From benchside to bedside and back again:

Evidence-based content drives the future of Precision Medicine

by Helena F. Deus*, Katie Scranton*, Rui Fa, Finlay Maclean, Thom Pijnenburg, Jabe Wilson and Mevan Samarasinghe

**These first two co-authors contributed equally to this paper*

The size and breadth of data available across the health care fields has led to vivid discussions around the possibilities of applying artificial intelligence (AI) to support decision making in clinical settings. While AI can perform incredibly well on discrete tasks, attempts to apply it to larger, more complex problems in medicine have fallen short. In this article, we propose that such challenges should instead be tackled with intelligence augmentation (IA), an approach that integrates AI methods with human expertise. Of paramount importance in this approach is that an AI-based system must be integrated within an existing physician workflow and empower the physician to be the central decision maker. In this application of intelligence augmentation, practicing precision medicine would not require a drastic shift in a physician's workflow, but instead would use modern, efficient alternatives to traditional clinical decision support.

Introduction

The most trusted and accurate clinical decision support system has historically been the network of physicians themselves. Physicians and other clinicians commonly consult colleagues who, through years of experience or specialized training, have developed an intuition for identifying the patterns that lead to the correct diagnosis or the correct treatment. The resources that physicians most commonly use such as textbooks and journal articles can also be used to develop reference clinical decision support tools. In fact, the tools that are most widely used today are created and curated by physicians (e.g. UpToDate, Dynamed, StatDx, ViaOncology, among others). These tools may be attractive because of the speed at which they operate, but they are trusted because of the wealth of experience, evidence and knowledge they bring in support of precision medicine.

Tools that rely on manual curation for content preparation, however, may be limited by the speed at which clinicians can find and extract relevant content from reference and research material. As long as preparing such content remains a manual endeavor, it is unlikely that clinical decision support tools will be able to scale to cover a wide array of specialties or incorporate recent breakthroughs. Another weakness is the inability to deliver actively the right content at the right time to the right person. These tools are references that may sort or filter content, but are inherently passive, requiring clinicians to enter queries or keywords and to know the sources to access for the right information.

There is an opportunity to address these challenges and vastly improve existing methods for clinical decision support by leveraging modern artificial intelligence (AI) methods,[1] but many previous implementations of AI tools have not earned the trust of physicians. Our vision for the next step for precision medicine is to integrate the *knowledge and experience* that forms the basis of trust with clinical decision AI tools that solve the problems of *scalability and delivery*. Our approach to this challenging set of problems is to employ intelligence augmentation (IA), the integration of AI tools to supplement and enhance, not replace, a human's decision-making process.

Big Data and Opportunities for AI in Health Care

A revolution in the practice of medicine is underway as large volumes of health care data become available and our capability to perform computational inference, statistical analysis, and machine learning on big data increases. This explosion in the availability of data across health care fields includes patient data (e.g., family history, EHR),

