

Unclassified

English text only

3 June 2021

**DIRECTORATE FOR EMPLOYMENT, LABOUR AND SOCIAL AFFAIRS
HEALTH COMMITTEE**

Health Working Papers

OECD Health Working Paper No. 128

Laying the Foundations for Artificial Intelligence in Health

Tiago Cravo Oliveira Hashiguchi¹, Luke Slawomirski² and Jillian Oderkirk¹

JEL classification: I10, F50, H51, H87, O38

Authorised for publication by Stefano Scarpetta, Director, Directorate for Employment, Labour and Social Affairs

(1) OECD, Directorate for Employment, Labour and Social Affairs, Health Division

(2) Independent Health Economist

All Health Working Papers are now available through the OECD Website at
<http://www.oecd.org/els/health-systems/health-working-papers.htm>

JT03477600

OECD Health Working Papers

<http://www.oecd.org/els/health-systems/health-working-papers.htm>

OECD Working Papers should not be reported as representing the official views of the OECD or of its member countries. The opinions expressed and arguments employed are those of the author(s).

Working Papers describe preliminary results or research in progress by the author(s) and are published to stimulate discussion on a broad range of issues on which the OECD works. Comments on Working Papers are welcomed, and may be sent to health.contact@oecd.org.

This series is designed to make available to a wider readership selected health studies prepared for use within the OECD. Authorship is usually collective, but principal writers are named. The papers are generally available only in their original language – English or French – with a summary in the other.

This document and any map included herein are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.

The statistical data for Israel are supplied by and under the responsibility of the relevant Israeli authorities. The use of such data by the OECD is without prejudice to the status of the Golan Heights, East Jerusalem and Israeli settlements in the West Bank under the terms of international law.

© OECD 2021

You can copy, download or print OECD content for your own use, and you can include excerpts from OECD publications, databases and multimedia products in your own documents, presentations, blogs, websites and teaching materials, provided that suitable acknowledgment of OECD as source and copyright owner is given. All requests for commercial use and translation rights should be submitted to rights@oecd.org.

Acknowledgements

This OECD Health Working Paper is based on a document prepared by the Organisation for Economic Co-operation and Development's Directorate for Employment, Labour and Social Affairs (ELS), and Directorate for Science, Technology and Innovation (STI), as an input for the discussions in the G20 Digital Economy Task Force in 2020, under the Saudi Arabian Presidency. The authors thank the country delegations and experts who provided comments and information during the drafting of this report.

The authors would like to thank Sarah Box (with STI) for her extensive comments, input and direction. A number of colleagues at both directorates provided further comments and direction, including Martin Wenzl, Francesca Colombo and Mark Pearson (with ELS), as well as Andrew Wyckoff, Dirk Pilat and Alistair Nolan (with STI). Isabelle Vallard, Caroline Berchet and Marie-Clémence Canaud provided assistance. The authors are grateful to Zaina Konbaz (with the National Digital Transformation Unit in the Kingdom of Saudi Arabia) and to Marc Wierzbitzki (with Kearney) for their comments and suggestions in preparation of the earlier document drafted for the G20 Digital Economy Task Force.

The views expressed in this document are the views of the authors and not necessarily the views of any OECD country or individual expert.

Abstract

Artificial intelligence (AI) has the potential to make health care more effective, efficient and equitable. AI applications are on the rise, from clinical decision-making and public health, to biomedical research and drug development, to health system administration and service redesign. The COVID-19 pandemic is serving as a catalyst, yet it is also a reality check, highlighting the limits of existing AI systems. Most AI in health is actually artificial narrow intelligence, designed to accomplish very specific tasks on previously curated data from single settings. In the real world, health data are not always available, standardised, or easily shared. Limited data hinders the ability of AI tools to generate accurate information for diverse populations with potentially very complex conditions. Having appropriate patient data is critical for AI tools because decisions based on models with skewed or incomplete data can put patients at risk. Policy makers should beware of the hype surrounding AI and identify and focus on real problems and opportunities that AI can help address. In setting the foundations for AI to help achieve health policy objectives, one key priority is to improve data quality, interoperability and access in a secure way through better data governance. More broadly, policy makers should work towards implementing and operationalising the OECD AI Principles, as well as investing in technology and human capital. Strong policy frameworks based on inclusive and extensive dialogue among all stakeholders are also key to ensure AI adds value to patients and to societies. AI that influences clinical and public health decisions should be introduced with care. Ultimately, high expectations must be managed, but real opportunities should be pursued.

Résumé

L'intelligence artificielle (IA) a le potentiel de rendre les soins de santé plus efficaces, efficaces et équitables. Le développement des applications de l'IA sont en augmentation, allant de l'aide à la décision clinique et de santé publique, la recherche biomédicale et au développement de médicaments, en passant par l'administration du système de santé et la réorganisation de l'offre de soins. La pandémie de COVID-19 agit comme un catalyseur, mais elle révèle également les limites des systèmes d'IA actuels. La plupart des applications de l'IA dans le domaine de la santé correspondent en pratique à des systèmes limités d'intelligence artificielle, conçus pour accomplir des tâches très spécifiques sur des données préalablement recueillies dans des environnements uniques. En réalité, les données de santé ne sont pas toujours disponibles, standardisées ou facilement partagées. Ces limites entravent la capacité des outils d'IA à générer des informations précises pour diverses populations dont les états de santé sont potentiellement complexes. Il est dès lors essentiel de disposer de données de qualité sur les patients pour informer les outils d'IA, car les décisions basées sur des modèles nourris par des données biaisées ou incomplètes peuvent nuire à la sécurité des soins. Les décideurs politiques devront se montrer prudents face à la surmédiation entourant l'IA et se focaliser sur les vrais problèmes et opportunités que l'IA peut générer. Pour faire en sorte que l'IA contribue à la réalisation des objectifs de politique de santé, il est essentiel d'améliorer la qualité des données de santé, leur interopérabilité et leur accès sécurisé, notamment à travers une meilleure gouvernance. Plus généralement, les décideurs publics devraient s'efforcer de mettre en œuvre et de rendre opérationnels les principes de l'OCDE en matière d'IA, mais aussi investir dans les technologies et le capital humain. Des cadres institutionnels et politiques solides, reposant sur un dialogue inclusif et approfondi entre tous les acteurs, sont également essentiels pour maximiser la contribution de l'IA. Une IA qui influence les décisions cliniques et de santé publique doit être introduite avec précaution. De manière générale, les grandes attentes doivent être tempérées, mais les réelles opportunités saisies.

Table of contents

OECD Health Working Papers	2
Acknowledgements	3
Abstract	4
Résumé	5
1 Introduction	7
2 Artificial intelligence in health: profound potential but also risks	9
2.1. AI has profound potential to transform health care for the better	9
2.2. Applications of AI in health raise legitimate concerns and anxiety	16
3 Priority for policy: beware the hype and lay the foundations for AI	19
3.1. Fostering a digital ecosystem for AI, starting with health data governance	20
3.2. Operationalising the OECD AI Principles will be challenging but fundamental	21
3.3. Putting in place regulation and guidance that promote trustworthy AI	22
3.4. Building human capacity and preparing the workforce for the change	23
3.5. Investing strategically and sustainably in AI research and development	23
4 Conclusion	25
References	26
OECD Health Working Papers	32
Recent related OECD publications	33
Figures	
Figure 2.1. Scientific research on AI in health is booming	9
Boxes	
Box 1.1. What is artificial intelligence?	8
Box 2.1. Spotlight on AI and COVID-19: diagnosis	11
Box 2.2. Spotlight on AI and COVID-19: early warning systems	12
Box 2.3. Spotlight on AI and COVID-19: accelerated drug discovery and development	14
Box 2.4. Spotlight on AI and COVID-19: limits to what can be achieved in short-term	18

1 Introduction

1. Artificial intelligence (AI) has the potential to deliver a genuine leap in global productivity, with an impact on humanity as profound as those of steam power and electricity, and the promise of improving human health and welfare. Yet the risks of harm are also high. AI can entrench and even amplify existing socio-economic problems, rather than ameliorating them. At worst, it can be deployed towards nefarious and destructive purposes. Which path will be taken will largely depend on policy choices.
2. The potential impact of AI on health and health care is profound, for many reasons: the growing volume of electronic health data can power AI algorithms; the inherent complexity of the health sector, its reliance on information to solve problems, and the variability and complexity of how disease interacts with individuals and populations make health a perfect ecosystem for applying AI to address challenges. AI can be deployed in just about any facet or activity of the health sector, from clinical decision-making and public health, to biomedical research and drug development, to health system administration and service redesign. There is particularly promising potential to: tackle unwarranted variation in care; reduce avoidable medical errors; lessen inequalities in access, health and health care; and cut down on waste and low-value care (OECD, 2017^[1]).
3. The COVID-19 pandemic has accelerated the development and deployment of AI applications in health care. This creates new opportunities but also risks. The pandemic is serving as an important catalyst for addressing long-standing barriers to applying AI to health data, and has accelerated the acceptance of AI tools by health care providers. However, this acceleration in AI development, has created the risk of insufficiently tested tools being fast tracked through regulatory processes, potentially harming patients and causing a backlash that sets back the digital transformation in health more broadly.
4. It is essential to keep the longer-run in mind, even when facing a crisis, so that fundamental principles such as data privacy protection and patient safety are not lost or weakened. Presently, AI tools remain quite rudimentary and limited in their ability to autonomously act on information or aid human decision-making (see Box 1.1 below). While AI could put a share of jobs at risk of automation in other industries (OECD, 2019^[2]), today's AI tools do not have the decision-making capacities (not to mention the tacit knowledge and soft skills) of health care professionals, and experts do not expect most health care jobs to be at risk of automation in the foreseeable future.
5. The COVID-19 pandemic has exposed weaknesses in health information systems' capacity to meet the data needs of a global health emergency, weaknesses that the OECD Council Recommendation on Health Data Governance seeks to address. These include issues around data timeliness, quality, suitability for dataset linkages and accessibility; ensuring public transparency of data collection, use and protection; developing efficient exchange and interoperability of health data, including global health data standards; and eliminating unnecessary barriers to public-private and cross-border data sharing and research collaborations (OECD, 2016^[3]). It is important to create a coherent ecosystem with agreed rules, without unnecessary barriers that prevent many countries from participating in AI tool development. Initiatives include data hubs, and also federations and networks with both public and private sector members, that ensure data are FAIR (meet the principles of findability, accessibility, interoperability, and reusability) and protected.

Box 1.1. What is artificial intelligence?

Artificial intelligence (AI) describes algorithms that enable a computer or a computer-controlled robot to make decisions based on new input, and to update themselves based on new information as it becomes available. These decisions are often informed by “real-time” data. Over the past several years, AI has advanced to perform more sophisticated tasks, in great part due to machine learning and the arrival of deep learning and convolutional neural networks, but also driven by more data and computational power. Machine learning uses statistical techniques to derive relationships and rules directly from the data. Deep neural networks, for example, can mimic human abilities, such as recognizing speech, images and patterns.

