our analysis. To avoid overfitting our model, we need to ensure that each step of the process is accurate. Dr Pea discussed an example of multiplexing fluorescence imaging for pancreatic cancer glands. The task was to accurately identify and separate individual cells within the image, which was challenging due to the irregular shape and simultaneous division of the nuclei. If a method is not working well, we need to change the method or the model being used to improve accuracy. For example, in order to build a classifier that differentiates tumour cells from other cells, the ratio between nucleus and cytokines is

one of the most biological meaningful characteristics, so this step needs to be accurate.

2) Identifying Relevant Features for AI Models in Oncology Research



From the clinician's perspective, there are two important points to consider when using machine learning or deep learning algorithms for medical imaging. The first point is that the desired outcomes of the analysis are important because they determine which algorithm is the most suitable for the task. The second point is that the model should not just work on a single dataset, but it should be effective across different datasets and institutions. This is important for improving clinical decision-making and should not just be for publication purposes. Therefore, independent datasets are necessary for validating the effectiveness of the model.

How can we bridge the gap between research and clinical trials?

1. **Patient stratification and selection:** By analysing patient data, including medical history, genomic information, and biomarkers, AI can identify specific subgroups of patients more likely to respond to a particular treatment. This enables the design of targeted and personalised clinical trials, leading to better outcomes and reduced costs.
2. **In Silico Trials:** AI algorithms can simulate various trial designs and predict the trial's success based on parameters such as sample size, treatment arms, and endpoints. This helps researchers design more efficient and effective trials, minimising required resources and maximising the chances of success.
3. **Real-time monitoring and adaptive trial designs:** AI can monitor trial data in real-time, allowing researchers to identify potential issues or trends early in the process.
4. **Predictive modelling and data extraction:** AI can be used to develop predictive models that estimate patient outcomes and to analyse and extract relevant information from unstructured data sources, such as electronic health records and medical literature.

Considering each phase of the clinical trial, AI can provide valuable insights. In phase I, it can analyse preclinical data to estimate the optimal dose of the drug, identify side effects, and identify biomarkers. In phase II, it can help in patient stratification and selection to identify the most appropriate patient cohorts. In phase III, it can optimise trial design, predict trial success, analyse interim results, and interpret the final trial results. In the post-marketing surveillance phase, AI can

monitor large-scale real-world data to look for adverse events, efficacy trends, and patient subgroups with better or worse response to the drug being studied.

The Role of AI-Pathology in Precision Medicine Clinical Trials

Eric Walk is Chief Medical Officer at PathAI in Boston, USA, and a pathologist with more than two decades of experience in the field of oncology drug and diagnostics development. Prior to his current role, he held the position of CMO at Roche Tissue Diagnostics and was involved in translational management at Novartis Oncology. The use of precision medicine has become the standard for developing cancer therapies and treating patients in real-world settings. The accurate and reproducible assessment of biomarkers is a critical aspect of such model. In clinical trials, whether using a biomarker-enrichment design or a biomarker-stratified design, it is assumed that biomarkers, particularly histopathology ones, are measured accurately and reproducibly. And this, of course, becomes even more important in the clinic, where we are making treatment decisions for patients. Unfortunately, at least in some cases, it is not a safe assumption that biomarkers are assessed in an accurate and reproducible way. As we can garner from [an article](#) published by David Rimm and colleagues in 2022, which compared four different FDA registered tests to detect PD-L1, when it comes to scoring immune cells such as CPS scoring, the reproducibility is very low. The interclass correlation coefficients can be as low as 0.2 for immune cell (e.g. CPS) scoring, which is concerning. TPS scoring is better but still needs significant improvement.