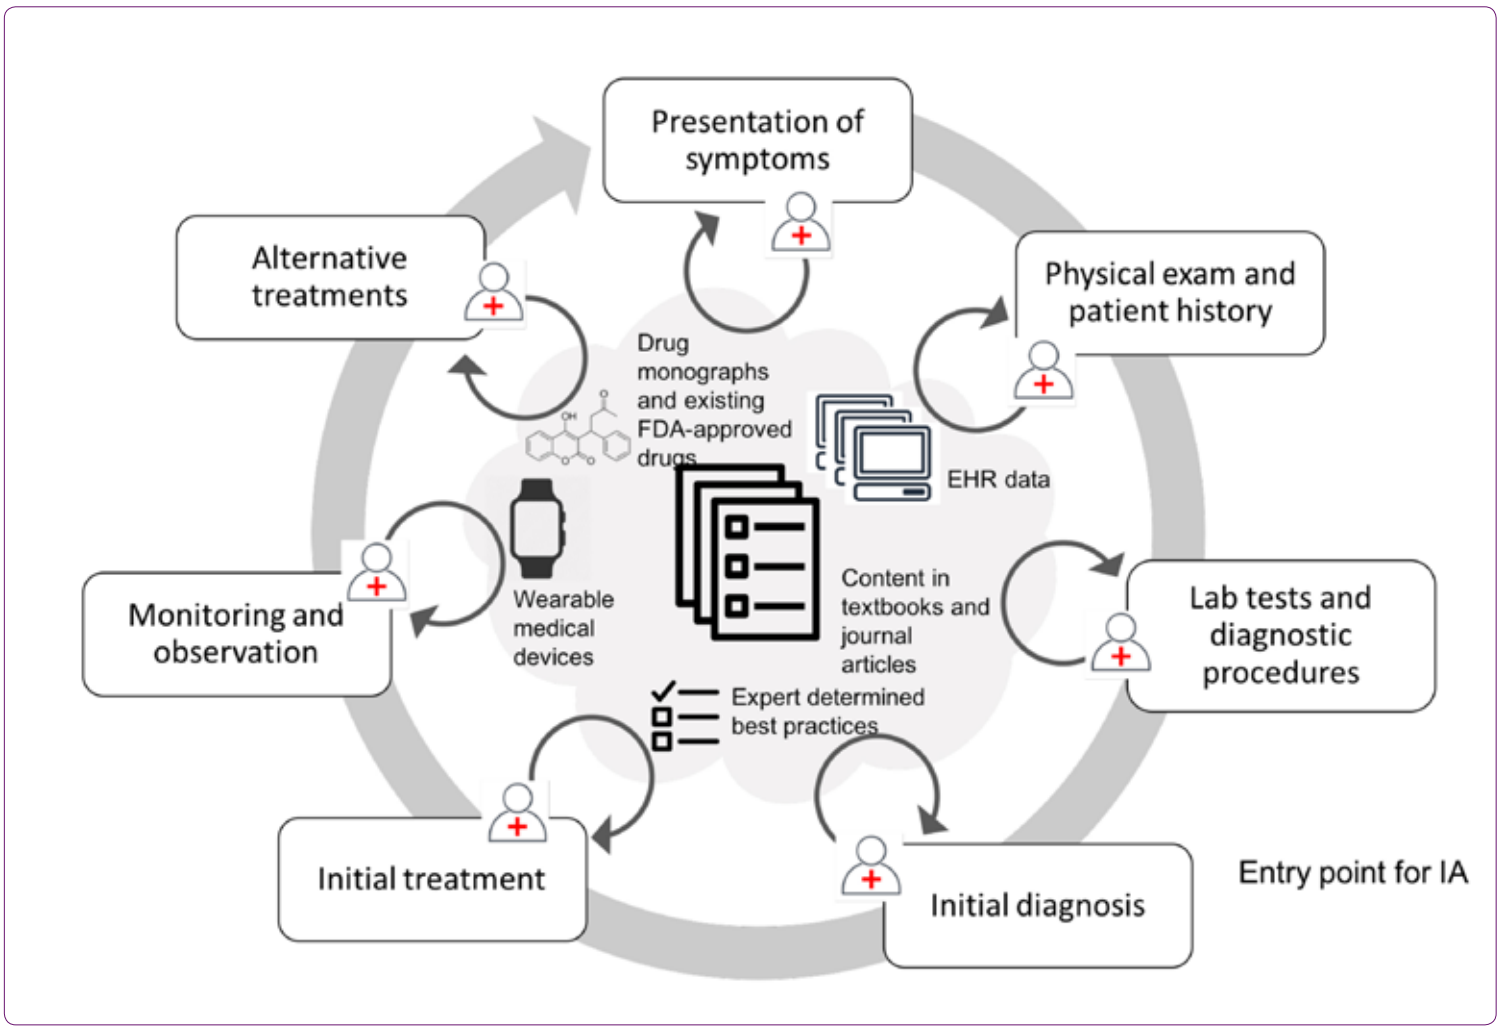

Figure 1: An illustration of the patient's journey towards diagnosis and treatment where, at each step a clinician incorporates information such as evidence-based content from textbooks and journals, expert- determined best practices, a patient's history from EHR data, observations from a procedure, clinical test or a wearable medical device, and drug monographs or FDA labels. The thin black arrows illustrate the opportunities for AI-based tools to make incorporation of knowledge more efficient and more complete. The entry point for IA is the initial diagnosis, which will lead to further testing and data collection in order to confirm the diagnosis and initiate treatment.

pharmacological data, and even medical knowledge and content.

Medical information is projected to double every 73 days in 2020 whereas in 2010 it was doubling only every 3.5 years.[2] Health care data is projected to increase in volume by 48% annually: from 153 exabyte (1 exabyte = 1 billion gigabytes) produced in 2013 to 2,314 exabytes in 2020.[3] At least 6 consumer physiological and emotional monitoring devices have been approved by the FDA as medical devices, including the Apple Watch EKG.[4] A reported 89% of health care providers already use or will use predictive analytics as part of the decision-making process within the next 5 years.[5] In fact, the size of the AI health market is poised to reach $6.6 billion USD by 2021.[6]

Most questions solved by AI and deep learning techniques are relatively straightforward and simple questions – that is, all information needed to solve the problem is present and at-hand. For example, a classic deep learning task is to train a model to recognize handwritten digits. Although it may require tens of thousands of samples of handwriting to learn the difference between a 1 and a 7, we know that we can provide the information needed to learn to perform the task. Similarly, deep learning can be incredibly efficient at solving discrete tasks in health care such as detection of lung nodules in chest CT radiology images, assuming all the information needed for that classification is available.[7] Clinical trials of new treatments can be made more efficient by recruiting and stratifying patients according to the presence of specific mutations for which the candidates are actively tested.[8,9]

Other problems in health care are significantly harder because the information used to answer the question might be vague, scattered across multiple sources, very hard to quantify, or simply missing or incomplete. For example, in diagnosing a patient the physician will likely incorporate information from the patient's medical history, a physical examination, answers to questions about their current illness, and tests ordered to confirm a suspected diagnosis. The physician builds up the full picture as they interact with the patient, relying on years of experience to filter out irrelevant information and probe for meaningful information. Obtaining this full picture for a patient is not a one step process but an iterative one, where the physician gathers new information based on what is already present, filling in the missing details.