Today, applications of AI in health are restricted to applications of artificial narrow intelligence, or “applied” AI, which for the most part seek to accomplish very specific tasks on previously curated data from one setting. An important concern exists about AI tools trained on data from a specific population and then applied to diverse patient populations. Predictions could fail to generalise to different populations or settings, and might exacerbate existing inequalities and biases that are present within the training data. With the AI industry being notoriously gender and race imbalanced, and with health professionals already overwhelmed by other digital tools, there may be little capacity to catch errors and biases.

Artificial general intelligence – autonomous machines capable of general intelligent action, capable of generalised and abstract learning across different cognitive domains – is not yet available and still some way off. There is broad agreement that artificial narrow intelligence will generate significant new opportunities, risks and challenges, and that the possible advent of an artificial general intelligence, perhaps sometime during the 21st century, would greatly amplify these consequences.

Source: (OECD, 2019^[41]).

6. OECD countries with the greatest capacity to use health information to support policy planning, decision-making and research during the COVID-19 pandemic are those that already had a digital strategy with three key elements:

- Mature health information systems containing key health data across the continuum of health care with national coverage; with data that are timely, of high quality, contain key contextual information and can be linked with one another;
- Standardised, coherent and accessible national electronic clinical record systems that address fragmentation of data across health care silos and enable “one patient, one record” for complete views of health care trajectories and outcomes; and
- Comprehensive health data governance with legislation and policies that protect privacy and data security while enabling data about health and health care to be generated, linked, accessed and analysed for uses within the public interest, including cross-border collaborations.

7. This paper describes current applications of AI in the health sector, including new applications responding to COVID-19, to showcase the potential of this technology to improve health and health care, but also to highlight concerns and risks. This report also provides policy guidance and best-practice principles for the assessment and adoption of AI technologies to advance health research, care and governance.


2 Artificial intelligence in health: profound potential but also risks

2.1. AI has profound potential to transform health care for the better

8. The use of AI in everyday health care practice is still extremely limited; however, the number of applications is growing rapidly, and this trend has accelerated since the onset of the COVID-19 pandemic. Applications range from those serving clinical settings, to biomedical research, health system administration and management. Virtually every aspect of health care delivery seems amenable to contributions from AI. The number of scientific publications relevant to AI in health in the top ten OECD countries and regions with most publications has grown from just 35 in 1980 to over 3 400 in 2019 (see Figure 2.1).

Figure 2.1. Scientific research on AI in health is booming

Number of relevant scientific publications in health, by top ten OECD country or region, from 1980 to 2019



Note: Please see methodological note (https://www.oecd.ai/assets/files/Methodology_20200219.pdf) for more information.

Source: OECD.AI (2020), visualisations powered by JSI using data from MAG, www.oecd.ai

2.1.1. AI promises to change clinical practice in the not so distant future

9. Much promise and research activity concerns the potential application of AI in the clinical setting, such as the automation of diagnostic processes and clinical decision-making, among others. A lot of activity to date has focused on **diagnostic imaging** (Neri et al., 2019^[5]), especially more recently in the context of

COVID-19 (see Box 2.1). There are many impressive examples¹ in the research setting where AI tools can perform as well as – or even better than – certain clinicians. These range from retinal scans to tumour detection in radiology. Another application of AI that has had success in the clinical research setting is **radiomics**: the extraction of certain features from multiple diagnostic images to produce a quantitative ‘picture’ of an anatomical region of interest. Such features may be used for prognosis and to predict response to treatment. By building patients’ radiological features, AI might enable analysis of and correlation between radiomics and other data (e.g. genomics, biopsies) at a level of accuracy that humans cannot achieve.

10. In the surgical field, AI can aggregate diverse sources of information – including patient risk factors, anatomic information, disease natural history, patient values and cost – to predict the consequences of surgical decisions. **AI-enhanced remote-controlled robotic surgery** can improve the safety of interventions where clinicians would otherwise be exposed to high doses of ionizing radiation, and makes surgery possible in anatomic locations not otherwise reachable by human hands (Svoboda, 2019^[6]). Robotic² surgery is already being used in 10 Irish hospitals (public and private), with St Vincent’s University Hospital having developed its own Robotic Surgery Programme. As surgical robots develop further, and come to embody greater intelligence, surgeons and robots could function in collaborative ways, for instance by recording steps in a surgery and providing feedback to the surgeon, providing cognitive support, or even taking over certain more routine steps in a procedure.

11. In many medical interventions, customising clinical diagnostic reference levels based on appropriateness criteria and on patient characteristics – such as age, body mass index, vital signs and prior exposure to disease or risk factors – is an important risk management process. AI can be an optimising tool for assisting clinicians in providing a **personalised patient protocol**, in tracking the patient’s dose parameters, and in providing an estimate of the risks associated with cumulative doses. In Queensland, Australia, a consortium of researchers, industry and health organisations is looking to use AI to analyse complex patterns in cancer patients’ genetic data. These patterns are then used to develop individual treatment plans, seeking to decrease the cost of treatment by reducing unnecessary interventions, as well as by recommending interventions with a higher likelihood of success.

12. AI tools can also affect the daily workflow of a health care practice by assigning treatment priority to patients based on appropriateness criteria. **Clinical decision support systems** can assist referring physicians to choose the most appropriate investigative procedure based on the level of evidence for appropriateness, and the level of emergency. Such a system can be enhanced by AI to improve the speed and accuracy of decision making speed and optimise clinical workflow. For example, researchers at Clalit Research Institute, the largest health payer-provider network in Israel, have developed a model, based on a combination of regression and machine learning, with superior discriminatory power to that of previously reported models, that can be used to identify high risk of readmission early in the admission process (Shadmi et al., 2015^[7]). Other AI systems have been used to predict the likelihood that patients will follow treatment, so that staff can follow up at-risk patients (Combs, 2019^[8]). A deep learning algorithm developed in Norway has been shown to accurately predict colorectal cancer outcomes, with the potential to target treatment depending on patient risk of severe outcomes (Skrede et al., 2020^[9]).

¹ Some of the numerous recent experimental applications of AI include determining skeletal age using pediatric hand radiographs **Invalid source specified.**, breast cancer detection in mammography and Magnetic Resonance Imaging (MRI), chest radiography interpretation, liver lesion characterisation on ultrasound and Computed Tomography (CT), brain tumour, and prostate cancer detection.

² Other current and emerging uses of AI-driven robots in health, not restricted to clinical care, include automation of laboratory testing (especially timely in the context of COVID-19), screening of patients (for example using swabs), and sterilisation of health care settings, among many others (Nolan, 2020^[76]).

Box 2.1. Spotlight on AI and COVID-19: diagnosis

There is promise in using AI to diagnose COVID-19 but some debate surrounding its use in imaging

Around the world, researchers and health care providers are working together to develop deep learning networks that can help diagnose COVID-19 more quickly and with fewer errors. They are also trying to help identify those patients most likely to develop acute respiratory failure. AI tools have been developed to support diagnoses of COVID-19 and to monitor disease progression through both **CT scan and radiology (X-ray) images**. AI software supporting interpretation of X-ray images has been made available as open-source software to countries and would be particularly useful in health care settings with limited access to CT scanners (Lunit, 2020^[10]; Hwang et al., 2019^[11]). AI tools to interpret and score CT images are being used to distinguish COVID-19 from other causes of pneumonia. A deep learning AI called COVNET was developed and tested in China and reports a high degree of sensitivity and specificity to detect COVID-19 (Li et al., 2020^[12]). The US Food and Drug Administration (FDA) recently announced that it will permit the use of cleared AI algorithms to detect COVID-19 and prioritise patients from CT scans of the lung or parts of a lung (ITN, 2020^[13]). AI tools for CT scan image interpretation may have uses beyond diagnosis. Studies point to the possible use of AI to support evaluation of the progression of the disease in patients and to contribute to the development of drug therapies (Zhang et al., 2020^[14]).

AI-powered diagnostic tools based on CT scans have received mixed reviews from frontline health care providers as well as researchers (Olson, 2020^[15]). Some health care providers believe the technology allows radiologists to produce more diagnostic reports more quickly, with a typical manual read of a CT scan taking 15 minutes compared to 10 seconds with AI. Another benefit is that the algorithms provide initial diagnoses that can help radiologists – especially younger less experienced ones – to consider elements they may have missed otherwise. Finally, a CT scan can also provide important prognostic information, helping providers to allocate intensive care to patients with worse prognoses.

Not all providers and researchers are as receptive though, questioning both the value of CT scans in diagnosing COVID-19 generally as well as the validity and accuracy of the AI algorithms used to read scans. The American College of Radiology – representing nearly 40 000 workers – has advised against relying on CT scans to diagnose COVID-19 stating the method is not specific (ACR, 2020^[16]). Hope and colleagues (2020^[17]) consider that “framing CT as pivotal for COVID-19 diagnosis is a distraction during a pandemic” that changes little in terms of workflows and care decisions, an opinion echoed by radiologists in Italy (Olson, 2020^[15]). There is a lack of high-quality evidence of effectiveness, as well as scepticism that the algorithms can be validated in a short timeframe, especially as similar algorithms in breast cancer have taken years to validate.

AI applications are also emerging to **protect health care providers from SARS-CoV-2 infection**. Applications for detecting fever using infrared sensors have been deployed in China. This technology is useable even in crowded areas, such as railway stations. The US Tampa General Hospital has deployed a similar technology at entrances to the hospital to conduct a facial thermal scan and detect other symptoms of COVID-19 including sweat and discoloration, as well as other autonomous monitoring of staff and patients that facilitates track and trace efforts and authorises and tracks access to restricted areas (HBR, 2020^[18]). A possible AI application for protecting health care providers involves intelligent robots developed by Boston Dynamics and MIT and deployed at the Brigham and Women’s Hospital and Massachusetts General Hospital to perform tasks such as obtaining vital signs or delivering medications to patients in surge clinics and inpatient wards (Boston Dynamics, 2020^[19]).

13. The process of reviewing the records of patients with unexpected outcomes to identify recommendations for improvement, known as a **clinical audit**, is an important but labour-intensive task for health professionals in all settings. It is therefore a prime target for automation. A recent study compared human performance in generating audits on a neurosurgical ward, with that of an AI algorithm (Brzezicki et al., 2020^[20]). The audits conducted by AI were significantly more accurate and logically consistent, and far less costly than a human audit. For example, the mean time to deliver a report was 5.80 seconds for AI compared to 10.21 days for humans. AI may also be helpful in automating pharmacovigilance (i.e. drug safety) case processing (Schmider et al., 2019^[21]).

2.1.2. AI is already making a mark in biomedical research and public health

14. AI is currently more commonly applied in **biomedical and population health research** compared to the clinical setting. The reasons for this disparity are unclear. However, the exciting potential of combining AI with large datasets was demonstrated recently when multiple organisations rapidly spotted information about a “pneumonia of unknown cause” in China. See Box 2.2 for more details.

