Domande estratte n. 2 e n. 4

Domanda Specifica 1

Cosa sono i Common Data Models?

Domanda Specifica 2

Cos'è il Natural Language Processing?

Domanda Specifica 3

Cos'è la Network Medicine?

Domanda Specifica 4

Quali sono le metriche di qualità di un sistema di Machine Learning / Apprendimento Automatico?

Domande estratte n. 3 e n. 1

Domanda Informatica 1

Cos'è il linguaggio HTML e in che contesto viene utilizzato?

Domanda Informatica 2

Cos'è un Browser? Puoi fare degli esempi?

Domanda Informatica 3

Cos'è il protocollo TCP/IP?

Domanda Informatica 4

Come si misurano le prestazioni di un microprocessore?

Porzione di testo tradotta: n. 3 e n. 4

Published in final edited form as:
Cancer Cell. 2021 July 12; 39(7): 916–927. doi:10.1016/j.ccell.2021.04.002.

Artificial Intelligence for Clinical Oncology

Benjamin H Kann[1,2], **Ahmed Hosny**[1,2], **Hugo JWL Aerts**[1,2,3,4]

[1]Artificial Intelligence in Medicine (AIM) Program, Mass General Brigham, Harvard Medical School, Boston, United States of America. [2]Department of Radiation Oncology, Brigham and Women's Hospital, Dana-Farber Cancer Institute, Harvard Medical School, Boston, United States of America. [3]Department of Radiology, Brigham and Women's Hospital, Dana-Farber Cancer Institute, Harvard Medical School, Boston, United States of America. [4]Radiology and Nuclear Medicine, CARIM & GROW, Maastricht University, Maastricht, The Netherlands.

SUMMARY

Clinical oncology is experiencing rapid growth in data that are collected to enhance cancer care. With recent advances in the field of Artificial Intelligence (AI), there is now a computational basis to integrate and synthesize this growing body of multi-dimensional data, deduce patterns, and predict outcomes to improve shared patient and clinician decision-making. While there is high potential, significant challenges remain. In this perspective, we propose a pathway of clinical, cancer care touchpoints for narrow-task AI applications and review a selection of applications. We describe the challenges faced in the clinical translation of AI and propose solutions. We also suggest paths forward in weaving AI into individualized patient care, with an emphasis on clinical validity, utility, and usability. By illuminating these issues in the context of current AI applications for clinical oncology, we hope to help advance meaningful investigations that will ultimately translate to real-world clinical use.

INTRODUCTION

Over the last decade, there has been a resurgence of interest for artificial intelligence (AI) applications in medicine. This is driven by the advent of deep learning algorithms, computing hardware advances, and the exponential growth of data that are being generated and used for clinical decision making (Esteva et al., 2019; Kann et al., 2020a; LeCun et al., 2015). Oncology is particularly poised for transformative changes brought on by AI, given the proven advantages of individualized care and recognition that tumors and their response

Corresponding author: Hugo Aerts, PhD, Director, Artificial Intelligence in Medicine (AIM) Program, Harvard-MGB, Harvard Institutes of Medicine – HIM 343, 77 Avenue Louis Pasteur, Boston, MA 02115 P - 617.525.7156, F - 617.582.6037, haerts@bwh.harvard.edu.

Declarations of Interests

H.A. is a shareholder of and receives consulting fees from Onc.Ai and BMS, outside submitted work. A.H. is a shareholder of and receives consulting fees from Altis Labs, outside submitted work.

rates differ vastly from person to person (Marusyk et al., 2012; Schilsky, 2010). In oncology, much like other medical fields, the overarching goal is to increase quantity and quality of life, which, from a practical standpoint, entails choosing the management strategy that optimizes cancer control and minimizes toxicity.

2

As multidimensional data is increasingly being generated in routine care, AI can support clinicians to form an individualized view of a patient along their care pathway and ultimately guide clinical decisions. These decisions rely on the incorporation of disparate, complex datastreams, including clinical presentation, patient history, tumor pathology and genomics, as well as medical imaging, and marrying these data to the findings of an ever-growing body of scientific literature. Furthermore, these datastreams are in a constant state of flux over the course of a patient's trajectory. With the emergence of AI, specifically *deep learning* (LeCun et al., 2015), there is now a computational basis to integrate and synthesize these data to predict where the patient's care path is headed, and ultimately improve management decisions.