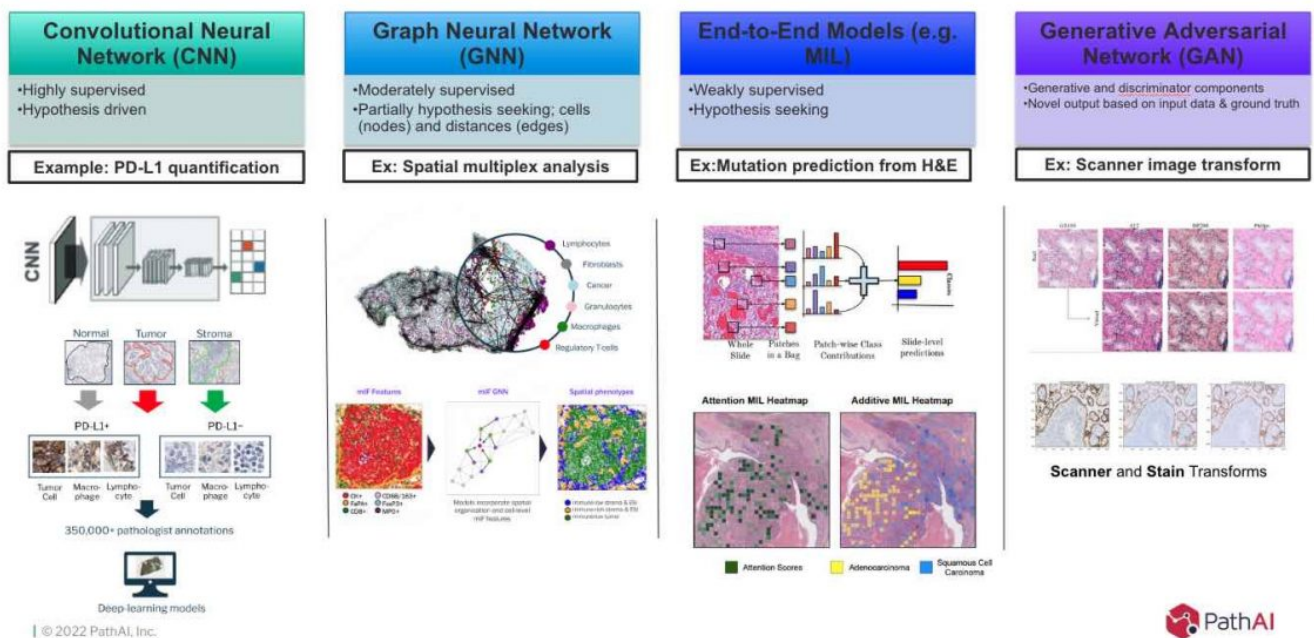
[Another study](#), by Aileen Fernandez, David Rimm, and colleagues, aimed to determine if the current ERBB2 (HER2) assays and interpretation methods can accurately differentiate between 0 and 1+ scores, which is necessary for the HER2-low companion diagnostic test used in clinical trials for trastuzumab deruxtecan. The significance of HER2-low is widely recognized now, particularly in light of the Destiny-04 data. The findings of [this study](#) indicate that while the interrater concordance at the 2+/3+ cut-off is acceptable, the scoring accuracy for HER2 low (0 and 1+) was poor (26% concordance) and could lead to mistreatment in the real world. This is clearly a challenge that we need to address because while the drug is clearly active, there is also a significant issue with interpretation accuracy and reproducibility. The same issue applies outside of oncology as well. Thinking, for example, of non-alcoholic steatohepatitis (NASH), there are biomarkers used such as the NAFLD activity score and the fibrosis score, measured by H&E and trichrome staining respectively. These biomarkers not only help select patients for clinical trials but also are used as endpoints to measure drug effectiveness. The issue, similar to PD-L1 and HER2-low, is that these important tissue biomarkers suffer from poor reproducibility. David Kleiner's studies, conducted 14 years apart, show that there has been no meaningful improvement in reproducibility of these key scoring parameters, despite major efforts to improve education and training. In recent years, there has been increased attention in research and literature on the significance of NASH pathology scoring variability in achieving accurate patient staging and endpoint efficacy evaluation in clinical trials. Liver biopsy remains the standard for diagnosing and tracking NASH progression, and is currently the primary inclusion criteria and endpoint in NASH clinical trials. The NASH Clinical Research Network (CRN) fibrosis staging system is the most validated approach for evaluating changes in disease stage in these trials. Despite being validated, the NASH-CRN manual scoring system is prone to both inter- and intra-observer variability. This variability can potentially reduce the likelihood of observing the true drug effect by up to 32%, which in turn hinders the success of clinical trials and prevents promising therapies from reaching patients.