Physicians and other clinicians leverage their medical knowledge and training when making these assessments, capturing data points but also making links and identifying patterns about the reasons why a combination of symptoms may be indicative of a diagnosis. This process may not end at a simple diagnosis but instead may be thought of as a loop (**Figure 1**). In the simplest case of this iterative process, the initial presentation of symptoms leads to identifying a diagnosis, a treatment, and an appropriate follow up schedule. A patient's journey continues as their treatment progresses and may change pending new data. At every step in this loop, the clinician is using information from a variety of sources: ongoing clinical tests, their own experiences, opinions of colleagues, textbooks, peer reviewed journal papers, guidelines, and drug monographs.

> *"Our approach to this challenging set of problems is to employ intelligence augmentation (IA), the integration of AI tools to supplement and enhance, not replace, a human's decision-making process"*

Artificial intelligence tools are not currently capable of filling in the missing details to answer these more complex problems in health care. Furthermore, a single technology product (AI powered or otherwise) or data set cannot address the clinicians' needs throughout their journey with the patient. For AI to make substantive contributions to precision medicine, multiple tools will need to be created for different steps in the process. Clinicians will then need to incorporate these tools into their existing workflow, to gain insight and information when they need it most.

The Possibilities of Intelligence Augmentation

Our approach starts with the assumption that clinicians need to be drivers – not passengers

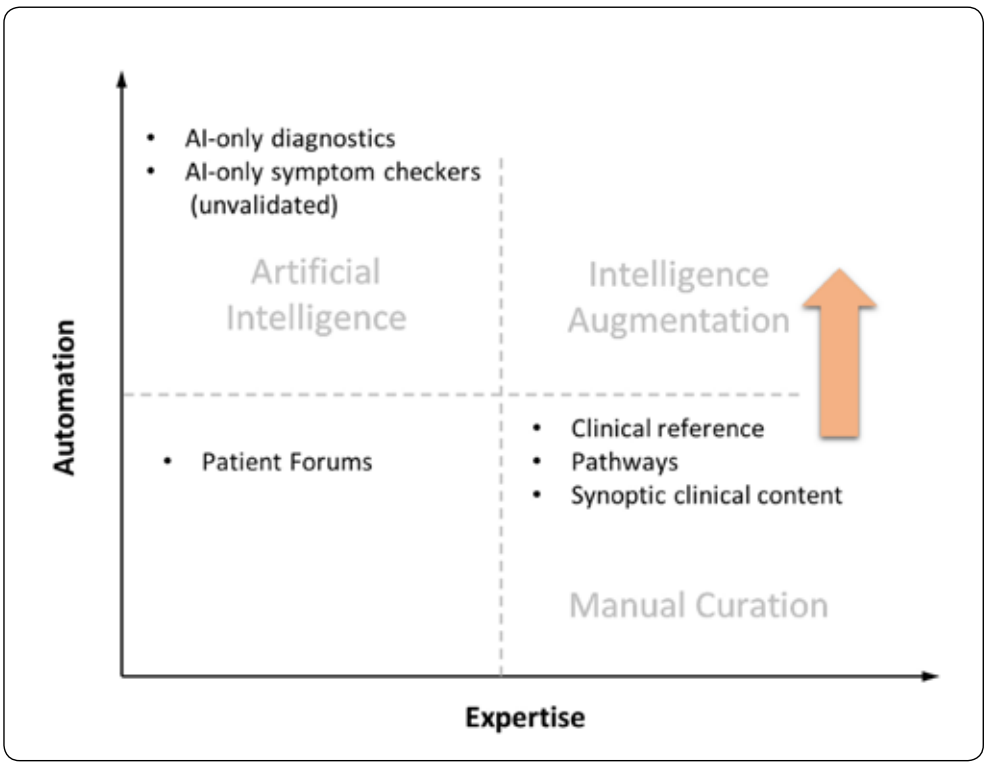

Figure 2: Today, the most trusted tools supporting clinical decisions are populated by content manually prepared by subject matter experts and therefore rely on very little or no automation (Manual Curation Quadrant). Some tools have attempted to automate this content preparation and recommendation by digesting all the information in literature and electronic health records but fallen short of physician trust expectations (Artificial Intelligence Quadrant). We propose that clinical decision tool builders shift to an approach of Intelligence Augmentation, where knowledge statistically extracted and inferred from highly trusted sources of content are integrated into the clinician's workflow at the point of care as an aid, not a replacement, for clinical decision – much like a GPS is an aid to navigation, not a replacement.

– of decision-support tools in precision medicine. As GPS navigation is useful in providing navigation information to the driver of a car, AI-based approaches should provide information that is relevant to guide the clinician in answering clinical questions instead of trying to answer the questions for them. This approach is known as intelligence augmentation (IA).[10] IA is not a separate field or technology from AI, but a different perspective on how best to employ AI, machine learning, and deep learning methods.[11]

For any successful clinical decision support tool, integration of information at each step of the patient care continuum is key. Today, relevant information mostly lives in different databases and file systems. The endorsement of HL7® Fast Healthcare Interoperability Resources (FHIR®) by the U.S. Health IT Policy and the Health IT Standards committees[12] will facilitate the integration of data stored in the electronic health records (EHR), in consumer genomic tests, and in wearable smart devices. Clinicians interpret a patient's data in the context of their knowledge, training, and experience in practicing medicine. Therefore, a complete IA approach requires tools based on collective expert knowledge and evidence-based content provided quickly and efficiently. Regardless of the data source, clinical decision support tools are built with one of two basic approaches: AI or manual curation. IA lives at the intersection of these two approaches (**Figure 2**). Below we describe IA systems for several use cases in precision medicine.

