HEALTHCARE, ARTIFICIAL INTELLIGENCE, DATA AND ETHICS – A 2030 VISION

How responsible innovation can lead to a healthier society

December 2018

Microsoft Rue Montoyer 51, Brussels, 1000, Belgium

Table of Contents

Foreword I

AI for Good

Technology is never neutral. We all know this. But sometimes we forget that the non-neutrality of technology comes with robust orientations. One may use a bayonet to cut some bread, but it is designed to kill a human being. A scalpel may be used as a weapon, but it is designed to alleviate human suffering and save lives. Digital technologies are biased like scalpels not like bayonets. Computers, Internet, the Web, Big Data, Cloud Computing, smart applications of all kinds, the Internet of Things, Artificial Intelligence … these are all great human developments that, by and large, have an intrinsic tendency to improve our lives, do things instead of us and better than us, freeing our time and capacities, and enable us to do more with less, or indeed achieve things otherwise impossible. This is particularly true in health care, where technological feasibility and ethical expectations can join forces productively, to achieve unprecedented levels of reach, in terms of population, and of tailoring, in terms of individualised care. So, I am delighted to see these crucial topics addressed in the following pages. Personally, I am convinced that the effort to inject an ethical vision into what can, in its absence, become a mere technological push for market uptake, is good for society and good for entrepreneurship. It does take a more long-term and insightful approach to engage with ethical and inclusive innovation. But compliance, or focusing on what may be done, is merely necessary yet insufficient, when compared to what more should be done over and above the legal requirements. The good news is that an ethical approach to digital innovation, especially in the health sector and when it comes to data management, is a collaborative enterprise. All stakeholders can help, not just to ensure that their voices are heard, but also to leverage all intelligences to design the right solutions. In all this, Europe, with its attention for human dignity and its care for human flourishing, can and should lead by example, supporting a vision and development of technology that is socially good and environmentally sustainable. Thus, this broad overview on health data ethics is timely, as the European election cycle begins. I hope that, in response, the political process will reaffirm Europe's mandate to be a global thought-leader for innovation, health, and well-being. The present conversation is crucial and needs to be supported. I hope it will join similar conversations in other, related areas, dealing with similar challenges, especially in the overlapping field of artificial intelligence.

Luciano Floridi
Professor of Philosophy and Ethics of Information
Director, Digital Ethics Lab
University of Oxford

Foreword II

Empowering Intelligent Health: unlocking the innovation potential in health data

Healthcare systems globally have been undergoing a profound digital transformation.

With that has come the creation of a wealth of data that has significant potential to help identify diseases earlier, create and improve treatments and improve the lives of patients across the globe. Unfortunately, even with advances in data protection and governance, health data is not easily accessible by the researchers, patients and doctors when they need it to help realize better outcomes.

Patient-data have been locked away in numerous silos, limiting the ability to combine data and leverage it to drive innovation. The causes are partly technical, with divergent systems holding data in formats that are not easily used by other systems, and partially based on out-dated laws and policies. And there are significant privacy and trust issues that we will need to overcome before we can effectively leverage large ecosystems of data for broader uses and drive healthcare benefits for us all. Without a proper trust foundation and painting a clearer portrait of the significant benefits that broader use of patient data is already delivering, we run the risk of missing a tremendous opportunity. That opportunity, to leverage significant advances in machine learning capabilities and the inexpensive yet massive compute power available from modern cloud computing platforms, brings healthcare to the forefront of discussions around applications of Artificial Intelligence (AI) to some of our most pressing challenges.

The use of AI in the healthcare context is already raising a series of important societal and ethical questions which we will need to address now, to ensure that Intelligent Health can deliver on its promise, respect existing norms and more importantly, helping us develop norms for some new issues that are starting to emerge. At Microsoft, we are confident that these new technological developments can be harnessed for social good, to deliver unprecedented improvements in many aspects of our healthcare.

But we also understand our obligation to play a role in the important conversations that must take place if we are to balance new opportunities with established and emerging social norms and regulatory frameworks. By working side-by-side with the healthcare industry's most pioneering players, we have the opportunity to advance this goal. Our mission at Microsoft is to empower every person and organization to achieve more, and with that in mind, our ambition is that health organizations can harness the benefits of AI to unlock biological insight and break data from silos for a truly personal understanding of human health and in turn, make Intelligent Health possible. This is how responsible innovation can lead to a healthier society, enable better access to care, lower costs and improved outcomes.

John Frank
Vice-President,
EU Government Affairs, Microsoft

Neil Jordan
Worldwide General Manager,
Health Industry, Microsoft

Introduction

The intersection between technology and health has been an increasing area of focus for policymakers, patient groups, ethicists and innovators. As a company, we found ourselves in the midst of many different discussions with customers in both the private and public sectors, seeking to harness technology, including cloud computing and AI, all for the end goal of improving human health. Many customers were struggling with the same questions, among them how to be responsible data stewards, how to design tools that advanced social good in ethical ways, and how to promote trust in their digital health-related products and services.

In November 2017, in an effort to facilitate greater sharing of the many questions and thoughtful responses we were hearing, we convened a year-long consultation with a series of discussions with stakeholders across Europe.[1] The goal was not to engage in an academic exercise. Instead, we hoped to exchange practical ideas to the design of a European framework that maximizes the benefits to all from the use of health data. We have also drawn from our collaboration and long-standing relation with the European Cloud in Health Advisory Council that was founded in 2015 to advocate for an environment which allows healthcare institutions- and patients- to reap the benefits of data-driven health care.[2]

We promised those who gave their time and energy to our discussions that the work would culminate in a single document. This is it: a synthesis of what we take away from those conversations.

[1] See acknowledgment at the end of this paper.
[2] Ibid.

Executive summary and recommendations

Intelligent Health - powered by advances in computing power, Artificial Intelligence (AI) and a rapidly expanding corpus of patient data - holds enormous potential to improve health care systems and patient health in Europe

The Commission's Communication on enabling the digital transformation of health and care in the Digital Single Market calls out the profound challenges that Europe's healthcare systems are facing and the need for new technologies and approaches that better leverage data [3]. By enabling the smart, efficient and safe use of patient data, AI-infused technologies are already transforming many aspects of contemporary healthcare across Europe, for the benefit of patients and the broader public.

As we spoke with stakeholders over the past year, we saw that many appreciate there are gaps between where we are now and the future healthcare system that effectively leverages data and AI.

What are those gaps and what will we need to do to overcome them? We found most of the gaps and challenges that stakeholders shared with us fit into three areas:

(1) Organizational and technical barriers to data sharing and data use;

(2) Insufficient public trust and lack of a regulatory framework that promotes more access to and use of patient data for research purposes, while addressing privacy and security concerns; and

(3) A lack of clear rules, or even a tentative discussion framework, governing the ethical and social implications of patient data, AI and its growing use in the field of healthcare.