While there is much reason to be hopeful, numerous challenges remain to the successful integration of AI in clinical oncology. In analyzing these challenges, it is critical to view the promise, success, and failure of AI not only in generalities, but on a clinical case-by-case basis. Not every cancer problem is a nail to AI's hammer; its value is not universal, but inextricably linked to the clinical use case (Maddox et al., 2019). The current evidence suggests that clinical translation of the vast majority of published, high-performing AI algorithms remains in a nascent stage (Nagendran et al., 2020). Furthermore, we posit that the imminent value of AI in clinical oncology is in the aggregation of narrow task-specific, clinically validated and meaningful applications at clinical "touchpoints" along the cancer care pathway, rather than general, all-purpose AI for end-to-end decision-making. As the global cancer incidence increases and the financial toxicity of cancer care is increasingly recognized, many societies are moving towards value-based care systems (Porter, 2009; Yousuf Zafar, 2016). With development of these systems, there will be increasing incentive for the adoption of data-driven tools - potentially powered by AI - that can lead to reduced patient morbidity, mortality, and healthcare costs (Kuznar, 2015).

Here, we will describe the key concepts of AI in clinical oncology and review a selection of AI applications in oncology from the lens of a patient moving through clinical touchpoints along the cancer care path. We will therein describe the challenges faced in the clinical translation of AI and propose solutions, and finally suggest paths forward in weaving AI into individualized patient cancer care. By illuminating these issues in the context of current AI applications for clinical oncology, we hope to provide concepts to help drive meaningful investigations that will ultimately translate to real-world clinical use.

Artificial Intelligence: from shallow to deep learning

The concept of AI, formalized in the 1950's, was originally defined as the ability of a machine to perform a task normally associated with human performance (Russell and Haller, 2003). Within this field, the concept of machine learning was born, which refers to an algorithm's ability to learn data and perform tasks without explicit programming (Samuel, 1959). Machine learning research has led to development and use of a number of "shallow"

learning algorithms, including earlier generalized linear models like logistic regression, Bayesian algorithms, decision-trees, and ensemble methods (Bhattacharyya et al., 2019; Richens et al., 2020). In the simplest of these models, such as logistic regression, input variables are assumed to be independent of one another, and individual weights are learned for each variable to determine a decision boundary that optimally separates classes of labelled data. More advanced shallow learning algorithms, such as random forests, allow for the characterization and weighting of input variable combinations and relationships, thus learning decision boundaries that can fit more complex data.

Deep learning is a newer subset of machine learning, which has the ability to learn patterns from raw, unstructured input data by incorporating layered neural networks (LeCun et al., 2015). In supervised learning, which represents the most common form within medical AI, a neural network will generate a prediction from this input data and compare it to a "ground truth" annotation. This discrepancy between prediction and ground truth is encapsulated in a loss function which is then propagated back through the neural network, and over numerous cycles, the model is optimized to minimize this loss function.

For the purpose of clinical application, we can view AI as a spectrum of algorithms, the utility of which are inextricably linked to the characteristics of the task under investigation. Thorough understanding of the data stream is necessary to choose, develop, and optimize an algorithm. In general, deep learning networks offer nearly limitless flexibility in input, output, architectural and parameter design, and thus are able to fit vast quantities of heterogeneous and unstructured data never before possible (Esteva et al., 2017). Specifically, deep learning has a high propensity to learn non-linear and high-dimensional relationships in multi-modal data including time series data, pixel-by-pixel imaging data, unstructured text data, audio/video data, or biometric data. Data with significant spatial and temporal heterogeneity are particularly well-suited for DLNNs (Zhong et al., 2019). On the other hand, this power comes at the expense of limited interpretability and a proclivity for overfitting data if not trained on a large enough dataset (Zhu et al., 2015). While traditional machine learning and statistical modeling can perform quite well at certain predictive tasks, they generally struggle to fit unprocessed, unstructured, and high dimensional data compared to deep learning. Therefore, despite its limitations, deep learning has opened the door to big data analysis in oncology and promises to advance clinical oncology, so long as certain pitfalls in development and implementation can be overcome.