How can AI help mitigate or solve some of these problems? AI pathology can be utilised in two ways to improve and enable precision medicine in 2023: firstly, it can be used to enhance the abilities of human pathologists by providing consistent and quantitative biomarker results at scale.

This is achieved by reducing the impact of intra- and inter-observer variability, improving patient selection strategies, and standardising the assessment of biomarkers, histologic scores and endpoints. For example, AI can accurately and reproducibly count all tumour cells and immune cells and indicate which fraction of each are positive for PD-L1, a challenging task for pathologists to perform manually. Secondly, AI can identify signatures or biomarkers that are indiscernible to the human eye. Spatial location-based cellular and tissue relationships are becoming increasingly relevant as predictive and prognostic biomarkers, but assessing spatial patterns of multiple cell types and their locations is impossible for humans without advanced analytics. This area is just now being explored in depth and has tremendous potential for the future of precision medicine.

Before moving on to some examples, Dr Walk provided a brief overview of different machine learning methods. These include convolutional neural networks, graph-neural networks, end-to-end models, and generative adversarial networks. The variations between these methods are quite distinct. Convolutional neural networks are highly supervised and require a pre-existing hypothesis, which makes them suitable for applications such as PD-L1 scoring. On the other hand, end-to-end models are hypothesis-seeking and weakly supervised, allowing the models to identify correlations that may be imperceptible to humans. Graph-neural networks fall in the middle of this spectrum, while generative adversarial network technology occupies the opposite end. Generative modelling refers to a type of machine learning task that falls under unsupervised learning, where the aim is to automatically identify and learn the patterns or regularities present in a given set of input data. The objective is to develop a model that can generate or produce new examples that are similar in nature to the original dataset. This is how ChatGPT is powered. We can use it to generate synthetic data and transform image quality. For instance, we can transform one scanner image into another, which can help during the training of machine learning models.

Matching Machine Learning Methods to Application Use Cases



Dr Walk then showed an example of a traditional convolutional neural network model that was developed by PathAI to address the PD-L1 scoring problem. This model was trained using 350,000 human pathologist annotations and was able to discern the difference between PD-L1 negative and positive cells as well as between tumour and normal tissue in lung cancer. The model was initially validated by comparing the AI output with human pathologist counts in 150 x 150 pixel squares of the image prior to larger scale validation. The workflow for a pathologist using this AI tool would be

to consider the AI score in the context of the tissue and cellular visualizations (heatmap overlays) indicating how the tool has assessed different tumour tissue regions (e.g. tumour, stroma, necrosis) and individual cells (tumour, immune, PD-L1 +/-, etc.) The final task for the pathologist is to accept or reject the score prior to reporting. PathAI presented a study at the AACR meeting in 2022, that validated the accuracy of the AI model for PD-L1 scoring in 350 cases of lung cancer. The study compared the consensus score of 12 human pathologists with the AI score, and found a high correlation of 0.93, indicating that the model is accurate. Interestingly, the study also revealed that some pathologists perform much higher against consensus than others, highlighting the variability in manual scoring that the AI model aims to address.

Another way to look at accuracy is to look at actual outcome data. A retrospective analysis of IO outcome data was conducted by PathAI in collaboration with Bristol Myers Squibb (BMS), using OPDIVO (nivolumab) clinical trials, specifically the Phase 3 CheckMate 57 and 26 studies. Compared to the original manually derived PD-L1 prevalence figures, AI scoring resulted in a substantial increase in the number of patients assessed as PD-L1 positive. The study aimed to determine if the AI-positive patients were falsely or truly positive. Recalculation of the survival data revealed that the recurrence-free survival of dual positive patients (positive both by AI and manual) and AI-only positive patients were similar. This suggests that the AI-positive patients may be biologically and clinically relevant and could respond well to treatment. Of course, the results are based on retrospective data and require careful interpretation, but they certainly are intriguing.