Few conversations on the use of patient data happen without reference to the ***technical and organizational challenges that impede better sharing and use of such data.***

Most commentators agree, our healthcare data has historically been locked up in silos, spread across different provider organizations. Many of these organizations use different systems that make it hard for data to be aggregated to get a single view of a patient, let alone do larger "population health" scale views. Moreover, these diverse organizations holding patient data often lack incentives to overcome these technical and organization barriers to aggregating and using patient data.

These data silo challenges are only becoming more acute as patients themselves begin to amass their own repositories of health data from an array of new wellness and healthcare devices that are landing in the market. We are confident that a mix of adoption of market-driven technical standards and incentives that reward provider

[3] *Communication on enabling the digital transformation of health and care in the Digital Single Market; empowering citizens and building a healthier society,* available at http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=51628 (April 25, 2018).

stakeholders for sharing and aggregating data could go a long way in reducing current barriers to data sharing and use in the healthcare context.

We heard the largest number of comments, and some of the most pressing concerns, on the topic of ***privacy and security of patient data.*** All stakeholders agreed that while we must make sure that patient data is stored in a secure manner, we cannot let privacy or security be barriers to better use of an individual's data for his or her own diagnosis and care. And most agreed that we need to better consider the genuine altruism among patients and enable research uses that allow broad societal benefits from research use of patient data. From our stakeholder conversations, it became evident that we must do more to show the value of data sharing for research use, showing citizens how their data already is being used for great benefit. The GDPR provides a timely opportunity and frame for these secondary use discussions, offering Member States discretion to set flexible rules for research uses of patient data.[4] But equally importantly, we must ensure that regulatory frameworks governing the use of patient data enhance trust with citizens. To that end, it will be important to not only have discussions about how and under what circumstances patient data can be used, but also to have discussions that outline certain prohibited uses of patient data and identify methods that give patients enhanced controls. To do this, we will have to move beyond a narrow transactional view of data protection and delve into wider ethical discussion about sharing, aggregating and extracting insight from sensitive patient health information.

Of the three buckets of topics we heard in our discussions, by far the ones in the most embryonic stages of development are those related to ***the ethical and social implications of patient data, AI and its growing use in the field of healthcare.*** While AI-infused technologies are already delivering benefits to patients around the world, these early applications of AI are only the front edge of a potentially much larger wave of healthcare AI technologies. With the advent of these AI healthcare solutions, the question has become, how do we respond to the ethical and legal challenges they will undoubtedly create? Society only will obtain the public health benefits of AI-infused technologies if these systems are developed and deployed responsibly. During our stakeholder meetings, Microsoft released a set of principles to guide the responsible creation and deployment of AI. We believe that, to protect people and maintain trust in technology, the development of AI should be rooted in a commitment to 6 key principles of fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. These guiding principles were in fact also raised by stakeholders in our discussions, and importantly, they are not new to the health care sector. For these principles to be effective, however, they must be integrated into ongoing operations. We are addressing this in part through our AI

[4] The GDPR allows Member States to adopt laws that permit processing of personal data for scientific research purposes in ways that derogate from the GDPR (Article 89). For example, Member State law can restrict certain data subject rights, such as rights of access and objection (which can be difficult to fulfill in the case of ongoing research) so long as organizations implement appropriate safeguards when processing the data for scientific research purposes.

and Ethics in Engineering and Research (AETHER) Committee, which brings together senior leaders from across Microsoft to develop engineering best practices, tools and guidelines.

Summary of Policy Recommendations:

To address organizational and technical barriers to data sharing and data use:

(1) **Promote the use of open standards to better enable technical interoperability and explore opportunities to create greater incentives for data sharing across organizations.**

(2) **Enable new technical solutions such as blockchain to improve data provenance, health information exchange and collaboration.**

(3) **Continue EU funding in digital health solutions to enable exchange of health information, and data provenance, including for PROMs.**

To address insufficient public trust and the need for a regulatory framework that promotes more access to and use of patient data for research purposes, while addressing privacy and security concerns:

(4) **Analyze the implementation of research provisions under the GDPR in Member States, and where needed, amend laws or create more clarity through interpretations and guidance, to ensure innovative research projects don't die on the vine.**

(5) **Demonstrate the value of a 'data commons' and build confidence in all stakeholders through visibility of success stories where data sharing and technological innovation have improved health outcomes.**

(6) **Explore and promote new models for data donation that encourage patients to more easily enable their data to be used for beneficial research purposes.**

(7) **Invest in technical solutions, including through research funding, to enable secure machine learning with multiple data sources/systems.**

(8) **Support commonly used global standards for the controls in national certification schemes for handling of patient health information and promote GDPR harmonized EU-wide certifications and accreditation schemes.**

To address the lack of clear rules, or even a tentative discussion framework, governing the ethical and social implications of the growing use of AI and patient data in the field of healthcare:

(9) **Utilize emerging frameworks that will help ensure AI technologies are safe and reliable, promote fairness and inclusion and avoid bias, protect privacy and security, provide transparency and enable accountability.**

(10) **Invest in more research to explore and enhance methods that enable intelligibility of AI systems.**

(11) **Advance a common framework for documenting and explaining key characteristics of datasets.**

I. Connected Health

Technical and Organizational Barriers to Data Sharing and Use

I. Connected Health - Technical and Organizational Barriers to Data Sharing and Use

When we talk to stakeholders in the healthcare community about better leveraging patient data not only to provide care to an individual patient but also to make advances in population health, one of the first things we hear about are the challenges caused by patient health data being siloed in so many different places. While technical interoperability challenges related to moving data between different **Electronic Medical Record Systems** (EMRs) are usually one of the most frequently cited causes of these data silos, there are a range of other semantic and organizational barriers and blockers that we will also need to address if we are to capitalize on the full potential of applying modern technologies to patient health data.