Cancer care as a mathematical optimization problem

To appreciate the promise surrounding AI applications for clinical oncology, it is essential to incorporate a mathematical lens to the patient care path through cancer risk prediction, screening, diagnosis and treatment. From the AI perspective, the patient path is an optimization problem, wherein heterogeneous data streams converge as inputs into a mathematical scaffold (i.e. machine learning algorithms) (Figure 1). This scaffold is iteratively adjusted during training until the desired output can be reliably predicted and an action can be taken. In this setting, an ever-growing list of inputs include patient clinical presentation, past medical history, genomics, imaging, and biometrics, and can be roughly subdivided as tumor, host, or environmental factors. The complexity of the algorithms is

often driven by the quantity, heterogeneity, and dimensionality of such data. Outputs are centered, most broadly, on increasing survival and/or quality of life, but are often evaluated by necessity as a series of more granular surrogate endpoints.

Datastreams for clinical oncology

The arc of research in oncology, increasing data generation, and advances in computational technology have collectively resulted in a frameshift from low-dimensional to increasingly high-dimensional patient data representation. Earlier data and computational limitations often necessitated reducing unstructured patient data (e.g. medical images and biopsies) into a set of human-digestible discrete measures of disease extent. One notable example of such simplification lies within cancer staging systems, most prominently the AJCC TNM classification (Amin et al., 2017). In 1977, with only three inputs commonly available - tumor size, nodal involvement, and presence of metastasis - the first edition AJCC TNM staging became standard of care for risk-stratification and decision-management in oncology. Over the subsequent decades, with the incorporation of other discrete data points, predictive nomograms could be generated using simple linear models, which have found practical use in certain situations (Bari et al., 2010; Creutzberg et al., 2015; Mittendorf et al., 2012; Stephenson et al., 2007). More recently, improved methods to extract and analyze existing data coupled with new data streams and a growing understanding of inter- and intra-tumoral heterogeneity, have all led to the development of increasingly complex and specific stratification models. Key examples of novel data streams introduced over the past two decades are the Electronic Health Record, The Cancer Genome Atlas (Weinstein et al., 2013), The Cancer Imaging Archive (Clark et al., 2013), and the Project GENIE initiative (AACR Project GENIE Consortium, 2017). Key examples of advanced risk-stratification and prediction models are the prostate cancer Decipher score (Erho et al., 2013) and breast cancer OncotypeDx score (Paik et al., 2004), which utilize discrete genomic data and shallow machine learning algorithms to form clinically validated predictive models. Useful oncology datastreams, roughly following historical order of availability, include: clinical presentation, tumor stage, histopathology, qualitative imaging, tumor genomics, patient genomics, quantitative imaging, liquid biopsies, electronic medical record mining, wearable devices, and digital behavior (Figure 1). Furthermore, as a patient moves along the cancer care pathway, the number of influxing, intra-patient datastreams grows. With each step through the pathway, new data is generated out of the pathway with the potential to be reincorporated at a later time back into the pathway (Figure 2).

3 As our biological knowledge base and datastreams grow in clinical oncology, machine learning algorithms can be deployed to learn patterns that apply to more and more precise patient groups and generate predictions to guide treatment for the next, "unseen" patient. As we assimilate more data, *optimal* cancer care, i.e. the care that results in the best survival and quality of life for a patient, inevitably becomes *precision* care, assuming we have the necessary tools to fully utilize the data. Here, at this intersection of data complexity and precision care in clinical oncology, is where the promise of AI has been so tantalizing, though as of yet, unfulfilled.

AI Applications and Touchpoints along the Clinical Oncology Care Path

We propose that AI development for clinical oncology should be approached from patient and clinician perspectives across the following cancer care touchpoints: Risk Prediction, Screening, Diagnosis, Prognosis, Initial Treatment, Response Assessment, Subsequent Treatment, and Follow-up (Figure 2). The clinical touchpoint pathway shares features with the "cancer continuum," (Chambers et al., 2018) though it consists of more granular, patient and clinician decision-oriented points of contact for AI to add clinical benefit. Each of these touchpoints involves a critical series of decisions for oncologists and patients to make and yields a use-case for AI to provide an incremental benefit. Furthermore, touchpoint details will vary by cancer subtype. Within these touchpoints, ideal AI use-cases are ones with significant unmet need and large available datasets. In the context of supervised machine learning, these datasets require robust and accurate annotation to form a reliable "ground-truth" on which the AI system can train.

Narrow tasks with high reliability

As clinical oncology datastreams increase in complexity, the tools needed to discern patterns from these data are necessarily more complex, as well. Amidst this flood of heterogeneous *intra-patient* data, there is a relative dearth of *inter-patient* data which is needed to train large scale models. Therefore, to accumulate the training data required for generalizable models, it will likely be more fruitful to target and evaluate individual AI models towards specific datastreams at a particular touchpoint along the care pathway.