Another machine learning technique, graph-neural network (GNN), is a pattern-based method which involves identifying nodes that represent nuclei or cells, and edges that represent the distance between those cells. The algorithm then analyses whether a spatial pattern connecting different cells on a slide corresponds to the endpoint of interest. In the example given by Dr. Walk, the technique was used to develop a model for CD8 cytotoxic T-cells that could automatically generate readouts for different immune phenotypes such as hot/inflamed, cold/non-inflamed, and immune excluded. These patterns have been shown to correlate with response to immune oncology therapies but suffer from subjectivity and poor reproducibility. When the PathAI CD8 immune phenotyping algorithm was applied to the CheckMate 67 study, comparing NIVO versus NIVO + IPI in melanoma, the algorithm identified 38% more responding patients who were PD-L1 negative/CD8-excluded compared with predicting response via PD-L1 status alone. In addition to increasing the responsive population, the benefit improved, with a hazard ratio of 0.46 vs. 0.37 for PD-L1 alone. This is an example of how AI pathology can be utilized to discover and reveal novel biomarkers and potentially benefit patient outcomes.

Outside of oncology, in conditions like NASH, for instance, AI algorithms can measure and evaluate biomarkers in a similar way, generating continuous biomarkers with a lot of potential for clinical trials. To illustrate this, Dr. Walk gave two examples of retrospective NASH clinical trials that were reanalysed with AI. In the first example, a phase II study of a drug called pegbelfermin, the human glass analysis showed no statistically significant difference between the placebo and treatment arms. However, when the slides were digitised and run through the AI-NASH algorithm, a highly statistically significant difference was generated, with more of a drug response curve as well. In the second example, a different NASH drug called semaglutide was analysed. The placebo response rate generated through human analysis was very high, but when reanalysed with AI, this dropped dramatically. This is important because an artificial increase in the placebo response rate can make it much more difficult to show the benefit of the drug. Lastly, AI has the potential to create continuous biomarkers in fields where ordinal biomarkers are currently used. For example, a machine learning-based continuous fibrosis score has the potential to increase sensitivity and resolution to a detect drug response in NASH clinical trials.

In sum, ensuring the accurate and consistent measurement of tissue-based biomarkers in clinical

trials is crucial for precision medicine in various fields, including oncology and NASH. However, existing biomarkers often have high inter-reader variability, which can be problematic. Fortunately, AI pathology can help address this issue by assisting pathologists in assessing biomarkers like PD-L1, HER2, and NASH NAS, thereby reducing reader variability and improving clinical trial enrolment and endpoint analysis. Moreover, AI pathology can aid in the development of novel predictive biomarkers by utilising advanced machine learning techniques like GNNs and CD8 spatial phenotypes. This shows the potential of AI pathology to contribute to the field of precision medicine by improving the accuracy and reproducibility of tissue-based biomarker measurements.

AI applications in healthcare

Peter Krusche, Director of Data Science in the Advanced Methodology and Data Science team at Novartis, Basel, CH, is a computer scientist who specialises in the development and improvement of clinical study planning at Novartis. His team is responsible for the development, evaluation, and deployment of new statistical, AI, and ML methodologies. While their work is primarily theoretical in nature, it is applied in the context of drug development. As we tend to bring AI together with other concepts from statistics, from analytics and so on, Dr. Krusche started with a definition of artificial intelligence: it is a technique used to train computers to perform complex tasks by providing them with examples. A simple example is using images of cats and dogs to teach the computer how to distinguish between them. This type of input could also be medical images, as discussed in previous talks. The model used in AI consists of a set of parameters that describe what the model is looking for, and a set of computational instructions that explain how the parameters are used to distinguish between cats and dogs. AI is not limited to this basic approach, however. It can also take high-dimensional data such as images and convert them to low-dimensional data, such as a single value indicating whether it is a cat or a dog. Conversely, AI can take a model and use randomness to generate new images of cats or other types of medical images. This capability makes it much more versatile, allowing for the combination of data from different modalities.