The reality is that patient health data is stored in many distinct and diverse places. Whether it is data from a single patient that is scattered across different healthcare providers who have treated that patient, in large clinical research data sets that have been developed outside traditional provider relationships with patients or, increasingly, by patients themselves as they leverage new technologies that allow them to collect their own data or record information about outcomes (Patient Reported Outcome Measures or PROMs, as this data is known). These data sets are often in different computer systems in unique technical formats, collected for different purposes and thus semantically divergent, and nominally controlled by entities in different places in the healthcare continuum (often entities with potentially divergent interests).[5]

At a foundational technical level, there have indeed been technical interoperability challenges to aggregating and leveraging the data of an individual patient. To overcome these obstacles and to support effective data exchange, market-driven, consensus-based standards are critical to data driven healthcare and technologies. Healthcare developers are tasked with the challenge to bring diverse data sets together and develop machine learning across those data sets. We believe the best way to support developers working with health data is to offer tools that allow them to come together – for collaboration, creation, sharing, and building on each other's work. Significant progress is being made on this front in the form of a new consensus-based global standard named the **Fast Healthcare Interoperability Resources** (referred to as **FHIR** and pronounced «fire»). This important standard describes data formats and an application programming interface (API) for exchanging electronic health records, and importantly, it has now been embraced by a range of large EMR vendors[7] and others in the technical community, including all major cloud computing vendors.[8]

---

[5] See discussion in EU Green paper on mHealth (published Oct 4, 2014) available at https://ec.europa.eu/digital-single-market/en/news/green-paper-mobile-health-mhealth

[6] A recent policy document released by the UK Department of Health and Social Care includes a principle on use of open standards, recommending that stakeholders "Utilise and build into your product or innovation, current data and interoperability standards to ensure you can communicate easily with existing national systems." *Initial code of conduct for data-driven health and care technology,* available at https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology (Published 5 September 2018).

[7] Rahul Patel, Forbes, *Open Standards And Health Care Transformation: It's Finally Delivering On The Value It Promised,* available at https://www.forbes.com/sites/forbestechcouncil/2018/10/25/open-standards-and-health-care-transformation-its-finally-delivering-on-the-value-it-promised/#649a671615cf

[8] Microsoft, Amazon, Google, IBM, Oracle, and Salesforce issue joint statement for healthcare interoperability, available at https://cloudblogs.microsoft.com/industry-blog/industry/health/microsoft-amazon-google-and-ibm-issue-joint-statement-for-healthcare-interoperability/
See also https://cloudblogs.microsoft.com/industry-blog/industry/health/fhir-server-for-azure-an-open-source-project-for-cloud-based-health-solutions/

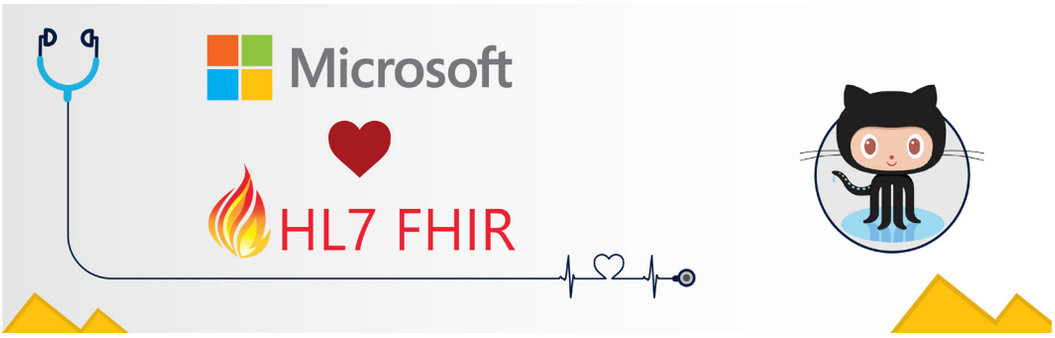

Fast Healthcare Interoperability Resources (FHIR) is rapidly gaining support in the healthcare community as the next generation standards framework for interoperability. Microsoft announced the release of FHIR Server for Azure, an open source project on GitHub. FHIR Server for Azure provides support infrastructure for immediate provisioning in the cloud, including mapping to Azure Active Directory (Azure AD), and the ability to enable role-based access controls (RBAC).

But EMR system interoperability is just the tip of the iceberg. As participants in our series of events called out, and as the EU mHealth Green Paper noted back in 2014, there are significant amounts of data outside of traditional EMR systems, increasingly data generated by patients themselves whether from mobile devices or through initiatives and tools aimed at gathering PROMs. Some participants noted a hesitancy for healthcare providers to incorporate this information into existing patient records systems or to use it their diagnosis and treatment decisions. There are important, valid issues that will need to be addressed for these data types to be mixed in with and considered alongside existing stores of patient data. Patient collected data has the potential to vastly augment existing data collected in the care setting, but we will need tools and frameworks to ensure that data collected by patients for their own personal health or wellness use is compatible with broader uses, whether by a care provider to treat that patient or in the context of broader population health research. Simply put, information collected in one context will not necessarily be useful or valid in another context.

For instance, questions around the accuracy and comparability of PROMs related to the ability of different patients to use similar criteria to report outcomes that are directly comparable have hindered broader use collection and use of PROMs. Clinicians have historically had questions around whether PROMs from a wide range of patients can be aggregated and compared in the same way that reported outcomes from them and other clinicians have typically been utilized. Organizations like the International Consortium for Health Outcome Measurement (ICHOM) have created frameworks that are starting to address these issues by working across stakeholder groups to create “Standard Sets” of outcomes that are most relevant to specific medical conditions. These Standard Sets, combined with recommendations from ICHOM about how to use various technologies and tools to facilitate capturing patient feedback offer the promise to create patient-reported data that is robust and involves a sampling rate far greater than anything clinician reporting of outcomes ever can deliver. These advances in collecting and using patient-generated data will also address another common theme we heard from stakeholders: the need to put the patient at the center of diagnosis and care.

Similarly, PROMs and other information collected by patients may be suitable for uses such as individual diagnosis and treatment, but certain stakeholders may view such data as less suitable for broader uses, such as for regulatory review and approval purposes. Regulators will have to grapple with these issues around accuracy and consistency of reporting before they will be comfortable relying more extensively on patient-generated data in making regulatory decisions.

Microsoft has joined forces with EFCCA-The European Federation of Crohn's & Ulcerative Colitis Associations, industry partners Takeda and SoftJam to launch Patient Voice, a patient-owned and led platform for sharing health data and enabling it to be used to report on outcomes (in line with the International Consortium for Health Outcomes Measurement standards). Hosted on Microsoft Azure, Patient Voice is designed to help patients report and measure their outcomes (PROMs) easily, effectively, and with confidence that their sensitive health data is secure. The platform will also give patients actionable insights related to this data, meaning they have more productive and effective conversations with their healthcare providers and systems.

The exchange of data across healthcare organizations is also a challenge the merits consideration here. Today, the healthcare industry suffers major inefficiencies due to diverse uncoordinated and unconnected data sources/systems. With digitized health data, the cross-organizational exchange of healthcare information is essential to support effective care collaboration. Traditional health information exchanges have had limited success. Blockchain offers new capabilities to greatly improve health information exchange. Blockchain holds great potential for healthcare consortiums to collaborate to improve the quality of care, lower costs, and improve the experience of patients and healthcare workers.