Box 2.2. Spotlight on AI and COVID-19: early warning systems

AI-based early warning systems could be key to identifying new hotspots as countries exit lockdowns

By linking real-time data from multiple sources (from insurance coverage databases to customs and immigration data), some countries have been able to raise alerts based on patients’ travel history and clinical symptoms when they present at health care facilities. Big data outside of the health system – from social media and web searches, to environmental sensors and satellites – can also help.

As experience around the world has shown, SARS-CoV-2’s high basic reproduction number, its relatively long incubation period, and the elevated incidence of asymptomatic patients have led to outbreaks quickly overwhelming some health care systems. Being able to rapidly identify hotspots and contain further spread is crucial. The United States’ Centers for Disease Control and Prevention (CDC) have partnered with academic researchers to feed machine learning algorithms with administrative data from CDC, Google searches and Twitter to predict the number of infections in real-time (Hao, 2020^[22]). Data from social networks alone are likely no substitute for epidemiological surveillance data, as past experiences of Google Flu Trends show (Olson et al., 2013^[23]).

Travel data, even when anonymised, can help machine learning models predict the most likely locations of new infections, which might inform border controls. It has been much-publicised that machine learning algorithms were used by various organisations – among them Canadian firm BlueDot and American-based HealthMap and Metabiota – to search global media for information on various infectious diseases, first spotting cases of COVID-19 in late December. BlueDot, for example, issued an alert to its clients about potential risks of travelling to cities like Shanghai, Tokyo and Hong Kong (China) a week before the United States CDC and the World Health Organisation (WHO) issued alerts.

It has been pointed out that human analysts also raised alerts around the same time as organisations using AI, and that human interpretation was still necessary to attribute relevance to the alerts raised by AI algorithms (Naudé, 2020^[24]). In this case, the added value of the AI technologies is in assisting humans to interpret millions of data points from mainstream news, online content and other information channels in multiple languages. Moreover, AI can also help policy makers and providers understand public sentiment. In Australia, the Department of Health is doing so, using social media data to understand public opinion on COVIDSafe (a digital contact tracing mobile application).

15. **Precision medicine** is an area where AI fed by large volumes of personal health data can produce information that can help researchers tailor medical treatment to the individual characteristics of each

patient (Eichler et al., 2018^[25]). Most medical treatments in use today work for a large majority of patients, but for some patients treatment either fails to deliver benefits or may even put them at risk of adverse events. Being able to pre-emptively identify those patients who are most likely to benefit, or otherwise, from therapy can lead to better clinical outcomes. Developing precision medicine turns the traditional research paradigm on its head – statistical noise and variation are the variables of interest. These cannot be determined by prospective, traditional clinical trials. AI models working with large and varied personal health data could make precision medicine a reality. To affect routine care, however, the data need to represent the diversity of patient populations.

16. AI might also increase the probability of **new drug discovery** including important drugs such as novel antibiotics and antivirals, as well as pharmacological therapies for COVID-19 (see Box 2.3). For example, researchers at the Massachusetts Institute of Technology (MIT) recently trained a deep learning algorithm to predict molecules' potential antimicrobial activity (Stokes et al., 2020^[26]). The algorithm screened over one billion molecules and virtually tested over 107 million, identifying eight antibacterial compounds that were structurally distant from known antibiotics. One of those antibacterial compounds effectively treated resistant infections in mice. In similar work, a deep learning model of small-molecule drugs was used to identify key biological mechanisms implicated in fibrosis and other diseases (Zavoronkov et al., 2019^[27]). AI can also assist with manufacturing and quality control, not only of drugs but also in the medical devices industry.

17. AI can also improve matching of individuals to **clinical trials** (Lee and Lee, 2020^[28]). Patients can be identified to enrol in trials based on more sources than clinical or administrative data. The criteria for including patients in a trial could take significantly more factors, like genetic information, into account in order to target more specific populations. This can enable trials to be smaller, shorter, to be set up more effectively and therefore be less expensive, all without sacrificing statistical power. AI may also help address the documented problem of underrepresentation of minorities in clinical trials.

18. Disease **prediction and prevention** are other promising areas for AI. Among other applications³, researchers have demonstrated the ability of an algorithm to accurately predict the risk of emergency admission based on an individual's electronic health record data (Rahimian et al., 2018^[29]) and even web browsing history, as well as more recently the risk of mortality for COVID-19 patients (Yan et al., 2020^[30]). Coupled with data from outside the health system (e.g. social media or e-commerce), such algorithms could be even more powerful. However, predictions of future health conditions raise concerns about privacy and ethical uses of data. Policy makers should consider building trust among data subjects and data custodians by restricting access to and use of identifiable data and authorising data uses by patient consent.

19. Several initiatives to support **AI tool development for biomedical COVID-19 research and public health** have emerged. These include the COVID-19 Open Research Dataset developed by the Allen Institute for AI to support semantic search of thousands of research papers related to COVID-19 (SemanticScholar, 2020^[31]). The COVID-net open-access neural network was developed to help researchers across the globe on AI detection of COVID-19 from X-ray images (Heaven, 2020^[32]). A German initiative aims to overcome barriers to share data for the purpose of developing AI algorithms by asking individuals to donate their health data to support machine learning applications (Delcker, 2020^[33]).

³ Machine learning algorithms using internet search and social media data have also been used by the City of Chicago, in the United States, to predict and pinpoint the source of foodborne disease outbreaks much faster than traditional inspection methods (OECD, 2019^[59]).

Box 2.3. Spotlight on AI and COVID-19: accelerated drug discovery and development

AI may help identify treatments and vaccines to end the pandemic

It is widely agreed among epidemiologists and public health officials that until an effective treatment and/or vaccine are developed and mass-produced, non-pharmaceutical interventions like social distancing will be the only way to contain COVID-19. Pharmaceuticals (especially vaccines) can take years to develop, test and distribute, and drugs need to undergo rigorous and time-consuming assessments of safety and cost-effectiveness before being widely used. There has thus been considerable interest in using AI techniques to accelerate drug development and use, by helping identify new drugs and repurpose older ones, but also by adapting RCTs so that they are conducted in a faster and more effective way.

A machine learning model from a London-based company has found that a drug used in rheumatoid arthritis (Baricitinib, commonly referred under its brand name Olumiant) may be effective against the virus (Mccall, 2020^[34]). Researchers from Korea have used deep learning to test whether old compounds could be effective against SARS-CoV-2 (Beck et al., 2020^[35]). Several other institutions are using AI to identify treatments and develop prototype vaccines. Google DeepMind has recently contributed to gaps in researchers' understanding of SARS-CoV-2 by predicting the structure of proteins associated with the virus (DeepMind, 2020^[36]).

While promising, there is a long way to go from computer models to human trials and market approval, and drugs identified using AI will still need to transition from the 'bench to the bedside'. AI may help, accelerating the clinical testing needed to move from the lab to the pharmacy (Woo, 2019^[37]). The use of AI in drug development and testing is still nascent and it is unclear whether AI will help deliver a vaccine or an effective treatment for COVID-19. The challenge is considerable, with or without AI.

2.1.3. The use of AI for administrative purposes could have the most immediate impact

20. Health systems are notoriously complex. Providing services to individuals and populations involves a wide range of actors and institutions: patients, professionals, health care facilities and organisations, laboratories, imaging facilities, pharmacies, administrators, payers, and regulators. In parallel to the care provided, administrative workflow includes **scheduling, billing, coding, managing workflow and payment**. One of the principal and immediate applications of AI is to perform these mundane, repetitive tasks in a more efficient, accurate and unbiased fashion (NAM, 2019^[38]). This can free up resources, especially staff time, for activities directly related to patient care. The back-office (e.g., scheduling, billing, coding and payment) is also a relatively safer testing ground for AI, as errors tend to mostly carry administrative or financial risk, and only in select cases could they jeopardise patient safety.

21. **Coding clinical terminology** is an obvious administrative use of AI-based automation. This describes the process of extracting information from clinical records and codifying it using classifications such as the International Classification of Diseases (ICD) or Diagnosis-related groupings (DRGs). Coding is a complex, labour-intensive process. Yet, accuracy is pivotal for reimbursement, administration and research. While computer-assisted coding has existed for more than a decade, AI can enhance the accuracy and transparency of clinical coding. AI tools are being developed⁴ to aggregate long, narrative clinical records with other data (e.g. medication orders, imaging and laboratory tests) for analysis using deep learning and other AI models. Human review is still required, but a part of the processing effort is performed by machines. In the short-term, AI may help human coders and create checks for policy makers

⁴ For example, in China, researchers developed a deep-learning based natural language processing system to comb through more than 1.3 million electronic health records to extract relevant information and generate diagnoses **Invalid source specified.**

and payers. In the long-term, near-full automation might be achieved, but will undoubtedly rely on data quality and comprehensiveness, algorithm transparency and accuracy, and the trust of those relying on the outputs. Coders' time can be reallocated to more advanced tasks, including the important work of developing, testing and refining algorithms.

22. The ability of AI to analyse free text can be particularly powerful where administrative decisions are based on narrative data. These decisions include **prior authorisation (PA)**, which is needed in most health systems to supply health services and products to patients (Accenture, 2018_[39]). PA requires the submission of patient information along with the proposed request and justification. Determination requires professional skill, analysis, and judgment. Automating this can improve the speed, consistency, and quality of decisions. With good governance, such a process could lead to fewer appeals and limit liability. A range of options exist to deploy AI in the context of PA. For example, AI methods could be used to triage cases to the appropriate level of reviewer. A more complex algorithm could find relevant information across one or more datasets to, for example, determine the eligibility of patient for a procedure, and estimate the costs, the benefits and the risks associated with it. Moreover, lessons from applying this process at scale could be particularly useful in deploying similar AI functions in clinical settings, such as triaging images for human review or automated reading.

23. Much like AI can be taught to spot irregularities in medical images, algorithms can also learn to look for **fraudulent activity** in health care⁵. A variety of previously curated claims data (i.e. identified as fraudulent, valid, or requiring investigation) can teach an algorithm to identify potentially fraudulent claims, including subtle systematic cases of over-servicing and upcoding (i.e. using a code for a more expensive medical service than was performed), as well as potential underuse of available services. Combining claims with other data (e.g. clinical and demographic) can increase the accuracy of the audit process. For example, in addition to age and co-morbidities, information on a patient's entire medical history may help ascertain if a procedure was in fact necessary, if additional payment adjustments for case complexity are justified, or if a potentially beneficial treatment may have been missed.

24. Managing **clinical workflow** is another example where AI can add value right now. In Canada, the Humber River Hospital has partnered with GE Healthcare to process real-time data from multiple sources and systems (GE Healthcare Partners, 2018_[40]). The so-called "command centre" makes use of advanced and predictive analytics to provide staff with alerts on, among other things, delayed patient care activity, unbalanced physician and staff workload and unusual situations that may increase risk to patients. The information is used for real-time decision support, including prioritising patient care activities and discharges, making short-term staffing decisions and mitigating potential bottlenecks before they occur. The system is estimated by the facility to save millions per year while improving patient experience.