Experiments and Use Cases

A Knowledge Graph for Health

A significant amount of highly valuable and relevant knowledge is contained in textbooks and articles, not readily available at the point of care. About 1.3M new peer-review medical and biological scientific articles are indexed in PubMed every year,[13] with nearly 29 million publications indexed in 2018. Searching and filtering for relevant articles is a time-consuming task and is not feasible for a physician or nurse interviewing a patient at the point-of-care. Evidence such as facts from textbooks and conclusions from scientific studies need to be extracted from text and integrated with the rest of the patient's health data.

A knowledge graph is a data platform that represents evidence and conclusions as a web of relationships between events or concepts. It emphasizes the relationship between two medical concepts such as diseases or symptoms. For example, a bit of medical knowledge like "Asthma has a clinical finding of wheezing" is a relationship between two concepts (asthma and wheezing). The relationship itself has a type ("has clinical finding") and pieces of metadata can be attached to the relationship (see **Figure 3**). Metadata could be the provenance of the knowledge (the source from which the fact was extracted) or statistical evidence supporting that statement. Metadata could also be geographic specificity, seasonality, or the context for one of the concepts such as location of a symptom in the body, severity of a symptom, or aggravating factors. Relationships can also be used to allow the creation of disease hierarchies (e.g. "Asthma is subtype of chronic disease of the airways"), which are useful constructs for propagation and inference of new relations.

The value of knowledge graphs as platforms to store and retrieve medical knowledge has been widely demonstrated.[14,15] We evaluated one health knowledge graph for its ability to retrieve medically relevant information. To do this, we investigated two methods for answering clinical questions: four common questions from Ely's taxonomy[16] and additional disease specific questions. For six cardiac diseases, we attempted to answer these questions using standard clinical overviews (medically appraised synoptic content that has been prepared to support point of care decision making). In parallel, we tested an alternative approach using a medical knowledge graph containing content retrieved from unstructured medical text. The standard clinical reference pages contained information to fully or partially answer 75% of the more general questions, but only 27% of the specific questions. When the health knowledge graph was used as an intermediary between the question and the content, 100% of the general questions and 96% of the specific questions could be easily answered, demonstrating the incredible increase in ease and access to medical knowledge.

Knowledge graphs not only make it possible to store and retrieve knowledge, they allow