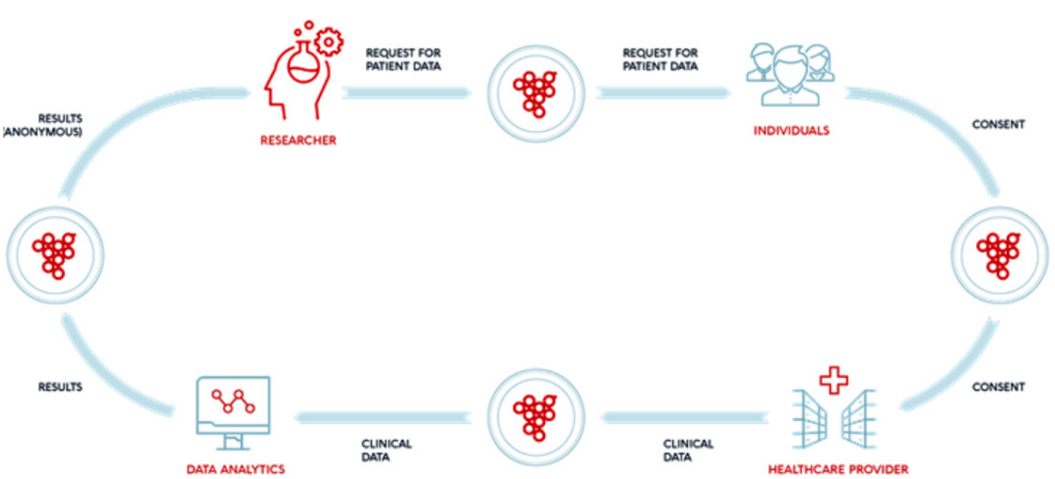

Finally, we would be remiss in not discussing government policy incentives which can spur activity to overcome the data silos from which the healthcare sector currently suffers.

The European Commission and Member States have identified and are working to tackle the shared and serious challenges currently facing Europe's health and care systems; policymakers have also recognized that data is a key enable for the transformation of these systems. There are active efforts underway to find solutions, including EU funding to support research and innovation in digital health solutions and improved infrastructure to enable the cross-border exchange of health information; the **eHealth network** to advance the interoperability of eHealth solutions; and public-private partnerships to promote innovation and strengthen interoperability.

Importantly, there is substantial overlap between the recommendations that emerged from our discussions and the EU's Digital Health and Care agenda, and in particular the EU goals of enabling citizens to securely access and share their health data across borders; to advance research, prevention and care through better use of data; and to empower individuals through better digital tools to manage health and care.[9] We look forward to continuing our work with EU institutions to achieve these shared objectives.

Few conversations on the use of patient data happen without reference to the technical and organizational challenges that impede better sharing and use of such data. That said, we are confident that a mix of adoption of market-driven technical standards and policymaker led incentives that reward various stakeholders for sharing, aggregating and using data could go a long way in reducing current barriers to data sharing and use in the healthcare context.

[9] The European Commission has set out its proposed plan of action in its April 2018 Communication on enabling the digital transformation of health and care in the Digital Single Market; empowering citizens and building a healthier society, available here: https://ec.europa.eu/digital-single-market/en/news/communication-enabling-digital-transformation-health-and-care-digital-single-market-empowering

II. Patient Trust & Data Commons:

Privacy, Security & Health Data Use

II. Patient Trust & Data Commons: Privacy, Security and Health Data Use

In our discussions with patients, researchers, technology innovators and regulators, the privacy and security of health data emerged as a constant theme. Stakeholders see challenges in balancing the clear need for privacy and security, and the associated patient trust, with regulatory frameworks that promote research and innovation to improve patient care through the use of an individual patient's health information.

Keeping healthcare data, and data more broadly, secure is paramount. In many healthcare scenarios, multiple parties would benefit from pooling their private datasets, training machine-learning models on the aggregate data, and sharing the benefits of using these models. This will only be achievable if the safeguards applied to that data are robust.

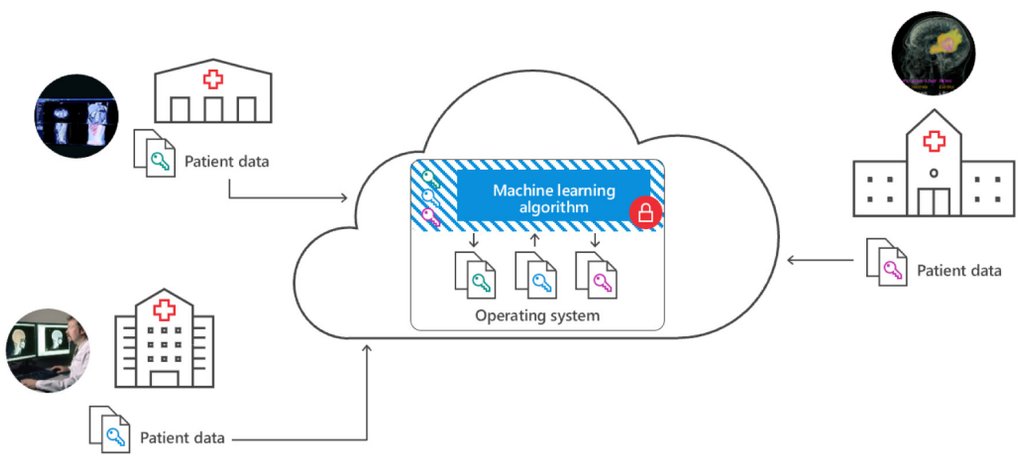

Multi-Party machine learning
Recent advances in hardware and software allow us to build a cloud service where multiple parties can now share data, but we can provide reassurance the analysis will run on the combined data, the learned model will be shared with each of the parties and no one will have access to data from P_1,P_2,..., P_n. The data will stay encrypted all the time.

Any discussion of data privacy and security in Europe (and, increasingly, beyond Europe) must begin with the **General Data Protection Regulation** (GDPR). The GDPR has set a high standard in the EU – and has emerged as a benchmark for third countries – for the protection of personal data, including health data. While the GDPR retains consent as a key control for data subjects over their health data, the GDPR provides Member States with important flexibilities to enable use of that health data *without consent*. In particular, the GDPR authorizes Member States to permit processing of health data without consent where it is necessary for scientific research purposes or for public interest in public health, subject to appropriate safeguards – facilitating "secondary processing" of health data (i.e. uses beyond those for the provision of primary care) in the Member States.