Generative AI, or GenAI, is a form of artificial intelligence that is capable of generating a diverse range of data, including images, videos, audio, text, and 3D models. It accomplishes this by studying patterns in existing data and applying this knowledge to create novel and distinctive outputs. For instance, the system generates a text embedding of an input sentence, followed by utilising an image generator to produce the corresponding image. The two models work in tandem to generate the desired output. Among real-world examples, the input for ChatGPT or Bing Chat is text, and the output is also text. DALL-E, StableDiffusion, Bing Image Creator, etc., use text as input and generate images as output. Segment Anything is a model with the ability to segment images in a highly effective and adaptable manner. And BioGPT is an example of an open-source tool that can make predictions using a training set derived from PubMed.

If all of these tools are out there, **why is applying AI in a medical or scientific setting difficult?** Firstly, in clinical decision making the stakes are very high; decisions affect human lives very directly. Moreover, patient data are personal and cannot be shared or combined easily. Secondly, clinical data are difficult. Even data from clinical trials contain missingness and training data are often sparse. Noise and uncertainty need to be modelled explicitly, and this is not easy. Lastly, data collection is often not standardised: even transferring medical images between doctors may not entirely be straightforward. There are also some methodological difficulties with machine learning models. The first one is that complex models can be made to behave in unexpected ways. A classic example, from Ian Goodfellow and colleagues' paper "Explaining and Harnessing Adversarial Examples", is the panda and the gibbon. The model is given the picture of a panda and recognises it as such with relatively high confidence. However, adding noise that is not visible to the human eye makes the model fairly confident that it is a gibbon. The implications of that in medical imaging can

be left to the imagination. Another issue is that models are not perfect, and when they generate output randomly, they might generate things that are not real. When we rely on models that build complex outputs based on randomness, we cannot always expect them to generate things that are rooted in reality without a lot of feedback.


There are other challenges to consider as well. Firstly, the training process itself may introduce biases or be inherently biased, particularly in the case of imaging data. For example, a model may focus on learning about skin markings rather than the diagnostic properties of medical images. Secondly, these methods may be misused by malicious actors, and without a complete understanding of their potential, it may be difficult to defend against such attacks. Finally, we must remember that methodology research does not always prioritise solving real-world problems effectively. Many of the metrics used to develop these methods may not align with the most practical applications. How can we address these issues? This is an ongoing effort, and the best approach would be to implement regulatory guidelines and ethical principles for AI. Dr. Krusche mentioned three sources that showcase the advancements made in this area: the FDA has released draft guidance on software as a medical device, which outlines the development and testing process that AI models must adhere to before being implemented. The European Commission has also developed guidelines for ensuring the trustworthiness of AI. Finally, an interesting talk from a major machine learning conference sheds light on how to assess the validity of AI models in the public domain from the perspective of the methodology field.

(<https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device>; <https://ec.europa.eu/futurium/en/ai-alliance-consultation.1.html>; <https://neurips.cc/virtual/2022/invited-talk/55868>)

Considering the potential of generative AI in the future, the first application that comes to mind is text generation, given the vast amount of text produced daily for tasks such as writing papers, abstracts, presentations, notes, and summaries. The second application is in **finding candidates for hidden complex relationships** through pathway analysis and smarter literature research. The final application is in finding biomarkers and **supporting diagnostic processes**. What all these implications have in common is that they shift the human role from task execution or creation to quality control. However, it is important to note that quality control is not always easier than creation and must be carefully considered when deploying generative AI. While AI can provide helpful summaries and aid in learning, subject matter expertise is still crucial, as AI may make errors or overlook important details.

What is AI?