25. **Clinical scheduling** more broadly can be enhanced by AI, and almost immediately. A no-show is not only an inconvenience but also a waste of highly qualified human resources who are typically in high demand. Algorithms fed on historical data can predict which patients may not attend and take proactive action to manage this. Beyond blanket or even targeted reminders, AI can address a patient's needs and queries⁶. It is not only patient appointments that need scheduling. An Irish company awarded a seal of excellence by the European Commission has developed an AI-based automated scheduler for

⁵ For example, SAS is working with DentaQuest, a health insurer in the United States with 24 million members, to detect fraud using SAS Enterprise Miner, a data mining and analytics platform that uses AI.

⁶ For example, Northwell Health – a US health network – has launched a text-based patient outreach programme aimed at reducing no-show rates for colonoscopy procedures. The colonoscopy health chat platform uses a flexible algorithm that aims to give patients the information they need so that they feel comfortable and motivated to get screened. The chatbot addresses questions and concerns related to the procedure, its benefits, and reminds the patient of the date and location as the procedure draws closer.

laboratories, which can connect to their supply chain, and prioritise tasks for completion to optimise productivity and cycle times between manufacturing and shipping (European Commission, 2019^[41]).

26. These administrative uses of AI are far from trivial, considering that around one fifth of all health spending is estimated to be wasted on inefficiencies, including fraud. A systematic reduction represents hundreds of billions of dollars each year that could be diverted towards better ends in health systems (OECD, 2017^[1]).

2.2. Applications of AI in health raise legitimate concerns and anxiety

27. Deploying AI has the potential to transform almost every aspect of the health sector, but its use in everyday health care is still very limited. Most models are still relatively basic, with frequent problems and errors based on overfitting of data (i.e. models do not generalise well outside that data) and mistaking correlation for causation, for example (NAM, 2019^[38]). This can be highly problematic. There are also challenges in scaling up projects to the level of health systems due to, among other things, questions concerning the robustness of algorithms in the real world, a lack of high quality health data, limited institutional and human capacity to realise the potential of AI.

2.2.1. AI in health is not yet robust: for every success story there is a cautionary tale

28. Unfortunately, AI applications in health have been beset by misfires and setbacks, with hype often clashing with reality. A recent study reviewed dozens of studies claiming an AI performs better than radiologists, finding that only a handful were tested in populations that were different from the population used to develop the algorithms (Reardon, 2019^[42]). The difficulty to scale certain AI applications is often due to trivial factors. For example, the way different facilities label their images can confuse an algorithm and prevent the model from functioning well in another institution with a different labelling system. This serves to highlight that most AI in health is very narrow, designed to accomplish a very specific problem-solving or reasoning task, and unable to generalise outside the boundaries within which the model was trained (see Box 1.1).

29. Most AI applications in health are still in research and development stages, and concentrated in a few countries and regions. Figure 2.1 illustrates that, in the OECD, the top five countries/regions (United States, EU27, United Kingdom, Canada and Australia) combined have more than ten times the research papers than the bottom five. As such, most of the data used to train these models is from high-income countries and are not representative of a global population in many important ways. It is likely that algorithms used to explain or predict human behaviours based mainly on care patterns for one population will produce biased predictions outside that group (NAM, 2019^[38]). For example, an AI algorithm used to identify patients with complex needs in the United States has been shown to suffer from racial bias, assigning lower risk to non-Caucasian patients compared to Caucasian patients. Using health costs as a proxy for health needs, the algorithm learned that since less money is spent on non-Caucasian patients who have the same level of need, non-Caucasian patients are healthier than equally sick Caucasian patients (Obermeyer et al., 2019^[43]).

30. A related challenge is **overfitting**, which occurs when an AI model learns statistical irregularities specific to the data on which it is trained. Unless the training data are vast (and therefore difficult and costly to create) the model may confuse irrelevant noise with the actual signal. An overfitted model will not generalise to different input data. Overfitting was one of the problems in the IBM Watson cancer initiative, where the model was trained on hypothetical data and then graduated to real clinical situations too quickly (Strickland, 2019^[44]). A key limitation of contemporary AI is that most machine-learning-based prediction models are based on **correlation, not causation** (NAM, 2019^[38]). Previous studies have identified counterintuitive associations that lead to nonsensical predictions. For example, a model that predicts risk of death for a hospitalized individual with pneumonia learned that patients who have asthma and

pneumonia are less likely to die than patients who only have asthma, because patients with asthma and pneumonia receive more aggressive treatment and thus have lower mortality rates. In another example (NAM, 2019^[38]), the time a lab result is measured can often be more predictive than the value itself (e.g. if it is measured at 2am).

31. Algorithms that learn from human decisions could also **learn human mistakes, biases and stereotypes**. Given that the AI sector is extremely gender and race imbalanced (AI Now Institute, 2019^[45]), biased predictions might not be flagged by developers working to validate model outputs, especially if those teams do not involve the users and subjects of their models. For example, Apple's HealthKit, an application to track intake of selenium and copper, neglected to include a women's menstrual cycle tracker until iOS 9; the development team reportedly did not include any women (NAM, 2019^[38]).

32. Finally, the outputs of an AI model must be presented to users as meaningful information: e.g. an alert or pop-up window within electronic health record software. Evidence suggests that the implementation of health information systems can result in unintended consequences. These include alert fatigue, imposition of additional workloads for clinicians, disruption of interpersonal (including doctor-to-patient) communication styles, and generation of specific hazards that require a higher level of vigilance to detect. A growing number of health workers are overwhelmed (Dyrbye et al., 2017^[46]), with some suffering from **change fatigue**: getting tired of new initiatives and how they are implemented (Garside, 2004^[47]). Against this backdrop, the black-box nature of AI algorithms may result in either resistance from clinicians to adopt and use their predictions, or a blanket acceptance of their outputs with little critical assessment of the potential for biased and suboptimal predictions. Getting the interface between the algorithmic world and the brick-and-mortar world right is key to success. This requires concerted effort to co-designing systems and their applications with the end user.

2.2.2. Poor governance of health data means AI requires a lot of human curation

33. For most methods and applications, AI typically needs large amounts of training data. If these data even exist – which is not guaranteed – it is likely they will need human **curation**. For example, the data need to be stratified by patient cohort, segmented to extract the region of interest for AI interpretation, filtered, cleaned and labelled. This can be very time- and labour-intensive. Moreover, it is not easy to automate and may not be automatable soon.

34. Curation is also essential to establish a 'ground truth', a fundamental concept in supervised learning. This means validating the output as a true positive or negative when the machine is in the learning phase. This can be problematic in health, where **data are notoriously messy**, and classifications or interpretations of the underlying data can be wrong to begin with. For example, in the clinical setting, determining the presence or absence of pathology often involves considerable uncertainty. For example, in oncology, there can be disagreement on what constitutes cancer requiring medical intervention. Experienced pathologists will often disagree about the histopathology and diagnosis, particularly in early-stage lesions. If the aim of an AI application is to help detect cancer early, this **disagreement presents a conundrum** concerning how to train the algorithm, as the resulting AI tool should not over- or under-diagnose tumours.

35. Even if there were perfect agreement on what the ground truth is, most existing medical data are not readily available for use in AI algorithm development. They are also rarely of high enough quality, meaning they are not easy to exchange, process or interpret, and may be riddled with errors and inaccuracies. While most people think health care is awash with big data, in many countries the reality is much closer to having **"a large number of disconnected small data"** (Lehne et al., 2019^[48]). A central source of patient-level data for AI development – the electronic health record – is emerging as a key resource for dataset research in only a handful of OECD countries that have implemented standardised, interoperable electronic health record (e-HR) systems offering "one patient, one record" national data (Oderkirk, 2017^[49]). Against this backdrop of insufficient data governance and data quality, it is unlikely

that AI algorithms can be used in the real world without extensive human input. More investment to improve the quality of routinely collected clinical data is needed. Deploying AI in the fight with COVID-19 illustrates some of the challenges relating to data quality and availability, as outlined in Box 2.4.

Box 2.4. Spotlight on AI and COVID-19: limits to what can be achieved in short-term

Scepticism that AI will help with this pandemic, but hope that it will be useful in the next crisis

While there is considerable optimism and excitement about the potential of AI in helping to tackle the COVID-19 pandemic, it is too early for AI to make a difference in this pandemic. A recent review of prediction models (many using machine learning methods) for COVID-19 found that, while performance was very high (almost perfect in some cases), there was a high risk of bias due to poor reporting and an unrealistic mix of patients with and without COVID-19 (Wynants et al., 2020^[50]). This is problematic because methods that are narrowly focussed on a specific task and trained using a specific set of data, may not work well outside those limits. Another recent review of the use of AI in computerised decision support systems (not limited to COVID-19) cautioned policy makers that evidence of effectiveness was limited (Cresswell et al., 2020^[51]). The lack of large curated data sets to train models is an important barrier, especially in the context of a fast-moving pandemic. No AI algorithm can make up for a complete lack of data, yet some data are often unavailable in health care. Even when data are available, people behave differently in times of crises, and certain algorithms trained on pre-crisis data will almost certainly run into trouble when predicting under extreme circumstances (Cho, 2020^[52]).

2.2.3. The carbon footprint of AI is still unclear

36. An **area of uncertainty** surrounding AI in general is its carbon footprint. Training a single AI model can emit as much carbon as five cars in their lifetimes, a likely minimum estimate because training a single model is only the beginning of an algorithm's lifecycle (Hao, 2019^[53]; Strubell, Ganesh and McCallum, 2019^[54]). While a growing appetite for digital services could mean the data centres that power them emit a rising share of global greenhouse gases, a shift to cloud computing has led to massive efficiency gains in recent years (Masanet et al., 2020^[55]). AI itself could be used to promote a circular economy and increase energy efficiency. Google DeepMind used machine learning in Google's data centres to reduce energy use for cooling and claims that efficiency gains reached 40% (DeepMind, 2016^[56]). With one estimate indicating health systems in the OECD, China and India already account for 5% of total emissions, more than aviation or shipping (Pichler et al., 2019^[57]), **AI could therefore have both positive and negative effects on carbon emissions.**

37. The ultimate net effect matters. According to the WHO, between 2030 and 2050, climate change could cause approximately 250 000 additional deaths per year, from malnutrition, malaria, diarrhoea and heat stress (WHO, 2014^[58]). The direct damage costs to health (i.e. excluding costs in health-determining sectors such as agriculture and water and sanitation), is estimated to be between USD 2-4 billion per year by 2030 (ibid). Areas with weak health infrastructure – mostly in developing countries – will be the least able to cope without assistance to prepare and respond.