Member States have leveraged this margin of discretion. Ireland, for example, has adopted specialized regulations on health research, which include a broad definition of what activities falls into that category, and which impose a range of safeguards on health data processing, including prior approvals by research ethics committees and compulsory data protection training for researchers.[10] Similarly, German law allows for the use of health data (without consent) for scientific research following a balancing of interest test and subject to safeguards, such as encryption, training, and the appointment of a Data Protection Officer.[11] In Belgium, the national law was updated to lift (subject to safeguards) certain rights of individuals in their personal data in order to better balance the interests of individuals with the specific needs of scientific research.[12]

These developments are promising. More remains to be done at national level to improve the framework for secondary use of health data to promote public health, however. Earlier this year, the **European Cloud in Health Advisory Council** issued a GDPR-focused call to action to Member States, encouraging them to adopt interpretations of GDPR terminology such as "scientific research" and "adequate safeguards" that would stimulate research and enable public health officials, healthcare providers and patients to benefit from advances that will improve health outcomes and reduce the cost of care.[13]

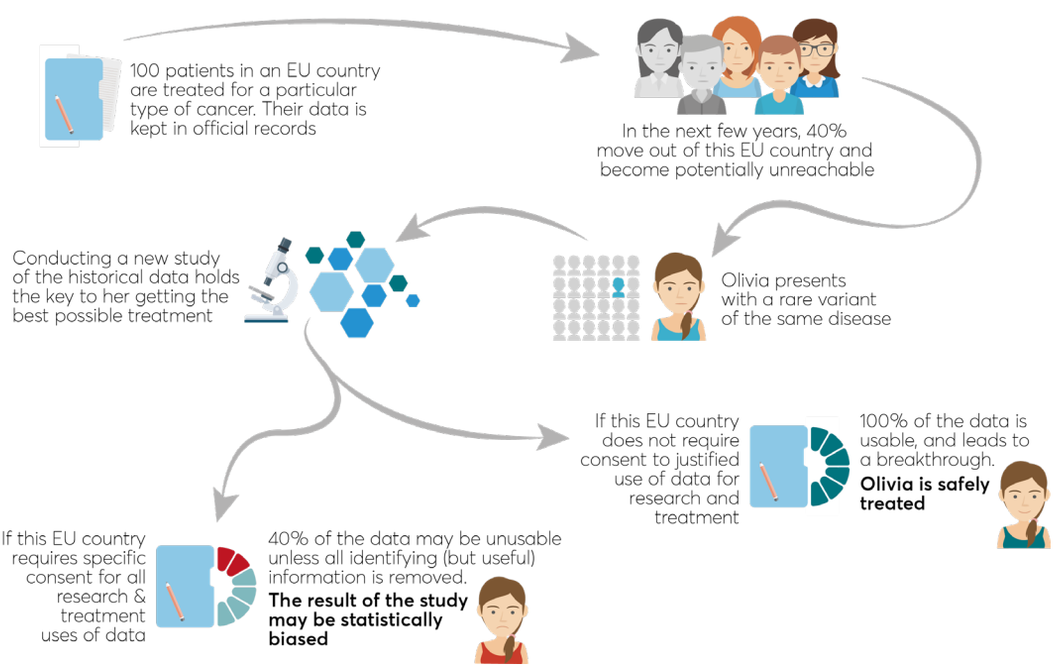

[10] *See* S.I. No. 314/2018 – Data Protection Act 2018 (Section 36(2)) (Health Research) Regulations 2018.
[11] *See* Bundesdatenschutzgesetz, Sections 27 and 22(2).
[12] *See* Wet betreffende de bescherming van natuurlijke personen met betrekking tot de verwerking van persoonsgegevens, Titel IV.
[13] EU Cloud in Health Advisory Council, Enable data-driven healthcare & research for citizen benefit while protecting patient privacy (Jan. 10, 2018), http://cloudinhealthadvisorycouncil.eu/papers/enable-data-driven-healthcare-research-for-citizen-benefit-while-protecting-patient-privacy/

We also heard in our discussions some expressions of confusion among compliance and legal teams regarding the current rules for secondary use of data and patient consent, leading to delays and bottlenecks in data sharing efforts. Similarly, researchers are struggling to determine whether anonymization of datasets can meet regulatory requirements while also retaining the necessary data of value to achieve research goals. Questions have also arisen on how to allocate responsibility for compliance with data protection law under GDPR concepts such as joint-controller, co-controller or processor. The mere existence of uncertainty in these areas, in particular in a risk adverse sector such as health, can reduce collaboration and data sharing. GDPR derogations for scientific research, specifically designed to achieve an acceptable balance between the rights of individuals and the needs of scientific research, risk being underutilized because of the lack of legal certainty and trust. Uncertainty related to basic concepts of data protection law, such as consent and the allocation of responsibility, can cause patient intake studies and similar initiatives to collapse.

To address this lack of clarity, stakeholders agreed that there is a need for guidance from regulators with respect to the applicable **"rules of the road"** for processing of health data. But at the same time, there is also a need for more public discussion so that patients better understand how, when and why their data will be used. To succeed, any effort to develop and implement new regulatory frameworks for research and secondary uses must go hand-in-hand with further clarification and explanation of – and enhanced trust in – those frameworks.

Our conversations with stakeholders also surfaced the need for a wider, ethics-based discussion about new models for use of health data – and for potentially supporting research projects that consider broadened concepts of "consent," as well as ideas around "data donation." Any such discussions will need to explore how these models can achieve the benefits of improved health outcomes through use of patient data, but also must be sensitive to the limitations on the use of such data to minimize risks to patients. There are a number of potential sources of inspiration here, among them the research project between the Digital Ethics Lab of the Oxford Internet institute, the Data Ethics Group at The Alan Turing Institute, and Microsoft, exploring the **"Ethics of Medical Data and Advanced Analytics."** Existing ethical frameworks, such as the **"Ethical Code for Posthumous Medical Data Donation"**, also provide groundwork for a discussion[14] (although the potential benefit of data donation likely lies in donation during one's lifetime, rather than after death – which raises additional considerations for patients, including that the data may be used in ways that compromise the patient's privacy or disadvantage the patient (e.g., to deny medical services).

More open dialogue about **"Data Commons"** is an important tool to promote trust. Greater transparency is also key to driving trust – transparency for patients, but also for healthcare providers, researchers and others innovating with health data. Under the GDPR, controllers that use third-party platforms (e.g., cloud computing services) must select technologies for hosting and sharing that data that meet robust privacy and security requirements. It can be difficult for "laypersons" without a technical background to understand what is required of them, or to meaningfully assess the privacy and security controls offered by third-party providers, however. To address this challenge in the health sector, some Member States are working to develop standards or accreditation processes that enable healthcare providers, researchers and others to more readily and easily assess the controls required for compliance.

[14] Enabling Posthumous Medical Data Donation: A Please for the Ethical Utilisation of Personal Health Data. Jenny Krutzinna, Mariarosaria Taddeo, Luciano Floridi, Oxford Internet Institute / The Alan Turing Institute.