It is tempting to think that, given the increasing data streams that encompass multiple patient characteristics and outcomes, one could develop a unifying, dynamic model to synthesize and drive precision oncology, developing a "virtual-guide" of sorts for the oncologist and patient (Topol, 2019). Analogies are often made to transformative technologies, such as self-driving cars and social media recommendations that leverage powerful neural networks on top of streams composed of billions of incoming data points, to predict real-time outcomes and continually improve performance. While in theory, this strategy could one day be deployed in a clinical setting, there are vast differences between these domains that question whether or not we *should* or even *could* pursue this strategy currently. One of the most glaring differences between the healthcare and technology domains, in terms of AI application, is the striking difference in data quality and quantity. While there has been a sea change in the collection of data within the healthcare field over the past decade, driven by the adoption of the Electronic Health Record, datasets still remain virtually siloed, intensely regulated, and, particularly in cancer care, much too small to leverage the most powerful AI algorithms available (Bi et al., 2019; Kelly et al., 2019). One of the most high-profile of these endeavors, IBM's Watson Oncology project, has attempted to develop a broad prediction machine to guide cancer care, but has been limited by suboptimal concordance with human oncologists' recommendations and subsequent distrust (Gyawali, 2018; Lee et al., 2018; Somashekhar et al., 2017).

4 As our biological perspective has evolved, we now know that cancer is made up of thousands of distinct entities that will follow different trajectories, each with different treatment strategies (Dagogo-Jack and Shaw, 2018; Polyak, 2011). In computational model

development, there is thought to be a bare minimum number of data samples required for each model input feature (Mitsa, 2019). As we seek to make recommendations more and more bespoke, it becomes more challenging to accrue the quantity of training data necessary to leverage complex algorithms. Fortunately, this data gap in healthcare is well-recognized, and a number of initiatives have been proposed to streamline and unify data collection (Wilkinson et al., 2016). However, given the innately heterogeneous, fragmented, and private nature of healthcare data, we in the oncology field may never achieve a level of data robustness enjoyed by other technology sectors. Therefore, strategies are necessary to mitigate the data problem, such as proper algorithm selection, model architecture improvements, data preprocessing, and data augmentation techniques. Above all, thoughtful selection of narrow use cases across cancer care touchpoints is paramount in order to yield clinical impact.

Once rigorously tested, these narrow-AIs could then be aggregated over the course of a patient's care to provide a measurable, clinical benefit. This sort of AI-driven dimensionality reduction of a patient's feature space allows for optimizing the development process and exporting of quality AI applications in the present environment of siloed data, expertise, and infrastructure. As of writing, there are approximately 20 FDA-approved AI applications targeted specifically for clinical oncology, and each of these performs a narrow task, utilizing a single data stream at a specific cancer care touchpoint (Benjamens et al., 2020; Hamamoto et al., 2020; Topol, 2019) (Table 1). We hypothesize that the future of AI in oncology will continue to consist of an aggregation of rigorously evaluated, narrow-task models, each one providing small, incremental benefits for patient quantity and quality of life. In the next sections, we will review select AI applications that have excelled with this narrow-task approach.

Narrow-task AI examples across the clinical oncology touchpoints

T1. Risk Prediction and Prevention.—Given the burden to people and healthcare systems of cancer diagnosis and management, there is a significant opportunity for AI to help predict an individual's risk of developing cancer, and thereby target screening and early interventions effectively and efficiently. In a mathematical sense, the patient's entire personal history up until diagnosis makes up a vast and extremely heterogeneous datastream to be evaluated, positioning deep learning to have an impact. This is evidenced by the steady development of tools that leverage computational modeling to refine cancer risk. In the past few years, several DL algorithms have been investigated to further tailor risk prediction beyond traditional models. Some of these algorithms utilize novel datastreams that were not available until recently: satellite imagery (Bibault et al., 2020), internet search history (White and Horvitz, 2017), and wearable devices (Beg et al., 2017). Others maximize the utility of pre-existing datastreams, including patient genomics, routine imaging, unstructured health record data, and deeper family history to improve predictions (Ming et al., 2020).

T2. Screening.—Cancer screening involves the input and evaluation of data at a distinct time-point to determine whether or not additional diagnostic testing and procedures are warranted. Datastreams can be in the form of serum markers, medical imaging, or visual or endoscopic examination. Each of these modalities provides opportunities for the integration