38. With current AI models trained in one setting needing to be re-trained in other settings, and with the number of potential applications of AI in health rising every day, the climate impact of AI in health may not be negligible. If AI in health turns out to be a net contributor to greenhouse gas emissions, this would mean that an activity within the health sector would itself be associated with an increase in the burden of disease. Moreover, given the high entry costs to developing AI models, it is likely that development will take place mostly in high-income settings, with most of the climate effects on health being felt in low-income settings. On the other hand, if AI leads to more energy-efficient health care systems, then the impact on vulnerable areas could be disproportionately beneficial. As uncertainty remains today, **energy use associated with AI in health should be monitored and studied.**

3 Priority for policy: beware the hype and lay the foundations for AI

39. While interest in applying AI in the health sector is growing and the private sector is moving fast, use remains limited. This provides an opening for policy makers to get ahead of the curve, and examine how best to capitalise on the real opportunities that AI affords, while considering mechanisms to ensure risks are managed. In the absence of a well-defined ready-to-use menu of policy options, countries should engage in multilateral discussions on a plan of action that promotes trustworthy, safe and reliable AI in health.

40. The **OECD Principles on AI** provide a framework to guide the discussion⁷. The values-based principles aim to foster innovation and trust in AI by promoting the responsible stewardship of trustworthy AI while ensuring respect for human rights and democratic values. They comprise five complementary components:

- AI should benefit people and the planet by driving inclusive growth, sustainable development and well-being.
- AI systems should be designed in a way that respects the rule of law, human rights, democratic values and diversity, and they should include appropriate safeguards – for example, enabling human intervention where necessary – to ensure a fair and just society.
- There should be transparency and responsible disclosure around AI systems to ensure that people understand AI-based outcomes and can challenge them.
- AI systems must function in a robust, secure and safe way throughout their life cycles and potential risks should be continually assessed and managed.
- Organisations and individuals developing, deploying or operating AI systems should be held accountable for their proper functioning in line with the above principles.

41. In addition to these, the OECD Principles on AI make five recommendations to guide policy-makers and international co-operation for trustworthy AI:

1. investing in AI research and development;
2. fostering a digital ecosystem for AI;
3. shaping an enabling policy environment for AI;
4. building human capacity and preparing for labour market transformation; and
5. international co-operation for trustworthy AI.

⁷ See <https://www.oecd.org/going-digital/ai/principles/>. In June 2019, the G20 adopted human-centred AI Principles that draw from the OECD AI Principles.

42. In the context of health and health care, it is especially important for policy to foster a digital ecosystem for AI, operationalise the OECD AI principles, establish appropriate regulation and guidance, build human capacity and invest strategically and sustainably.

3.1. Fostering a digital ecosystem for AI, starting with health data governance

43. High-quality, representative data, and data generated in “real-world” situations, are essential to minimising the risk of error and bias. Creating an environment where such data – especially personal health data – are available to AI researchers and developers in a secure way that respects individuals’ privacy and autonomy is fundamental. This requires frameworks for **strong health data governance**, within and across countries, as well as developing better digital infrastructure and technological capacity.

44. The **OECD Council Recommendation on Health Data Governance** recommends that governments enable the efficient exchange and interoperability of health data, including codes, standards and the standardisation of health data; and remove barriers to effective cross-border cooperation in the processing of personal health data for data uses that serve the health-related public interest. Internationally agreed standards for health data terminology and exchange would enable a digital transformation in the health sector that fosters AI development and the interoperability of AI tools, so that benefits of investments might be shared within and across countries⁸. Without these investments, health data will be underused and returns and benefits from AI tool development will be restricted (OECD, 2019_[59]).

45. In line with the OECD Health Data Governance and AI Principles, frameworks for health data governance should emphasise transparency, public communication and stakeholder engagement, explicitly highlighting **the importance of trust** (OECD, 2016_[3]). Lack of trust among patients, the public, data custodians and other stakeholders, in how data are used and protected is a major impediment to data use and sharing. Personal health data are very sensitive, and privacy is understandably one of the most frequently cited barriers to using them. Yet, the potential benefits of using personal health data to generate new knowledge cannot be minimised, for example in the context of testing much needed drugs and vaccines (as currently highlighted by the COVID-19 crisis). Health care leaders should work to advertise the benefits of using health data, changing the discourse that sees use of data as the only risk and that ignores the foregone benefits to individuals and societies of failing to put data to work (OECD, 2019_[59]). It is also essential to dispel the idea that there is a trade-off between data protection and secondary use of health data. A risk management approach and careful implementation of good practices can enable both data protection and its use. Updated periodically, formal risk management processes could include unwanted data erasure, re-identification, breaches or other misuses, in particular when establishing new programmes or introducing novel practices.

46. For a number of reasons (e.g. representativeness and breadth of input data), many applications of AI in health would gain considerably from **cross-border collaboration** in the processing of personal health data for purposes that serve the public interest. This includes identifying and removing barriers to effective cross-border collaboration in the processing of personal health data, as well as engaging with relevant experts and organisations to develop mechanisms to enable the efficient exchange and interoperability of health data, including by setting standards for data exchange and terminology (OECD, 2016_[3]). The European Commission’s Recommendation on a European Electronic Health Record exchange format is a step in that direction (European Commission, 2019_[60]). Sharing data across jurisdictions is central to advance AI research in areas such as cancer and rare diseases, as it requires sufficiently large, representative and complete data (and could potentially reduce the carbon footprint

⁸ AI investments are often not transferable because of the high diversity of data exchange and terminology standards. Because of this, investments are wasted, with, for example, each hospital re-inventing similar tools.

associated with AI). Cross-border data sharing is also crucial during pandemics (e.g. COVID-19), when infection spreads globally and concerted action is needed.

47. The latest evidence suggests that countries are lagging in their implementation of robust, consistent governance frameworks for the use of personal health data (OECD, 2019^[59]). Given the fundamental importance of good data in AI, failure to implement strong governance models will hinder and ultimately stall the potential benefits of this powerful technology. There is **need for global coordination** in this regard, as recommended in the OECD Principles.

48. Harmonisation and **interoperability of the laws and policies governing health data** enables cross-country collaboration. OECD countries are divided regarding legal permission to share data across borders for research and statistical uses in the public interest, even if the data are de-identified first. Among 18 high-income countries, only seven had laws and policies that could permit the majority of national health datasets to be shared with a foreign researcher working in the non-profit or governmental sectors for a project within the public interest (OECD, 2019^[59]). The European Union's *General Data Protection Regulation (GDPR)* fosters improvement in this regard among EU countries. The *GDPR* is a central feature of the Union's ambition to make health data more structured, interoperable, portable across borders, secure and respectful of individual's privacy. These laws and policies have the potential to advance the availability of high-quality, representative data for the development of AI models. Importantly, the *GDPR* puts health data in a special category that can be used for secondary purposes such as research deemed to be in the public interest and sets out the conditions for approval of cross-border data sharing. The *GDPR* has the potential to influence legal reforms outside the EU, particularly among countries seeking research partnerships and data exchanges with EU countries, and thus it has the potential to promote harmonisation in laws and policies regarding cross-border sharing beyond the EU.

49. In parallel to improving health data governance, countries must ensure that **digital infrastructure and technological capacity** keeps pace with the growing technical demands of AI. Governments can help health care providers, academics and technology companies (especially small and medium enterprises) access the specialised and expensive resources that are often needed to train AI systems. Small and medium enterprises may be given access to computing capacities and cloud platforms (OECD, 2019^[4]), for example through a one-stop shop for high performance computing (HPC) services (OECD, 2018^[61]). One way in which resource- or scale-constrained countries can access and provide HPC services is through multilateral cooperation. For example, the European High Performance Computing Joint Undertaking (EuroHPC JU) pools resources from the European Union and 32 countries to build a world-class supercomputing and data infrastructure, and to develop a competitive HPC ecosystem in relevant technologies and applications. Already, Exscalate4CoV, a European project, is using supercomputing resources to search for potential drugs to fight COVID-19.

3.2. Operationalising the OECD AI Principles will be challenging but fundamental

50. Agreement on the values-based elements of the AI Principles was an important achievement but it represents only the start of a longer journey. Operationalising them consistently across countries will be challenging in terms of transparency, accountability, accuracy, security and safety, as well as equity and fairness.

51. In addition to the technical challenge of accuracy, AI actors in health should commit to **transparency and interpretability** and responsible disclosure regarding AI systems (Vollmer et al., 2020^[62]). Of particular relevance to health care is the principle that those affected by AI – adversely or otherwise – should be aware that AI was used, and be able to understand the outcome and potentially challenge it, based on plain and easy-to-understand information on the factors, and the logic that served as the basis for the prediction, recommendation or decision (OECD, 2019^[63]). Implementing this in

practice can be technically challenging, as illustrated in the following excerpt from a recent report of the United States National Academy of Medicine (NAM, 2019^[38]):

Many contemporary AI systems are deployed on cloud-based, geographically distributed, nondeterministically parallelized, spatially arbitrary computing architectures that, at any moment, are physically unidentifiable. To create and maintain a full log of each processor that contributed in some part to the execution of a multi-element ensemble model AI is possible in principle but would likely be cost-prohibitive and too cumbersome to be practical. Therefore, the limited traceability and fundamental non-recreatability and non-retestability of a patient's or clinician's specific execution of an AI system that may have contained a fault or that produced errors or failures—untoward, unexpected deviations from its specifications, validation testing, and hazard analysis—may pose particular problems for regulators, courts, developers, and the public.

52. AI actors should also be accountable for the proper functioning of their algorithms, within the scope of their own roles. Legal clarification regarding **accountability** and responsibility for AI model outputs is important. The European Union recently published a report stating that manufacturers of products or digital content incorporating emerging digital technology should be **liable for damage** caused by defects in their products. Importantly, this applies if the defect was caused by changes made to the product under the producer's control after it had been placed on the market, for example if algorithms are updated on a regular basis (European Union, 2019^[64]). Similarly, in 2018, the United States Food and Drug Administration (FDA) approved the first algorithm that can make a medical decision without the need for a physician to look at a diagnostic image (Reardon, 2019^[42]). Because no doctor is involved, the company that developed the algorithm has assumed legal liability for any adverse outcomes.

53. Ensuring the **robustness, security and safety** of AI algorithms and applications is paramount, and the FDA has recently proposed a set of guidelines to manage algorithms that evolve over time. Among them is an expectation that developers **monitor how their algorithms are changing** to ensure they continue to work as designed and notify the agency if they see unexpected changes that might require re-assessment (Reardon, 2019^[42]).

54. Even ideological differences can present obstacles to guaranteeing AI in health is guided by **human-centred values and fairness**. For example, some advocate that personal health data are like any commodity owned by the data subject who should have the freedom to sell or exchange them. While the question of ownership can be debated, there is little doubt that such commodification of health data will incentivise poorer, disadvantaged people to sell theirs. Ethical underpinnings of such a policy position aside, purely from a technical standpoint this will create sample bias in the data used to train AI models. While this could increase representation of patients of lower socio-economic status in AI algorithms, there are other ways to increase representation that do not involve having a group of patients sell their medical data.