Training data
(examples)



Training data contains a large set of examples for inputs and expected outputs.

Example: labelled pictures.

Dr. Krusche concluded his presentation highlighting some areas that his group works on: **good practice** for data science, which is crucial for achieving other objectives as well; **conditional data synthesis** that combines medical images and clinical data to create interpretable machine learning methods. This method can bridge the gap between clinical trials, real-world data, and images. Finally, **privacy and synthetic data**. We must ensure models preserve individual privacy. Overall, we need to acknowledge that AI is a new tool that needs to be differentiated from other tools and operationalised. The debate surrounding AI deployment is reminiscent of the introduction of calculators or computers in schools, which were initially controversial but ultimately unstoppable. We need to develop the appropriate skills, practices, and legal and regulatory foundations to ensure that AI tools and methods are deployed safely and do not cause harm.

Patient Involvement and Empowerment in AI Clinical Dimensions

Elliot K. Fishman is Professor of Radiology, Surgery, Oncology and Urology at Johns Hopkins Hospital, Baltimore, US. There is no doubt that artificial intelligence has the potential to revolutionise medicine, improving patient care and the physician-patient relationship. Every aspect of clinical trials can be enhanced with AI, from trial design to patient recruitment, outcome monitoring, and reducing patient dropout. But **what about the patient perspective?** For the past seven years Prof. Fishman has been working with AI for the detection of pancreatic cancer, and in his experience, patients are very enthusiastic when they learn that AI can detect pancreatic cancer at an early stage, leading to a potentially considerable improvement in survival rates. From AI and deep learning to radiomics, the multitude of possibilities in the field of medicine emphasises the importance of analysing, collecting, and understanding data. Whether it is detecting disease, monitoring patients' suitability for specific therapies, or predicting outcomes, these factors are all significant, and patients tend to be very supportive when presented with the potential benefits.

An article published in March this year in the European Journal of Radiology looked at liver metastasis and found that AI could detect half of the metastases that were missed by radiologists reading the studies. It is very easy for patients to get on board when shown these significant improvements. However, they are also concerned about the failures of AI. For instance, in recent years a number of articles were published about using a sepsis prediction algorithm. Being able to

predict which patients are at risk for sepsis would enable healthcare providers to proactively manage those patients. On the basis of this, one of the large EHR vendors, EPIC, developed a program. However, when implemented, it only picked up 7% of the patients with sepsis who could have been treated earlier, and failed to identify 1709 patients with sepsis already identified by the hospital. It was very unsuccessful and had to be withdrawn. Although there was no published evaluation of that program, it still had been adopted by a large number of hospitals because of its convenience and availability.

Proper evaluation and publication of results is fundamental to ensure the success and safety of AI implementations. Most of the 500 AI algorithms that the FDA has approved have limited datasets from only one institution, which makes it difficult to generalise the results to other institutions. This lack of data concerns patients, who want to be part of the conversation about AI's implementation. In a recent paper, Macrì and Roberts argue that patient involvement is crucial for AI to be successful, and shared decision-making is essential. Physicians need to have a clear dialogue with patients to build trust and ensure that patients see AI as a valuable tool in their treatment prospects. Collaboration and shared decision making between physicians and patients is critical in the implementation of AI in clinical trials and patient care. Patients must be informed of their options and the potential risks involved in the use of artificial intelligence. It is important for physicians to explore patient-specific values associated with the implementation of AI and apply them to clinical decision making.

An article published earlier this year by Alexandra Derevianko and her group discusses the importance of patient-doctor communication in AI-aided cancer diagnosis. The results again point to the fact that without a clear understanding and communication between patients and doctors about the risks and benefits of AI, the technology will not be able to fulfil its potential. We also need to keep in mind that the rapidly changing AI landscape can lead to confusion and fear among patients, thus communication and education should be ongoing to help patients understand the benefits and limitations of the technology. Particularly in the context of clinical trials, companies developing AI algorithms should focus not just on technical development, but also on how to effectively communicate progress to patients. The field of AI is constantly evolving, and new discoveries can change our understanding and progress. It is important for patients to trust that healthcare professionals will make the right choices.