The **NEN 7510 standard** in the Netherlands is one example; the standard, developed by the Dutch Standardization Institute, sets out controls for healthcare information security and serves as a benchmark against which organizations that process health information can assess compliance. The NEN standard borrows concepts and controls from global standards, in particular **ISO 27001**, adjusted to suit the sector. France has taken a slightly different approach, requiring service providers hosting certain types of health or medical data to be accredited. Previously, this required accreditation by the French Ministry of Health, but as of April of this year, service providers can be certified by an accredited body against criteria that again borrow from ISO 27001.[15] While a single EU wide certification may still be a way off in the future, the recent trend by Member States in looking to commonly used global standards for the controls in their national certification schemes for handling of patient health information is a positive one. Certification processes such as these help to simplify security and compliance for healthcare entities who want to use third-party technologies to unlock the potential of their data; increased harmonization in this context, both across Europe but also in relation to international standards, means assessments can be done more quickly. The GDPR now explicitly foresees harmonized EU-wide certifications, an option that should be explored and promoted by Member States – national laws should not stand in the way of utilizing the full potential of such certification / accreditation schemes.

On 6 November, Microsoft obtained the French HDS:2018 certification for hosting health data. French authorities have embraced a modern approach to information security that has its foundation in a set of well-established global standards, including ISO/IEC 27001, ISO/IEC 20000 and ISO/IEC 27018. The approach also includes a process which relies on existing third party private sector auditors, who must first be accredited by the Comité français d'accréditation (COFRAC) before undertaking the process to review and certify particular cloud services against the new requirements. Microsoft is the first hyperscale public cloud provider in France to obtain this certification following a documentary audit by the British Standards Institution (BSI) and an onsite audits of our Data Centers in the France Central and France South Regions. The HDS:2018 certification applies to all Azure Core Services available in our French Regions, all Office365 Core Services available in our French Regions and all Dynamics365 Core Services available in our French Regions.

.

[15] *See* Stephanie Faber, *France Issues New Rules for the Accreditation of Health Data Hosting Services Providers,* the National Law Review, available at https://www.natlawreview.com/article/france-issues-new-rules-accreditation-health-data-hosting-services-providers (May 3, 2018).

III. Ethical Considerations Related to AI in Healthcare

III. Ethical Considerations Related to AI in Healthcare

Artificial Intelligence offers incredible promise in the field of healthcare. From AI-infused tools that make more accurate analysis of various medical images to systems that look across large healthcare systems to determine how to more efficiently and promptly treat patients, the future certainly looks promising. But it is also critical to understand AI is already delivering innovation and benefits to healthcare today.

AI is an integral part of a new of tools that help clinicians more accurately detect and diagnose ailments such as diabetic retinopathy and cardiac irregularities. These diagnostic aids capitalize on highly advanced image-based machine learning capabilities that allow even non-specialist clinicians to detect and diagnose obscure disorders that they might previously have missed. Imagine the benefits that are possible when general practitioners have widespread access to tools that impart diagnostic capabilities such as ML models for detection of Diabetic Retinopathy[16] to find ailments that today only a small number of specialists can detect.[17]

But there are a host of other applications for AI in healthcare, including in systems that can better make sense of the millions of archived and real time data points that a modern medical facility generates and stores. A new artificial intelligence tool launched by Ochsner Health System enables doctors to do just that, by analyzing thousands of data points to predict which patients will deteriorate in the near future instead of waiting for patients to "code" and face a life threatening crisis.[18]

Although AI is already delivering clear benefits to the healthcare sector, it is clear that we have some critical work to do to ensure AI-infused healthcare technologies continue to be developed and deployed responsibly. We heard from stakeholders that beyond the data silo, trust and privacy challenges discussed above, there are important legal and ethical questions emerging around the development and use of AI technologies in the health sector. Achieving the public health benefits of AI-driven technologies will require us to work through these issues.

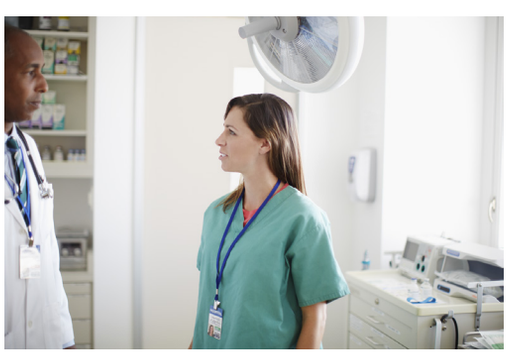

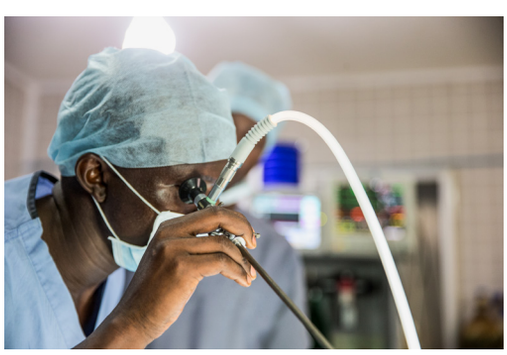

[16] https://blogs.technet.microsoft.com/machinelearning/2018/06/25/building-a-diabetic-retinopathy-prediction-application-using-azure-machine-learning/
[17] https://enterprise.microsoft.com/en-us/customer-story/industries/health/iris-increases-patient-engagement-help-microsoft-azure/
[18] https://news.microsoft.com/transform/ochsner-ai-prevents-cardiac-arrests-predicts-codes/

How to ethically and responsibly develop and deploy artificial intelligence technologies in health requires further discussion grounded in certain key principles, but our initial discussions suggest the following conclusions:

- **Trust must be built directly into the technology. Technologies must be reliable and safe.**
- **We must infuse technology with protections for privacy, and security.**
- **Recommendations or decision-making by artificial intelligence should be transparent to those relying on them for health decisions.**
- **Technology should be inclusive and respectful to everyone.**
- **Finally, since these are learning technologies, devices must be designed to detect new threats and devise appropriate protections as they evolve.**

During the course of our year long series of stakeholder meetings, Brad Smith, President and Chief Legal Officer, and Harry Shum, Executive Vice President of Microsoft AI and Research Group, released *The Future Computed: Artificial Intelligence and its role in society,*[19] which articulates principles to help guide the responsible creation and deployment of AI. The principles identified in the book overlap in a number of ways with those raised by stakeholders in our discussions. Repeated themes included the need to ensure AI technologies are safe and reliable, promote fairness and inclusion and avoid bias, and protect privacy and security. Underpinning these themes was a clear call for transparency and accountability in the use of AI technologies. These concepts are organized in the diagram below:

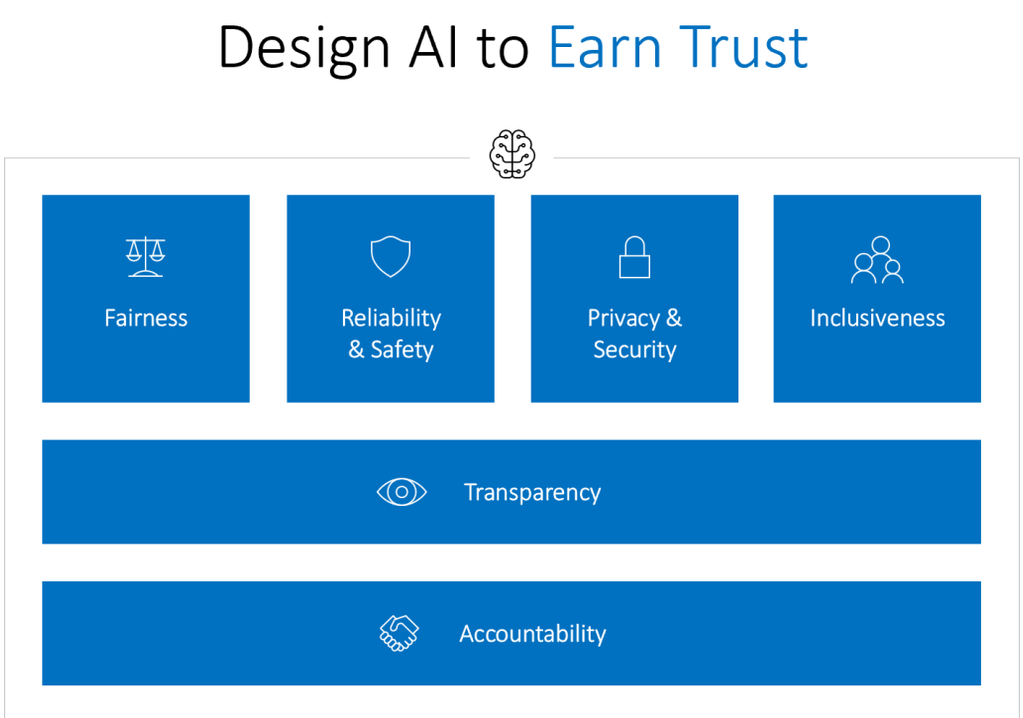

[19] https://blogs.microsoft.com/blog/2018/01/17/future-computed-artificial-intelligence-role-society/

There is wide consensus that these principles must serve as the foundation for the development and deployment of AI regardless of sector. Most recently, these principles have been endorsed by the data protection community in a Declaration on ethics and data protection in artificial intelligence, adopted in October 2018 during the 40th International Conference of Data Protection and Privacy Commissioners.[20]

Importantly for our stakeholder discussions, most of these principles are not new to the health care sector: ensuring that solutions are reliable, safe, representative, fair and secure have been core principles of the development of medical technologies for quite some time. The healthcare community therefore may well be better prepared to integrate these concepts into real world practice than other sectors.

One important new dimension (both in and outside of the health sector) is the increasingly **"black box"** nature of many AI technologies, however. Ensuring that users sufficiently understand these technologies to deploy them safely and effectively can be difficult. Because decisions made utilizing AI will impact patients' health and care, it is particularly important that everyone relying on these technologies understands how they interact with and on data, and their potential limitations. If users are not clear about the limitations of a technology or misunderstand the role of a technology in decision-making, use of the technology may create unfairness or negatively impact patient care. Adequate transparency therefore must involve transparency not only about how the AI system explains its results – teaching users what to expect from systems, how to interpret their results, and the extent to which they can be meaningfully relied upon in clinical decision-making is equally essential.

The GDPR addresses transparency of AI technologies in several ways, although at this point those references raise significant questions.[21] The GDPR restricts **"automated decision-making"** (machine-driven decision-making without any human intervention),[22] including when based on health data. Such decision-making is allowed only with consent of the individual concerned or on the basis of a Union or Member State law. The impact of this provision on the use of AI in healthcare is unclear. For example, what is the level of review by a clinician that is necessary to ensure that AI is no longer purely automated decision-making? The GDPR also provides that where automated decision-making is used, users should be notified and provided with "meaningful information" about the logic involved and the consequences and significance of such automated decision-making for the data subject.[23] This provision creates a framework for a conversation about what level of transparency is required for AI technologies in healthcare and how to implement that transparency, but again leaves open questions that will need to be worked through together, such as what constitutes "meaningful information" about the logic involved in AI-based decisions? How do you provide intelligibility?

[20] *See* https://icdppc.org/wp-content/uploads/2018/10/20180922_ICDPPC-40th_AI-Declaration_ADOPTED.pdf
[21] *See* Andrew Burt, *How will the GDPR impact machine learning?,* available at https://www.oreilly.com/ideas/how-will-the-gdpr-impact-machine-learning (May 16, 2018).
[22] *See* GDPR Article 22.
[23] GDPR Article 13(2)(f).

Intelligibility can improve the robustness of an AI system by making it easier to identify and fix bugs. It can also help users decide how much to trust an AI system, as well as uncover potential sources of bias or unfairness. Last but not least, intelligibility can help demonstrate compliance with regulatory obligations. A number of promising technical approaches to achieving intelligibility of both system components, including data and models, and entire systems have begun to emerge. Those that facilitate understanding of key characteristics of datasets, the overall behavior of trained models or entire systems, or why individual outputs or predictions were made are in particular relevant to the health care sector. Because intelligibility is a fundamentally human concept, more research is needed to understand which approaches do and do not help people achieve the end goals for which they need intelligibility. Early research involving human-subject experiments suggests that the landscape is not as straightforward as originally expected. Indeed, some system design choices commonly thought to influence intelligibility do not in fact have an appreciable effect on human understanding. Researchers have even shown that, at least in some contexts, literal exposure of model internals can prevent people from noticing when a model will make a mistake due to an overload of information.[24] Accordingly, as more methods for enabling human understanding of AI systems are developed and refined, it will be even more important to consider the full context in which a system is used and the reasons for needing intelligibility, as well as the utility of that method to real people, before selecting a particular method.

Beyond the GDPR, the ability of AI technologies to continuously learn and evolve raises more general questions about how these technologies are to be evaluated, and how determinations about safety and efficacy can be made as the technologies change over time in response to data fed in to them. While there is an existing regulatory regime for evaluating the safety and effectiveness of medical devices (including software), this regime is premised on a static product that can be reviewed at a point of time and will perform consistently after that point.[25] Any significant modifications to the product are incremental, and once made, are again static. The existing model for assessment is clearly in tension with the concept of a technology that continually learns by analyzing new data, modifying and improving the recommendations delivered to users. We will likely need to develop new frameworks for evaluating and ensuring the safety and reliability of AI technologies - frameworks that address the risk that dynamic changes introduce errors or bias that could negatively impact patient care and safety, while at the same time avoiding the need for near-constant validation and on-going regulatory review.