3.3. Putting in place regulation and guidance that promote trustworthy AI

55. AI is new territory for health policy makers, providers and the public. In terms of fostering trustworthy AI that delivers for patients and communities, an **enabling policy environment for AI** that includes a risk management approach is needed (in line with the AI Principle of robustness, security and safety). One method involves **regulatory sandboxes** – contained testing grounds for new approaches and business models. Regulatory sandboxes allow developers and health care providers to test and evaluate innovative products, services and business models in a live environment with the appropriate oversight and safeguards (ring-fencing wider health systems from risks and potential unintended consequences). The United Kingdom's Care Quality Commission and the Singaporean government are using regulatory sandboxes to test new (digital) health models.

56. Japan, for example, has developed AI Utilization Guidelines to enhance the outcomes of AI models (OECD, 2019^[4]). The Guidelines also specify that AI actors should create and publish an AI usage policy

and notify users, consumers, and others, so that they are aware of the use of AI. While Japan has traditionally been near the bottom of the pack in making personal health data available for secondary purposes, the government has recently set in motion legislative reforms to address this (OECD, 2019^[59]).

57. It is encouraging to see wide agreement regarding the need for AI principles and values, with at least 84 public-private initiatives describing high-level principles, values and other tenets to guide the ethical development, deployment and governance of AI (Jobin, Ienca and Vayena, 2019^[65]). However, a multitude of frameworks poses a risk to international cooperation. The onus is on countries to draw on a set of value statements to help develop and implement the necessary policies, **regulations and legal frameworks**. Consistency across jurisdictions will be to everybody's advantage, and the practical and technological challenges to several of the AI principles can perhaps be better overcome through **international co-operation for trustworthy AI**.

3.4. Building human capacity and preparing the workforce for the change

58. To date, there is **no evidence to suggest that AI will replace humans in health care**, but there is plenty to suggest that it will fundamentally change human tasks and augment skills and responsibilities. Given the scale at which AI could change the healthcare landscape, the way health workers – and indeed the entire workforce – are educated, trained and socialised will need to adapt. The approach will need to be multidisciplinary, involving AI developers, implementers, health care system leadership, frontline clinical teams, ethicists, humanists, and patients and patient caregivers, as each provides their own point of view and lived experience – each one should be heard (NAM, 2019^[38]). Health care workers – especially clinical staff – who make decisions based on AI algorithms, should receive practical training in how to use these products, in the same way that health care professionals are trained in the use of digital health technologies.

59. **New jobs and professions** will be needed to realise the potential benefits of AI in health: trainers, explainers and sustainers (Wilson, Daugherty and Morini-Bianzino, 2017^[66]). Trainers will provide meaning, purpose, and direction; explainers will use their knowledge in both the technical and application domains to explain how AI algorithms can be trusted to support decisions; and sustainers will help maintain, interpret, and monitor the behaviour and unintended consequences of AI systems.

60. A number of **countries are already preparing**. In France, a collaboration between Gustave Roussy, one of Europe's leading cancer hospitals, and two engineering schools in Paris, École des Ponts ParisTech and CentraleSupélec, aims to train young computer scientists to understand medicine, and conversely, to train medical researchers to better understand the basics of artificial intelligence. In the United States, about a dozen fellowships are offered to train budding doctors in a range of engineering approaches, including artificial intelligence. Australia, Canada, Norway, Switzerland, New Zealand, and the UK have all completed reviews or established regular processes to assess how technological developments will change skill requirements, roles and functions of health care staff.

3.5. Investing strategically and sustainably in AI research and development

61. Preparing health systems to manage the risks and make the most of AI requires long-term investment. Strategic, **coordinated and sustained resourcing is needed** to ensure that AI leads to desirable health, social and economic outcomes and takes a similar trajectory to successful industrial revolutions of the past. Public resources are and always will be scarce, but they need to be found given the profound opportunities for better and more equitable health outcomes brought by AI.

62. Private capital is piling into the AI field (OECD, 2018^[67]). **Private investment** doubled from 2016 to 2017, reaching USD 16 billion in 2017. AI start-ups attracted 12% of worldwide private equity

investments in the first half of 2018. For example, in pharma, AI-based drug discovery start-ups raised more than \$1 billion in funding in 2018. At least 20 separate partnerships have been reported between major pharma companies and AI-drug-discovery tech companies (Freedman, 2019^[68]). Pfizer, GlaxoSmithKline and Novartis are among the pharma companies said to have also built substantial in-house AI expertise. Besides potentially developing AI tools themselves, governments and public institutions should also devote resources towards developing the guardrails that ensure this technology does maximum good and minimise harm – and the checks and balances to steer the private sector in the right direction. This includes establishing and maintaining sound policy frameworks, building institutional and human capacity to verify AI technology where needed, and use it in an appropriate and cost-effective manner. Indeed, it is challenging for the public sector to attract AI talent.

63. Gauging the economic benefits of AI, and its superiority over conventional and cheaper techniques must also be an important consideration given the comparative costliness of AI. A recent systematic review found that very few studies of AI's use in the health sector assess **economic impact**, with none conducting a methodologically complete cost impact analysis (Wolff et al., 2019^[69]). Public investment is needed to measure the operational costs of AI infrastructure and services, and to compare the results to existing alternatives.

64. The prognosis is not favourable. National **health systems typically underinvest** in information systems, given the paramount importance of information, communication and knowledge in this sector (OECD, 2019^[59]). The COVID-19 pandemic is revealing important gaps and deficiencies in health data. Few countries have data that are timely enough to support decision making in real time. Legal and policy barriers to data accessibility, sharing and linkage further impede the management of the pandemic. Lack of health data standards for terminology and exchange limit the ability to share data for monitoring and research and for the development of AI tools. The absence of underlying standards also make it difficult to share AI tools among health organisations within countries and are significant barriers to cross-border sharing of data and tools. Plainly put, fiscal space to invest in guiding AI in health must be found. Fortunately, intelligent deployment of AI (and digital technology more broadly) provides opportunities to make existing clinical and administrative processes more effective, efficient and equitable. With around one fifth of health spending being wasteful or even harmful (OECD, 2017^[1]), this is a major opportunity to improve outcomes and value. In the medium-term, the investment may pay for itself.

4 Conclusion

65. Artificial intelligence is a tool with the potential to transform health sector activities. At present, AI is applied to various degrees and success in clinical practice, public health, biomedical research and health care administration. Often these applications are still in the research and development stage. Most are based on AI “narrow intelligence”, meaning that the tools are quite simple and limited in their ability to inform decisions and recommend actions, despite the complexity of the computer programmes underpinning them.

66. AI may progress in the future to more human-like reasoning and prediction abilities, however policy makers should **beware of hype regarding current AI applications** (NAM, 2019^[38]). Rather than approaching AI as a panacea to all challenges in health and health care, policy makers are advised to **identify and focus on real problems and opportunities** where AI could help and address the fundamental barriers to its successful development and adoption in the health industry. Policy makers can do this by ensuring quality data are available and secure, implementing and operationalising the OECD AI Principles, and investing in technology and human capital.

67. A key to realising and sharing useful AI tools in health care is to improve data quality, interoperability and access in a secure way. AI applications require large, up-to-date datasets to be trained on. As the data become narrower or more distant from the problem, the utility of the AI decreases and the probability of biases increases. The COVID-19 pandemic has showcased new AI tools that are increasingly accepted by health care practitioners and the public. At the same time, these new tools have raised concerns about the utility of AI developed with small, limited datasets that may not be applicable to diverse populations or to future time periods when the pandemic is under control.

68. Ensuring data quality, availability and security can be achieved by implementing the **OECD Health Data Governance Recommendation**. This Recommendation sets out principles for national health data governance frameworks that improve data quality and accessibility while protecting privacy and data security. The Recommendation calls on governments to address unnecessary barriers to the efficient exchange and interoperability of health data, particularly those that are blocking public-private and cross-border monitoring and research. The Recommendation calls for international cooperation to develop global standards for data exchange and data terminology; and harmonising health data governance frameworks that protect data privacy and security.

69. More broadly, the **OECD AI Principles** need to be operationalised in the health sector, which could inform any assessed need for industry-specific policy or regulatory reforms. Important steps include ensuring transparency, interpretability and accountability of AI outputs; regulatory oversight and guidance that encourages innovation in trustworthy AI; and, building human capacity to utilise AI tools among health workers and also patients and the public.

70. Finally, investments are needed in most countries to develop the data and human resources necessary to use AI tools in the public interest. Operationalising data governance and AI Principles will also require international collaboration to harmonise approaches to health data development, use and governance and to **ensure that the application of AI is ethical and helpful**, with benefits that outweigh its costs.

References

- Accenture (2018), *The Intelligent Payer: A survival guide*, [39]
https://www.accenture.com/_acnmedia/pdf-82/accenture-intelligent-payer-survivor-guide.pdf.
- ACR (2020), *ACR Recommendations for the use of Chest Radiography and Computed Tomography (CT) for Suspected COVID-19 Infection | American College of Radiology*, [16]
<https://www.acr.org/Advocacy-and-Economics/ACR-Position-Statements/Recommendations-for-Chest-Radiography-and-CT-for-Suspected-COVID19-Infection> (accessed on 14 May 2020).
- AEI (2020), *Erik Brynjolfsson: Can AI help us overcome the productivity paradox?*, [71]
<https://www.aei.org/multimedia/erik-brynjolfsson-can-ai-help-us-overcome-the-productivity-paradox/> (accessed on 4 March 2020).
- Beck, B. et al. (2020), “Predicting commercially available antiviral drugs that may act on the novel coronavirus (2019-nCoV), Wuhan, China through a drug-target interaction deep learning model”, *bioRxiv*, p. 2020.01.31.929547, [35]
<http://dx.doi.org/10.1101/2020.01.31.929547>.
- Boston Dynamics (2020), *Blog Post | Boston Dynamics*, [19]
<https://www.bostondynamics.com/COVID-19> (accessed on 14 May 2020).
- Brzezicki, M. et al. (2020), “Artificial intelligence outperforms human students in conducting neurosurgical audits”, *Clinical Neurology and Neurosurgery*, Vol. 192, p. 105732, [20]
<http://dx.doi.org/10.1016/j.clineuro.2020.105732>.
- Cho, A. (2020), “Artificial intelligence systems aim to sniff out signs of COVID-19 outbreaks”, *Science*, [52]
<http://dx.doi.org/10.1126/science.abc7698>.
- Combs, V. (2019), *South African clinics use artificial intelligence to expand HIV treatment - TechRepublic*, [8]
<https://www.techrepublic.com/article/south-african-clinics-use-artificial-intelligence-to-expand-hiv-treatment/> (accessed on 10 March 2020).
- Cresswell, K. et al. (2020), “Investigating the use of data-driven artificial intelligence in computerised decision support systems for health and social care: A systematic review”, *Health Informatics Journal*, p. 146045821990045, [51]
<http://dx.doi.org/10.1177/1460458219900452>.
- David, P. (1990), *The dynamo and the computer: An historical perspective on the modern productivity paradox*, American Economic Association, [73]
<http://dx.doi.org/10.2307/2006600>.