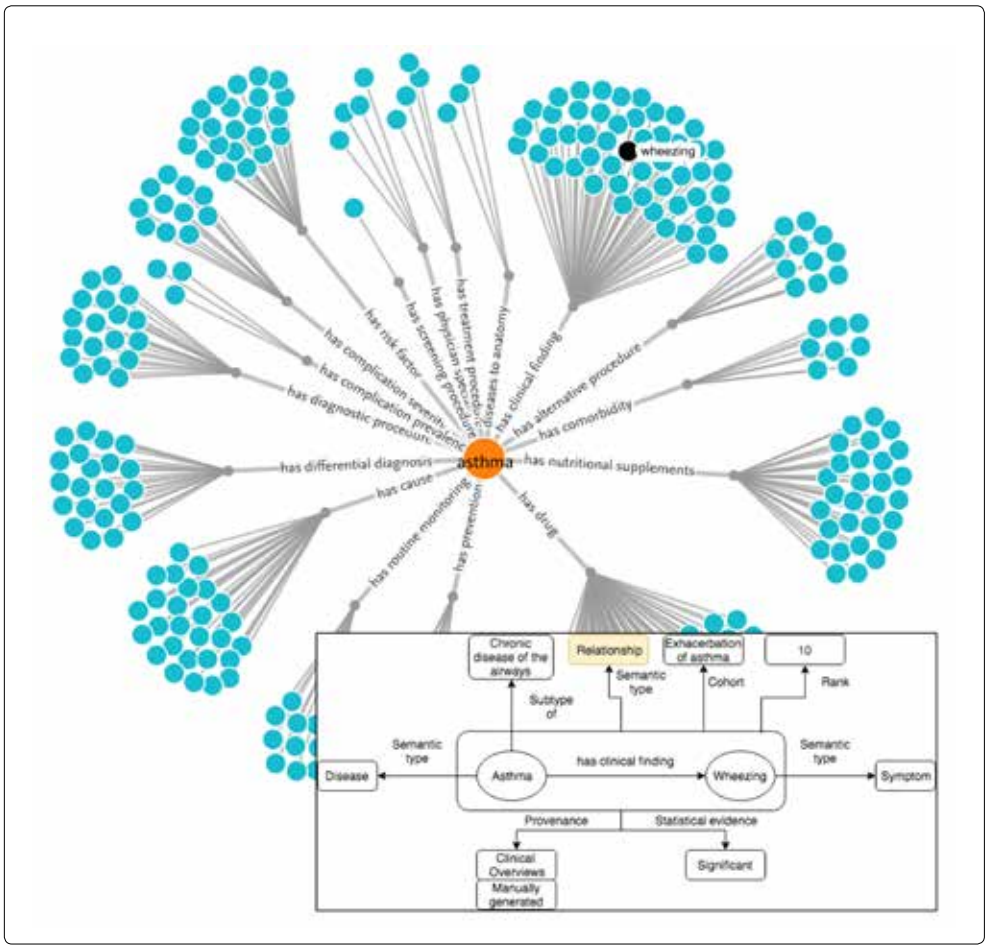

Figure 3: Example of the knowledge graph relationships associated with the concept "Asthma" and detailed inlet of metadata associated with each relationship. A given relationship can be expressed about a specific cohort, thus allowing for the specification of contextual information – for example, in a cohort of patients diagnosed with asthma, wheezing probably indicates an exacerbation of asthma with high confidence (rank 10 out of 10). Cohorts can be subgroups of patients defined by any variable or combination thereof (e.g. presence of a biomarker or group of biomarkers, recurrence of disease, previously treated with a certain drug, etc).

for new inferences and connections. Because a knowledge graph is a network of linked concepts it is possible to infer new knowledge by "walking the graph" or connecting previously unconnected concepts through a common link. A classic case is inferring that a relationship of a disease to a symptom would also be true for the diseases higher up in the hierarchy (e.g."Airway inflammation has clinical finding of wheezing"). Knowledge graphs built using standards can also be linked to other health care databases (e.g. UMLS, ICD10, SNOMED CT, RxNorm) that use the same standards, allowing for the integration of knowledge and data across disparate sources.

Medical Text Mining

Knowledge graphs are a powerful way of representing evidence, but they rely on a large complex web of existing relationships. Manually extracting those relationships from textbooks and journal articles would be effective at creating highly trusted content but would be incredibly costly and time consuming (**Figure 2** quadrant - manual curation). At the other end of the spectrum, artificial intelligence approaches offer the potential of automation and scalability but lacks the expert knowledge to determine what relationships are relevant (**Figure 2** quadrant - artificial intelligence). An IA approach to the problem uses manually curated content to train algorithms to retrieve the sentences and paragraphs from existing documents that are relevant (**Figure 2** quadrant - intelligence augmentation).

The first step to incorporating IA into the clinical decision-making process is to develop methods to extract actionable bits of knowledge from literature. For example, advanced analytic techniques are needed to extract content from cursive text and heterogeneous sources. The field of natural language processing (NLP) provides methods that use supervised or semi-supervised approaches to automate and scale the extraction of content, starting with a base of evidence manually curated from literature and physicians' notes. For example, named entity recognition analysis, part of speech analysis, and dependency parsing can be used as part of a semi-supervised machine learning pipeline to recognize and extract knowledge from sentences in medical documents that match grammatical structures automatically learned from human curated knowledge relationships. Combining knowledge curated by physicians with NLP techniques that automate extraction of knowledge from unstructured text allows us to extract exponentially more information from textbooks and articles than would be possible by hand.

We used a machine learning technique called a convolutional neural network (CNN) to learn grammatical patterns that indicate a "clinical finding" relationship. Named entities such as diseases or symptoms were identified in sentences using a dictionary of more than 400K words and phrases available in the health knowledge graph mentioned above. We filter these sentences to those that mentioned both a disease and a symptom, separated by a certain number of words. We compared those sentences to the existing relationships that have been expressed and evaluated by subject matter experts (SMEs) to label the extracted sentences as to whether the symptom and disease in the sentence had a real symptom-disease relationship. The CNN model was trained to learn the grammatical patterns of the sentences with the symptom-disease relationships, allowing us to identify new relationships from sentences with similar grammatical structures.

> *"Our approach starts with the assumption that clinicians need to be drivers - not passengers - of decision-support tools in precision medicine"*

Subsets of the new, learned relationships were then validated with the SMEs to ascertain the accuracy of the methods. In this way, manual curation by SMEs was combined with AI machine learning techniques to create a knowledge graph of medical knowledge that can be quickly and efficiently queried by clinicians and clinical decision support tools.

A Knowledge Graph for Life Sciences

Data bottlenecks that block the implementation of IA can occur at various points in the process. The health knowledge graph described in the previous section bypasses the bottleneck created by the inaccessibility of information in textual content. In pharmacological research and drug development, the high fail rate of candidate drugs means that only 1 in 5000 compounds identified as potential treatments (those that enter the clinical trial stage) make it to market. There is an opportunity to use AI methods to more efficiently filter candidate drugs, but high performing models require many different sources of data. The bottleneck of identifying potential therapies could be solved by linking knowledge graphs to create a data platform that integrates heterogeneous sources of data and perform data harmonization to create life science data sets prepared for analysis with AI methods.