Wrapped up in both these challenges are concerns regarding **fairness, inclusiveness and bias.** Given the scale of the data and computational processing utilized in these technologies, often the algorithm may not be able to be fully explained, or fully reviewed by clinicians or regulators. Many of the current AI-based technologies in healthcare provide recommendations based on clinical images. In these cases, a clinician can review the same image, and make an independent determination as to whether the features identified by the technology are clinically relevant or not.

---

[24] *See* https://arxiv.org/abs/1802.07810

[24] Jonathan Kay, *How Do You Regulate a Self-Improving Algorithm?*, The Atlantic available at https://www.theatlantic.com/technology/archive/2017/10/algorithms-future-of-health-care/543825/ (October 25, 2018).

But as technologies develop to analyze incredibly large datasets and variables, it may not be possible for a clinician to understand the data analyzed by the technology, let alone understand how the data results in the recommendation offered by the technology. This will make it more difficult for the healthcare community to assess whether the technologies propagate biases or underrepresentation inherent in the underlying dataset, and whether the technology is clinically accurate in addition to being technically accurate. Helping people understand the relevant aspects of a dataset's characteristics and origins can help them better understand the behavior of models and systems involving that dataset.

A recent project initiated by a group of researchers at Microsoft seeks to advance a common framework for documenting and explaining key characteristics of datasets.[26] Called "datasheets for datasets," the project replicates the common practice in the electronics industry of accompanying every component, no matter how simple, with a datasheet detailing standard operating characteristics, test results, recommended usage, and other information. The datasheets for datasets project similarly recommends that every machine learning dataset be accompanied by a datasheet that describes and explains its motivations, its composition, how it was collected and pre-processed, and any limitations that could result in unintended outcomes, such as known biases or violations of privacy restrictions. Work has also started to develop similar datasheets for documenting critical information about models and systems.

Further discussion is needed to develop appropriate mechanisms for technology developers and users to identify whether a correlation identified by the technology represents a true causal clinical effect or whether other variables not assessed by the technology or in the dataset confound the finding. In light of these challenges, there is an emerging consensus that when AI is deployed for healthcare, it should augment the skills and experience of clinicians, rather than replace those skills – in effect, these tools should offer an "augmented intelligence" rather than true "artificial intelligence."

In these early days of AI technologies, there are numerous publicized examples of where AI technologies have not lived up to their promise due to failures in reliability, safety, fairness or inclusiveness, many in sectors other than healthcare. As we continue to have discussions regarding the responsible development and deployment of AI technologies in healthcare, it is also important to make more visible the stories of success. Broader discussion of successful AI technologies in healthcare will enable broader trust in these technologies by all stakeholders – patients, healthcare professionals, healthcare systems, and regulators – and enable more productive development of frameworks for the responsible creation and deployment of AI.

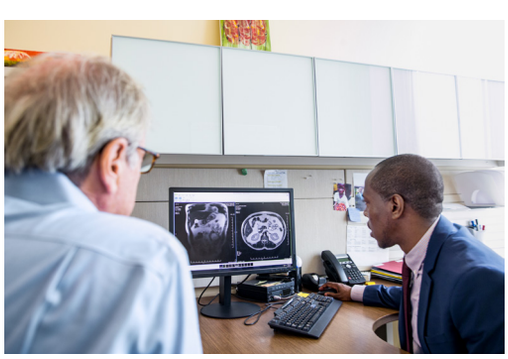

[26] *See* https://arxiv.org/abs/1803.09010

Conclusions and policy recommendations

As we spoke with stakeholders over the past year, we saw that many appreciate there are gaps between where we are now and the future healthcare system that effectively leverages data and AI. What are those gaps and what will we need to do to overcome them? We found most of the gaps and challenges that stakeholders shared with us fit into three buckets: (1) Organizational and technical barriers to data sharing and data use; (2) Insufficient public trust and lack of a regulatory framework that promotes more access to and use of patient data for research purposes, while addressing privacy and security concerns; and (3) A lack of clear rules, or even a tentative discussion framework, governing the ethical and social implications of patient data, AI and its growing use in the field of healthcare.

We distilled and formulated policy recommendations that reflect the discussions.

To address organizational and technical barriers to data sharing and data use:

(1) **Promote the use of open standards to better enable technical interoperability and explore opportunities to create greater incentives for data sharing across organizations.**

(2) **Enable new technical solutions such as blockchain to improve data provenance, health information exchange and collaboration.**

(3) **Continue EU funding in digital health solutions to enable exchange of health information, and data provenance, including for PROMs.**

To address insufficient public trust and the need for a regulatory framework that promotes more access to and use of patient data for research purposes, while addressing privacy and security concerns:

(4) **Analyze the implementation of research provisions under the GDPR in Member States, and where needed, amend laws or create more clarity through interpretations and guidance, to ensure innovative research projects don't die on the vine.**

(5) **Demonstrate the value of 'data commons' and build confidence in all stakeholders through visibility of success stories where data sharing and technological innovation have improved health outcomes.**

(6) **Explore and promote new models for data donation that encourage patients to more easily enable their data to be used for beneficial research purposes.**

(7) **Invest in technical solutions, including through research funding, to enable secure machine learning with multiple data sources/systems.**

(8) **Support commonly used global standards for the controls in national certification schemes for handling of patient health information and promote GDPR harmonized EU-wide certifications and accreditation schemes.**

To address the lack of clear rules, or even a tentative discussion framework, governing the ethical and social implications of patient data, AI and its growing use in the field of healthcare:

(9) **Utilize emerging frameworks that will help ensure AI technologies are safe and reliable, promote fairness and inclusion and avoid bias, protect privacy and security, provide transparency and enable accountability.**

(10) **Invest in more research to explore and enhance methods that enable intelligibility of AI systems.**

(11) **Advance a common framework for documenting and explaining key characteristics of datasets.**

Acknowledgements

This paper has been developed with contributions from a year-long initiative titled '*More Trust, More Data, Better Health: How does Europe grasp the innovation opportunity?*', sponsored by Microsoft and conducted in partnership with Fipra International.

The initiative was launched at the European Parliament in November 2017 at an event hosted by Seán Kelly MEP and with participation of a wide range of stakeholders including patient representatives, regulators, researchers, academics, policymakers and industry. The following 12 months saw events held across Europe; in Stockholm, Milan and at the European Health Forum Gastein in Austria.

At each event, stakeholders discussed the most pressing issues they face in their efforts to help push forward the digital transformation of healthcare and shared success stories showing what could already be achieved with today's technology and in the current regulatory environment.

We would like to extend our gratitude to these participants for joining us on this journey and helping us to better understand the environment they operate in. The knowledge we have gained has contributed to the development of this paper and we believe that it will ultimately help to make our healthcare systems more fit to face the challenge of the future and better able to respond to the needs of citizens, patients, and society at large.

Below is an overview of the events held as part of this project and those individuals who participated in its different stages.