- DeepMind (2020), *Computational predictions of protein structures associated with COVID-19* | DeepMind, <https://deepmind.com/research/open-source/computational-predictions-of-protein-structures-associated-with-COVID-19> (accessed on 14 May 2020). [36]
- DeepMind (2016), *DeepMind AI Reduces Google Data Centre Cooling Bill by 40%*, <https://deepmind.com/blog/article/deepmind-ai-reduces-google-data-centre-cooling-bill-40> (accessed on 9 March 2020). [56]
- Delcker, J. (2020), *Donate data to health authorities to fight virus, says German epidemiologist*, Politico, <https://www.politico.eu/article/coronavirus-hand-over-data-to-health-authorities-to-fight-virus-says-german-epidemiologist/>. [33]
- Dyrbye, L. et al. (2017), "Burnout Among Health Care Professionals: A Call to Explore and Address This Underrecognized Threat to Safe, High-Quality Care", *NAM Perspectives*, Vol. 7/7, <http://dx.doi.org/10.31478/201707b>. [46]
- Eichler, H. et al. (2018), "Data Rich, Information Poor: Can We Use Electronic Health Records to Create a Learning Healthcare System for Pharmaceuticals?", *Clinical Pharmacology and Therapeutics*, Vol. 105/4, pp. 912-922, <http://dx.doi.org/10.1002/cpt.1226>. [25]
- European Commission (2019), *LeanLab (PlanDomino)*, <https://ec.europa.eu/eipp/desktop/en/projects/project-11297.html> (accessed on 3 November 2020). [41]
- European Commission (2019), *Recommendation on a European Electronic Health Record exchange format | Shaping Europe's digital future*, <https://ec.europa.eu/digital-single-market/en/news/recommendation-european-electronic-health-record-exchange-format> (accessed on 4 November 2020). [60]
- European Union (2019), *Liability for Artificial Intelligence and other emerging digital technologies: Report from the Expert Group on Liability and New Technologies – New Technologies Formation*, <http://dx.doi.org/10.2838/573689>. [64]
- FDA (2019), *FDA approves new treatments for heart disease caused by a serious rare disease, transthyretin mediated amyloidosis*, <https://www.fda.gov/news-events/press-announcements/fda-approves-new-treatments-heart-disease-caused-serious-rare-disease-transthyretin-mediated> (accessed on 4 March 2020). [70]
- Freedman, D. (2019), *Hunting for New Drugs with AI*, Nature Research, <http://dx.doi.org/10.1038/d41586-019-03846-0>. [68]
- Garside, P. (2004), "Are we suffering from change fatigue?", *Quality & safety in health care*, Vol. 13/2, pp. 89-90, <http://dx.doi.org/10.1136/qshc.2003.009159>. [47]
- GE Healthcare (2015), *GE Healthcare and Humber River Hospital launch first in North America Managed Equipment Service solution*, <https://www.ge.com/news/press-releases/ge-healthcare-and-humber-river-hospital-launch-first-north-america-managed-equipment> (accessed on 3 November 2020). [75]
- GE Healthcare Partners (2018), *Humber River Hospital and GE Healthcare Building First Hospital Command Centre for Quality Healthcare in Canada*, <https://emea.gehealthcarepartners.com/insights/156-latest-news/521-humber-river-hospital-and-ge-healthcare-building-first-hospital-command-centre-for-quality-healthcare-in-canada> (accessed on 3 November 2020). [40]

- Hao, K. (2020), *This is how the CDC is trying to forecast coronavirus's spread* | MIT Technology Review, MIT Technology Review, <https://www.technologyreview.com/2020/03/13/905313/cdc-cmu-forecasts-coronavirus-spread/> (accessed on 14 May 2020). [22]
- Hao, K. (2019), *Training a single AI model can emit as much carbon as five cars in their lifetimes*, MIT Technology Review, https://www.technologyreview.com/s/613630/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/?utm_source=newsletters&utm_medium=email&utm_campaign=the_algorithm.unpaid.engagement (accessed on 4 March 2020). [53]
- HBR (2020), *How Hospitals Are Using AI to Battle Covid-19*, <https://hbr.org/2020/04/how-hospitals-are-using-ai-to-battle-covid-19> (accessed on 14 May 2020). [18]
- Heaven, W. (2020), *A neural network can help spot Covid-19 in chest x-rays* | MIT Technology Review, MIT Technology Review, <https://www.technologyreview.com/2020/03/24/950356/coronavirus-neural-network-can-help-spot-covid-19-in-chest-x-ray-pneumonia/> (accessed on 14 May 2020). [32]
- Hope, M. et al. (2020), *A role for CT in COVID-19? What data really tell us so far*, Lancet Publishing Group, [http://dx.doi.org/10.1016/S0140-6736\(20\)30728-5](http://dx.doi.org/10.1016/S0140-6736(20)30728-5). [17]
- Hwang, E. et al. (2019), "Development and Validation of a Deep Learning-Based Automated Detection Algorithm for Major Thoracic Diseases on Chest Radiographs", *JAMA network open*, Vol. 2/3, p. e191095, <http://dx.doi.org/10.1001/jamanetworkopen.2019.1095>. [11]
- ITN (2020), *FDA Approves Use of Aidoc's AI Algorithms for Incidental CT Findings Associated with COVID-19* | *Imaging Technology News*, <https://www.itnonline.com/content/fda-approves-use-aidoc%E2%80%99s-ai-algorithms-incident-ct-findings-associated-covid-19> (accessed on 14 May 2020). [13]
- Jobin, A., M. Ienca and E. Vayena (2019), "The global landscape of AI ethics guidelines", *Nature Machine Intelligence*, Vol. 1/9, pp. 389-399, <http://dx.doi.org/10.1038/s42256-019-0088-2>. [65]
- Lee, C. and A. Lee (2020), "How Artificial Intelligence Can Transform Randomized Controlled Trials", *Translational Vision Science & Technology*, Vol. 9/2, p. 9, <http://dx.doi.org/10.1167/tvst.9.2.9>. [28]
- Lehne, M. et al. (2019), "Why digital medicine depends on interoperability", *npj Digital Medicine*, Vol. 2/1, pp. 1-5, <http://dx.doi.org/10.1038/s41746-019-0158-1>. [48]
- Li, L. et al. (2020), "Artificial Intelligence Distinguishes COVID-19 from Community Acquired Pneumonia on Chest CT", *Radiology*, p. 200905, <http://dx.doi.org/10.1148/radiol.2020200905>. [12]
- Lunit (2020), *Lunit Releases Its AI Online to Support Healthcare Professionals Manage COVID-19*, <https://lunit.prezly.com/lunit-releases-its-ai-online-to-support-healthcare-professionals-manage-covid-19#> (accessed on 14 May 2020). [10]
- Masanet, E. et al. (2020), "Recalibrating global data center energy-use estimates", *Science*, Vol. 367/6481, pp. 984-986, <http://dx.doi.org/10.1126/science.aba3758>. [55]
- Matheny, M. et al. (eds.) (2019), *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril*, National Academy of Medicine, Washington DC. [38]

- Mccall, B. (2020), “News COVID-19 and artificial intelligence : protecting health-care workers and curbing the spread”, *The Lancet*, Vol. 2019/20, pp. 2019-2020, [http://dx.doi.org/10.1016/S2589-7500\(20\)30054-6](http://dx.doi.org/10.1016/S2589-7500(20)30054-6). [34]
- Naudé, W. (2020), “Artificial Intelligence against COVID-19: An Early Review”, *IZA Discussion Papers*, No. 13110, IZA - Institute of Labor Economics, <https://www.iza.org/publications/dp/13110> (accessed on 14 May 2020). [24]
- Neri, E. et al. (2019), “What the radiologist should know about artificial intelligence – an ESR white paper”, *Insights into Imaging*, Vol. 10/1, p. 44, <http://dx.doi.org/10.1186/s13244-019-0738-2>. [5]
- Nolan, A. (2020), “Why accelerate the development and deployment of robots?”, in *OECD Science, Technology and Innovation Outlook 2020 (forthcoming)*, OECD Publishing, Paris. [76]
- Obermeyer, Z. et al. (2019), “Dissecting racial bias in an algorithm used to manage the health of populations”, *Science*, Vol. 366/6464, pp. 447-453, <http://dx.doi.org/10.1126/science.aax2342>. [43]
- Oderkirk, J. (2017), *Readiness of Electronic Health Record Systems to Contribute to National Health Information and Research*, <http://dx.doi.org/10.1787/9e296bf3-en>. [49]
- OECD (2019), *Artificial Intelligence in Society*, OECD Publishing, Paris, <https://dx.doi.org/10.1787/eedfee77-en>. [4]
- OECD (2019), *Health in the 21st Century: Putting Data to Work for Stronger Health Systems*, OECD Health Policy Studies, OECD Publishing, Paris, <https://dx.doi.org/10.1787/e3b23f8e-en>. [59]
- OECD (2019), *OECD Employment Outlook 2019: The Future of Work*, OECD Publishing, Paris, <https://dx.doi.org/10.1787/9ee00155-en>. [2]
- OECD (2019), *Recommendation of the Council on Artificial Intelligence*, OECD/LEGAL/0449. [63]
- OECD (2018), *OECD Science, Technology and Innovation Outlook 2018: Adapting to Technological and Societal Disruption*, OECD Publishing, Paris, https://dx.doi.org/10.1787/sti_in_outlook-2018-en. [61]
- OECD (2018), *Private Equity Investment in Artificial Intelligence: OECD Going Digital Policy Note*, OECD, Paris, <http://www.oecd.org/going-digital/ai/private-equity-investment-in-artificial-intelligence.pdf>. [67]
- OECD (2017), *Tackling Wasteful Spending on Health*, OECD Publishing, Paris, <http://www.oecd-ilibrary.org/docserver/download/8116241e.pdf?expires=1518450288&id=id&accname=ocid84004878&checksum=8647E938E2C1B896ECB03B16256A576B> (accessed on 12 February 2018). [1]
- OECD (2016), *Recommendation of the Council on Health Data Governance*, OECD/LEGAL/0433, <http://legalinstruments.oecd.org> (accessed on 7 May 2019). [3]
- Olson, D. et al. (2013), “Reassessing Google Flu Trends Data for Detection of Seasonal and Pandemic Influenza: A Comparative Epidemiological Study at Three Geographic Scales”, *PLoS Computational Biology*, Vol. 9/10, p. e1003256, <http://dx.doi.org/10.1371/journal.pcbi.1003256>. [23]