Integrated knowledge graphs are needed to tackle the challenging task of drug repurposing to predict novel treatments. The power of drug repurposing (that is, applying existing therapies to treat rare and orphan diseases) has been well established.[17] Existing drugs on the market have known safety and interaction profiles, easing the burden of clinical trials for alternative indications. Repurposing a drug offers the promise of taking less time and being considerably less expensive than de novo drug development.[18] The use of drug repurposing has been recently extended to precision medicine, including *in vitro* approaches such as personalized drug screening[19] and *in silico*, or computationally based approaches.[20]

One such computational approach undertaken by Mission:Cure and CuresWithinReach used a life science data platform to highlight novel therapeutic drugs for Chronic Pancreatitis (CP). CP is a fibro-inflammatory disease caused by multiple genetic and environmental factors. The complexity of the genetic drivers and the diversity in allelic variations has obfuscated understanding of disease pathology and led to low success rates of medication therapy. A data platform was constructed from linked knowledge graphs to overcome such challenges, using both proprietary and open data sources encompassing disease mechanisms, gene expression, phenotypic expression, proteomics, small molecule-protein interactomes, and chemical and drug physicochemical properties. Machine learning methods were used to highlight the relationships between CP-related phenotypes and FDA-approved drugs. Combining the manually curated data sets with ML methods allowed this IA system to identify several novel repurposing candidates.

The novel repurposing candidates were then validated, resulting in a list of novel drugs for treatment, predicted to bind to important gene-coded proteins. In addition, a number of drugs were identified for pre-clinical testing that were predicted to either bind to or modify multiple genes important in the pathogenesis of the disease. Through automated data integration and harmonization, IA systems such

as this knowledge-based data platform could save valuable time in the drug discovery process.

Precision Medicine in Practice

One of the final challenges of implementing Precision Medicine in health care is transferring knowledge from benchside to bedside, delivering evidence-based content quickly and efficiently at the point of care. The heterogeneity of the data available requires an arsenal of professionals such as research geneticists, genetic counsellors, data scientists, pharmacologists, and systems biologists, all working to gain insight and knowledge from disparate sources of data. According to Dr. Mark Haupt of Ariel Precision Medicine, "success in [precision medicine] is translating the amassed knowledge and information into a clinical context to clinicians and care providers at the point of care". The aim is simply not to report on most effective treatments - the field of Precision Therapeutics - the aim is "to use individual biology rather than population biology at every stage of a patient's journey."

The methodology of Ariel Precision Medicine provides an example of successful integration of Precision Medicine into clinical practice. Ariel developed mechanistic models of complex chronic diseases and disorders, based on animal models, single-cell, genomic and transcriptomic data, amongst other data sources. The employment of Next Generation Sequencing (NGS) genetic tests provide a greater understanding of gene variants in more than 700 gene targets. Alongside additional patient data, these insights are used at every stage of complex disease progression – screening, monitoring, diagnosis and formation of treatment plans – from determining biomarkers to monitor disease progression to determine therapeutic efficacy of treatments. Ariel also actively "[encourages] patients to own the data"; overcoming interoperability issues and aiding the integration with both the healthcare provider and patient. The outcome is greater patient involvement, greater biological understanding, and ultimately greater clinical information to assist decision-making at the point of care.

Future directions in Precision Medicine

Next-generation clinical decision support will need to combine knowledge curated by subject matter experts with the power of AI-driven analysis in order to utilize big data in health care. This vision of the future of precision medicine is based on an intelligence augmentation (IA) approach that incorporates artificial intelligence (AI) within a human-centered workflow.

This IA approach to precision medicine is not without its challenges. Machine text comprehension methods that are needed to extract knowledge relationships from medical text are not as advanced as natural language processing methods, hindering approaches to using medical content. Security and privacy issues will be paramount for tools that, by definition, need to acquire and combine sensitive medical information from many different data sources. Strategies for regulating, evaluating, and validating the clinical safety of tools that use predictive analytics will have to keep up with the breakneck pace of analytical development. These will also create a need to re-train and certify clinicians on proper use of IA systems – although many users would likely prefer self-taught programs, some users will require formal training. Furthermore, user experience should be an integral part of the development and design of IA tools. Those tools should be not only intuitive to use but also protect against user error and accidental misuse.

> *"An IA approach to the problem uses manually curated content to train algorithms to retrieve the sentences and paragraphs from existing documents that are relevant"*

Patient-facing tools that return information to patients who have access to their own data but lack the medical expertise to interpret it will also need to be built with extra caution and regulation. The Achilles' heel of AI-tools is the inability to realize that many patients may not have the full picture and need to compensate for missing information. For these patients, finding and incorporating all medically relevant data will be crucial. Even for capable patients armed with this information, the focus needs to remain on the physician as the driver of the process, using AI-based clinical decision support tools to supplement and enhance their decision making.

Despite these challenges, IA systems that incorporate evidence-based content have the potential to provide insight into very challenging medical problems. The identification of candidate compounds that may be effective therapies relies on information in patents, scientific articles, and books. Identifying groups of similar patients for retrospective analysis or for recruitment into clinical trials requires knowledge from text in order to distinguish clearly between events that co-occur and an event that was caused by a previous event (e.g. broken bone and glioblastoma diagnosis versus treatment with Afatinib and reduction of tumor size). Extending the grouping of patients by considering their full longitudinal records would allow us to generalize from individual patient journeys to patient pathways. Combining rule-based and machine learning methods to analyze patient histories based on outcomes could identify ideal treatment pathways.