In terms of clinical trials, A [Deloitte article posted in 2020](#) talked about how AI can facilitate the development of patient-centric clinical trial designs by optimising and accelerating the process. It can also drive novel methods of data collection that minimise reliance on traditional in-person trial sites. For instance, body sensors and wearable devices like heart monitors, patches, and sensor-enabled clothing can be used to remotely monitor patients' vital signs and other data, reducing the need for invasive methods. AI can be combined with robotic process automation to link and harmonise data across various modalities of data collection. By applying machine learning to clinical data, it can illuminate complex relationships between different data domains and facilitate automated data management. Additionally, natural language generation can be used to auto-generate content for trial artifacts, thus streamlining and accelerating the process of creating regulatory documents.

One of the strengths of AI is its ability to predict things that may not have been considered before. For example, Prof. Fishman's team conducted a study on detecting and managing patients with cystic pancreatic lesions, which will shortly be submitted for publication. A few years ago, they published the results of a trial with over 800 patients across multiple countries where it was found that 40% of patients required surgery while the remaining 60% did not. Later, using a program called CompCyst that analysed the data, they were able to predict accurately in 60% of cases. However, they recently developed a new program using transparent AI, which increased accuracy to

90% using the same data. This demonstrates that computers and the constantly changing technology present new opportunities to think differently and achieve higher levels of accuracy.

Despite the potential benefits of AI in healthcare, a survey from the Pew Research Center revealed that 6 out of 10 adults feel uncomfortable with the idea of AI being used to diagnose them. Only 38% of respondents believed that AI could lead to better health outcomes, while 33% believed that it would worsen their health and 27% were unsure. The survey also found that men, younger individuals, and those with higher levels of education were more receptive to the use of AI in clinical settings. However, the level of comfort varied depending on the purpose of the AI application. Patients were more comfortable with AI making minor decisions, such as examining chest x-rays, but less comfortable with AI making cancer diagnoses. The survey revealed that patients who reported having little to no knowledge about AI were more likely to feel uncomfortable with their provider using it compared to those who were familiar with it. Concerns were raised by three quarters of respondents that healthcare providers are implementing AI tools too quickly without understanding the risks, while only 23% felt they were moving too slowly.

The results of a survey published by Dhruv Khullar in JAMA indicate that a significant proportion of patients are uncomfortable with receiving a diagnosis from an AI algorithm, even if it is 90% accurate. This is despite the fact that many patients acknowledge that AI can improve healthcare outcomes. Patients have expressed that they trust their clinician more than AI, even if the clinician's accuracy is lower. It is important to recognise that patients' comfort with AI is dependent on its specific application, and work must be done to ensure that patients understand how it can be used in conjunction with clinicians to provide the most accurate diagnoses and treatments. The survey revealed that a majority of respondents expressed worries about potential negative consequences of AI use in healthcare such as misdiagnosis, privacy violations, reduced interaction with clinicians, and increased healthcare expenses. These concerns demonstrate the need to actively involve patients in the process of AI implementation and to address their apprehensions. Patients are often afraid of "black boxes", so it is important to have explainable AI that is transparent in its decision-making process. Many tech companies such as Microsoft, Google, Apple, and Facebook are focussing on developing explainable AI.

The New England Journal of Medicine has recognised the significance of AI by introducing a dedicated column that examines its potential impact on healthcare. Their point is that AI is changing how we are practising medicine. Just as computer acquisition of radiographic images did away with the x-ray file room and lost images, AI machine learning can transform medicine. To the question, "Will AI put radiologists out of business?" the answer will be that the technology will not put health professionals out of business, but will enable them to work better. However, radiologists who refuse to use AI risk to become obsolete. The challenge lies in implementing AI and connecting with patients to ensure that human-human interactions remain a vital aspect of medicine. It is crucial to address this issue to ensure that patients can take advantage of the advancements in healthcare.