The potential value of AI methods that provide a differential diagnosis has been widely recognized.[21,22] However, an IA based approach could implement many different sub-tasks, intuitively following a human physician's diagnostic procedure. AI tools could each perform a discrete task to analyze a fraction of the larger picture such as synthesize a patient's entire EHR history, make predictions based on lab test results, and screen radiology images. In fact a number of studies have implemented predictive models using demographic factors and retrospective medical history,[23] automatic disease coding from clinical notes,[24] deep learning models using medical images,[25] and modelling of longitudinal clinical events in EHR.[26] Integrating tools that perform discrete tasks well with evidence based content in a knowledge graph would improve the accuracy of a differential diagnostic tool by filling in a larger part of the picture one tool at a time.

Though our focus has been on IA systems that bring information from benchside to bedside, a powerful possibility for IA systems is the ability to create a two-way street that brings information gleaned from real patient experiences back to researchers (also called reverse translational medicine). IA systems that bring insights from bedside to benchside open the possibility of new innovations in precision medicine. Data from patients who are diagnosed and treated can be linked with heterogenous sources of health care data, e.g., their other EHR records, data from wearable devices, their genomic data. With access to the millions of real-time observational studies occurring every day in patient care, clinical researchers would have a new avenue to study biomarker

identification, drug discovery, treatment decisions, among others.

Our thesis is that advancements in IA systems require the clinician to be in the driver's seat. AI tools should supplement, not substitute, the health care practitioner. For example, when the physician recommends a pathway treatment plan and later decides to deviate from that plan, any recommendation tool should still offer step-by-step support on the effective treatment plan. Similarly, if a radiologist is using an image recognition algorithm to identify problematic regions, he/she should still be empowered to request an analysis in off-target regions of the image, which can be a source of secondary prognosis. A differential diagnostic tool should inform the physician of the data that factored into the recommendation and allow the physician to edit the input, creating an interactive exploratory tool instead of a drop-in replacement for diagnosis by a physician. In order to become part of the medical workflow, IA systems need to demonstrate their worth and build the trust with healthcare practitioners, ensuring that doctors and the health care teams making use of them are still compliant with their Hippocratic Oath: "First do no harm". JoPM

Acknowledgements:

Sima Reichenbach contributed to the initial discussions, helping to shape the direction of the paper. Dimitrios Alivantros contributed to the drug repurposing study in the knowledge graph for life science section. We wish to acknowledge Anna Perini, who devised and carried out the experiment evaluating the health knowledge graph and Craig Stanley, who developed Figure 2 and contributed with useful discussions. We would also like to acknowledge Will Dowling, who worked on the text mining use case.

Helena F Deus has a PhD in Bioinformatics from Lisbon, Portugal and has been involved in developing technologies for over 13 years that aid and accelerate translational research and precision medicine for oncology at the MD Anderson Cancer Center in Houston, Texas and at Foundation Medicine in Cambridge, Massachusetts. She now works as a technology research director at Elsevier, seeking to optimize the use of machine learning and deep learning techniques in supporting point of care advanced clinical decision making.

Katie Scranton has a PhD from University of California, Berkeley and is a Senior Data Scientist in Precision Medicine at Elsevier. Her work focuses on predictive modeling using patient data for applications in clinical decision support.

Rui Fa has a PhD from University of Newcastle, UK, in 2007 and has worked in bioinformatics and computational biology for more than 10 years in many UK top universities and research institutes, including The University of Liverpool, University College London (UCL) and The Francis Crick Institute. In February 2018, He joined Elsevier as a senior research scientist in precision medicine. His current research interests include healthcare data analytics, bioinformatics and computational biology, machine learning/deep learning and network science. Dr Fa has authored and co-authored more than 60 peer reviewed journal and conference papers, and one research monograph.

Finlay MacLean completed a Masters of Engineering in Biomedical Engineering at the University of Nottingham. He is currently working as a data scientist at Elsevier, primarily involved in drug repurposing. He is passionate about aiding understanding of rare diseases through the use of novel computational methods.

Thom Pijnenburg is machine learning specialist at Elsevier. He earned his master's degree cum laude in Theoretical Physics at the University of Amsterdam and was a graduate research trainee at McGill University specializing in High Energy Theoretical Physics. With a fascination for discovering underlying patterns, he now works on developing machine learning based solutions to Life Sciences problems.

Jabe Wilson is the Consulting Director for Text and Data Analytics within the Professional Services organisation at Elsevier. He has a background in Artificial Intelligence (AI), which he has been working with for almost 30 years, since studying at the School of Cognitive and Computing Sciences at Sussex University. He has broad experience in industries delivering digital content, and has participated in several start-ups as well as teaching interaction design at the Royal College of Art (RCA). Since the late 1990s he has worked on a wide range of innovative projects instrumental in driving the transformation of Elsevier from an established scientific print publisher through to its transition into a predominantly digital information analytics company. Projects include: Elsevier's first electronic reference books; early collaborative social networks in biomedical science with BioMedNet; the indexing for adverse drug reactions with the creation of Pharmapendium.com; development of early semantic knowledge bases for therapeutic drug design, introduction of adaptive learning courses into the National Health Service in the UK with Elsevier Clinical Skills; and, developing the semantic search product Elsevier Text Mining.

Mevan Samarasinghe is VP of Knowledge & Analytics Platforms at Elsevier Health. His group is responsible for organizing Elsevier's vast medical & scientific content and developing machine learning and content discovery platform services that enable applications for clinical decision support, life sciences and medical education. Prior to Elsevier, Mevan designed and implemented highly-scalable knowledge and content discovery services for multiple companies from startups to Fortune 50 organizations. Mevan holds a master's degree in computer science from Johns Hopkins University.

References

1. Bertalan Mesko. "The role of artificial intelligence in precision medicine". In: Expert Review of Precision Medicine and Drug Development 2.5 (Sept. 2017), pp. 239–241.
2. Peter Densen. "Challenges and Opportunities Facing Medical Education". In: Transactions of the American Clinical and Climatological Association 122 (Jan. 2010), pp. 48–58.
3. Stanford Medicine 2017 Health Trends Report: Harnessing the Power of Data in Health. Tech. rep. Stanford Medicine, 2017.
4. Mohammed Taj-Eldin et al. "A Review of Wearable Solutions for Physiological and Emotional Monitoring for Use by People with Autism Spectrum Disorder and Their Caregivers." In: Sensors (Basel, Switzerland) 18.12 (Dec. 2018).
5. Jennifer Bresnick. "93% of Payers, Providers Say Predictive Analytics are the Future". In: Health IT Analytics (June 2017).
6. Stanford Medicine 2018 Health Trends Report: The Democratization of Health Care. Tech. rep. Stanford Medicine, 2018.
7. Yaniv Bar et al. "Chest pathology detection using deep learning with non-medical training". In: (2015), pp. 294–297.
8. Donald A Berry. "The Brave New World of clinical cancer research: adaptive biomarker-driven trials integrating clinical practice with clinical research". In: Molecular oncology 9.5 (2015), pp. 951–959.
9. Michael E. Goldberg et al. "Multiple configurations of EGFR exon 20 resistance mutations after fi and third-generation EGFR TKI treatment affect treatment options in NSCLC". In: PLOS ONE 13.11 (Nov. 2018), pp. 1–8.
10. Aaron Masih. "Why businesses must focus on intelligence augmentation over artificial intel- ligence". In: CIO Dive (Sept. 2018).
11. Nan-ning Zheng et al. "Hybrid-augmented intelligence: collaboration and cognition". In: Fron- tiers of Information Technology & Electronic Engineering 18.2 (Feb. 2017), pp. 153–179.
12. Steven Posnack and Wes Barker. "Heat Wave: The U.S. is Poised to Catch FHIR in 2019". In: Health IT Buzz (Oct. 2018).
13. MEDLINE PubMed Production Statistics. Tech. rep. 2019.
14. Helena F Deus. "Big semantic Processing in the Life Sciences Domain." In: Encyclopedia of Big Data Technologies. 2018.
15. Jamal Atif et al. "From Generic Knowledge to Specific Reasoning for Medical Image Inter- pretation Using Graph based Representations." In: IJCAI. 2007, pp. 224–229.
16. John W Ely et al. "A taxonomy of generic clinical questions: classification study". In: Bmj 321.7258 (2000), pp. 429–432.
17. Ted T Ashburn and Karl B Thor. "Drug repositioning: Identifying and developing new uses for existing drugs". In: Nature Reviews Drug Discovery 3.8 (2004), p. 673.
18. Hongming Chen et al. "The rise of deep learning in drug discovery". In: Drug Discovery Today 23.6 (2018), pp. 1241–1250.
19. Kirill Gorshkov et al. "Advancing precision medicine with personalized drug screening". In: Drug Discovery Today 24.1 (2019), pp. 272–278.
20. Eric March-Vila et al. "On the integration of in silico drug design methods for drug repur- posing". In: Frontiers in pharmacology 8 (2017), p. 298.
21. Peter B Jensen, Lars J Jensen, and Søren Brunak. "Mining electronic health records: towards better research applications and clinical care". In: Nature Reviews Genetics 13.6 (2012), p. 395.
22. Cao Xiao, Edward Choi, and Jimeng Sun. "Opportunities and challenges in developing deep learning models using electronic health records data: a systematic review". In: Journal of the American Medical Informatics Association 25.10 (2018), pp. 1419–1428.
23. Jionglin Wu, Jason Roy, and Walter F Stewart. "Prediction modeling using EHR data: chal- lenges, strategies, and a comparison of machine learning approaches". In: Medical care (2010), S106–S113.
24. James Mullenbach et al. "Explainable prediction of medical codes from clinical text". In: arXiv preprint arXiv:1802.05695 (2018).
25. Jeffrey De Fauw et al. "Clinically applicable deep learning for diagnosis and referral in retinal disease". In: Nature medicine 24.9 (2018), p. 1342.
26. Alvin Rajkomar et al. "Scalable and accurate deep learning with electronic health records". In: NPJ Digital Medicine 1.1 (2018), p. 18.