- Olson, P. (2020), *AI Software Gets Mixed Reviews for Tackling Coronavirus - WSJ*, The Wall Street Journal, <https://www.wsj.com/articles/ai-software-gets-mixed-reviews-for-tackling-coronavirus-11588597013?mod=searchresults&page=1&pos=2> (accessed on 14 May 2020). [15]
- Pichler, P. et al. (2019), "International comparison of health care carbon footprints", *Environmental Research Letters*, <http://dx.doi.org/10.1088/1748-9326/ab19e1>. [57]
- Rahimian, F. et al. (2018), "Predicting the risk of emergency admission with machine learning: Development and validation using linked electronic health records", *PLoS Medicine*, Vol. 15/11, <http://dx.doi.org/10.1371/journal.pmed.1002695>. [29]
- Reardon, S. (2019), *Rise of Robot Radiologists*, Nature Research, <http://dx.doi.org/10.1038/d41586-019-03847-z>. [42]
- SAS (2020), *What is deep learning? | SAS*, https://www.sas.com/en_us/insights/analytics/deep-learning.html (accessed on 14 May 2020). [74]
- Schmider, J. et al. (2019), "Innovation in Pharmacovigilance: Use of Artificial Intelligence in Adverse Event Case Processing", *Clinical Pharmacology & Therapeutics*, Vol. 105/4, pp. 954-961, <http://dx.doi.org/10.1002/cpt.1255>. [21]
- SemanticScholar (2020), *CORD-19 | Semantic Scholar*, <https://www.semanticscholar.org/cord19> (accessed on 14 May 2020). [31]
- Shadmi, E. et al. (2015), "Predicting 30-day readmissions with preadmission electronic health record data", *Medical Care*, Vol. 53/3, pp. 283-289, <http://dx.doi.org/10.1097/MLR.0000000000000315>. [7]
- Skrede, O. et al. (2020), "Deep learning for prediction of colorectal cancer outcome: a discovery and validation study", *The Lancet*, Vol. 395/10221, pp. 350-360, [http://dx.doi.org/10.1016/S0140-6736\(19\)32998-8](http://dx.doi.org/10.1016/S0140-6736(19)32998-8). [9]
- Stokes, J. et al. (2020), "A Deep Learning Approach to Antibiotic Discovery", *Cell*, Vol. 180/4, pp. 688-702.e13, <http://dx.doi.org/10.1016/j.cell.2020.01.021>. [26]
- Strickland, E. (2019), *How IBM Watson Overpromised and Underdelivered on AI Health Care - IEEE Spectrum*, IEEE Spectrum, <https://spectrum.ieee.org/biomedical/diagnostics/how-ibm-watson-overpromised-and-underdelivered-on-ai-health-care> (accessed on 4 March 2020). [44]
- Strubell, E., A. Ganesh and A. McCallum (2019), "Energy and Policy Considerations for Deep Learning in NLP", pp. 3645-3650, <http://arxiv.org/abs/1906.02243> (accessed on 4 March 2020). [54]
- Svoboda, E. (2019), *Your robot surgeon will see you now*, NLM (Medline), <http://dx.doi.org/10.1038/d41586-019-02874-0>. [6]
- Tiwari, P. et al. (2020), "Assessment of a Machine Learning Model Applied to Harmonized Electronic Health Record Data for the Prediction of Incident Atrial Fibrillation", *JAMA network open*, Vol. 3/1, p. e1919396, <http://dx.doi.org/10.1001/jamanetworkopen.2019.19396>. [72]
- Vollmer, S. et al. (2020), "Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness", *The BMJ*, Vol. 368, <http://dx.doi.org/10.1136/bmj.l6927>. [62]

- West, S., M. Whittaker and K. Crawford (eds.) (2019), *Discriminating Systems: Gender, Race and Power in AI*, <https://ainowinstitute.org/discriminatingystems.html>. [45]
- WHO (2014), *Quantitative risk assessment of the effects of climate change on selected causes of death, 2030s and 2050s*, World Health Organization. [58]
- Wilson, J., P. Daugherty and N. Morini-Bianzino (2017), “The Jobs That Artificial Intelligence Will Create”, *MIT Sloan Management Review*. [66]
- Wolff, J. et al. (2019), “A Systematic Review of Economic Impact Studies of Artificial Intelligence in Healthcare (Preprint)”, *Journal of Medical Internet Research*, Vol. 22/2, p. e16866, <http://dx.doi.org/10.2196/16866>. [69]
- Woo, M. (2019), *An AI boost for clinical trials*, Nature Publishing Group, <http://dx.doi.org/10.1038/d41586-019-02871-3>. [37]
- Wynants, L. et al. (2020), “Prediction models for diagnosis and prognosis of covid-19 infection: Systematic review and critical appraisal”, *The BMJ*, Vol. 369, <http://dx.doi.org/10.1136/bmj.m1328>. [50]
- Yan, L. et al. (2020), “An interpretable mortality prediction model for COVID-19 patients”, *Nature Machine Intelligence*, pp. 1-6, <http://dx.doi.org/10.1038/s42256-020-0180-7>. [30]
- Zhang, K. et al. (2020), “Clinically Applicable AI System for Accurate Diagnosis, Quantitative Measurements and Prognosis of COVID-19 Pneumonia Using Computed Tomography”, *Cell*, <http://dx.doi.org/10.1016/j.cell.2020.04.045>. [14]
- Zhavoronkov, A. et al. (2019), “Deep learning enables rapid identification of potent DDR1 kinase inhibitors”, *Nature Biotechnology*, Vol. 37/9, pp. 1038-1040, <http://dx.doi.org/10.1038/s41587-019-0224-x>. [27]

OECD Health Working Papers

A full list of the papers in this series can be found on the OECD website:

<http://www.oecd.org/els/health-systems/health-working-papers.htm>

No. 127 – SURVEY RESULTS: NATIONAL HEALTH DATA INFRASTRUCTURE AND GOVERNANCE (April 2021) Jillian Oderkirk

No. 126 – INTERNATIONAL MIGRATION AND MOVEMENT OF DOCTORS TO AND WITHIN OECD COUNTRIES - 2000 TO 2018 - DEVELOPMENTS IN COUNTRIES OF DESTINATION AND IMPACT ON COUNTRIES OF ORIGIN (February 2021) Karolina Socha-Dietrich and Jean-Christophe Dumont

No. 125 – INTERNATIONAL MIGRATION AND MOVEMENT OF NURSING PERSONNEL TO AND WITHIN OECD COUNTRIES - 2000 TO 2018 - DEVELOPMENTS IN COUNTRIES OF DESTINATION AND IMPACT ON COUNTRIES OF ORIGIN (February 2021) Karolina Socha-Dietrich and Jean-Christophe Dumont

No. 124 – SKILLS FOR THE FUTURE HEALTH WORKFORCE - PREPARING HEALTH PROFESSIONALS FOR PEOPLE-CENTRED CARE (February 2021) Akiko Maeda and Karolina Socha-Dietrich

No. 123 - CHALLENGES IN ACCESS TO ONCOLOGY MEDICINES: POLICIES AND PRACTICES ACROSS THE OECD AND THE EU (November 2020) Suzannah Chapman, Valérie Paris and Ruth Lopert

No. 122 - EXCESS MORTALITY: MEASURING THE DIRECT AND INDIRECT IMPACT OF COVID-19 (October 2020) David Morgan, Junya Ino, Gabriel Di Paolantonio and Fabrice Murtin

No. 121 – THE ECONOMICS OF PATIENT SAFETY PART III: LONG-TERM CARE - VALUING SAFETY FOR THE LONG HAUL (September 2020) Katherine de Bienassis, Ana Llana-Nozal and Nicolaas S. Klazinga

No. 120 – SYSTEM GOVERNANCE TOWARDS IMPROVED PATIENT SAFETY - KEY FUNCTIONS, APPROACHES AND PATHWAYS TO IMPLEMENTATION (September 2020) Ane Auraaen, Kristin Saar and Nicolaas S. Klazinga

No. 119 – CULTURE AS A CURE: ASSESSMENTS OF PATIENT SAFETY CULTURE IN OECD COUNTRIES Katherine de Bienassis, Solvejg Kristensen, Magdalena Burtscher, Ian Brownwood and Nicolaas S. Klazinga.

No. 118 REASSESSING PRIVATE PRACTICE IN PUBLIC HOSPITALS IN IRELAND: AN OVERVIEW OF OECD EXPERIENCES Michael Mueller and Karolina Socha-Dietrich.

No. 117 - THE EFFECTIVENESS OF SOCIAL PROTECTION FOR LONG-TERM CARE IN OLD AGE (May 2020) Tiago Cravo Oliveira Hashiguchi and Ana Llana-Nozal

No. 116 - BRINGING HEALTH CARE TO THE PATIENT: AN OVERVIEW OF THE USE OF TELEMEDICINE IN OECD COUNTRIES (January 2020) Tiago Cravo Oliveira Hashiguchi

Recent related OECD publications

PREVENTING HARMFUL ALCOHOL USE (2021)

OECD REVIEWS OF PUBLIC HEALTH: LATVIA (2020)

HEALTH AT A GLANCE: EUROPE (2020)

HEALTH AT A GLANCE: ASIA/PACIFIC (2020)

EMPOWERING THE HEALTH WORKFORCE - STRATEGIES TO MAKE THE MOST OF THE DIGITAL REVOLUTION (2020)

HEALTH AT A GLANCE: LATIN AMERICA AND THE CARIBBEAN (2020)

WHO CARES? ATTRACTING AND RETAINING CARE WORKERS FOR THE ELDERLY (2020)

REALISING THE POTENTIAL OF PRIMARY HEALTH CARE (2020)

WAITING TIMES FOR HEALTH SERVICES: NEXT IN LINE (2020)

IS CARDIOVASCULAR DISEASE SLOWING IMPROVEMENTS IN LIFE EXPECTANCY? OECD AND THE KING'S FUND WORKSHOP PROCEEDINGS (2020)

ADDRESSING CHALLENGES IN ACCESS TO ONCOLOGY MEDICINES (2020)

OECD REVIEWS OF PUBLIC HEALTH: KOREA - A HEALTHIER TOMORROW (2020)

OECD HEALTH STATISTICS 2020 - Online Database available from:

<https://www.oecd.org/health/health-statistics.htm>

COUNTRY HEALTH PROFILES (2019)

HEALTH IN THE 21ST CENTURY: PUTTING DATA TO WORK FOR STRONGER HEALTH SYSTEMS (2019)

THE SUPPLY OF MEDICAL ISOTOPES: AN ECONOMIC DIAGNOSIS AND POSSIBLE SOLUTIONS (2019)

HEALTH AT A GLANCE (2019)

THE HEAVY BURDEN OF OBESITY – THE ECONOMICS OF PREVENTION (2019)

HEALTH FOR EVERYONE? - SOCIAL INEQUALITIES IN HEALTH AND HEALTH SYSTEMS (2019)

RECENT TRENDS IN INTERNATIONAL MIGRATION OF DOCTORS, NURSES AND MEDICAL STUDENTS (2019)

PRICE SETTING AND PRICE REGULATION IN HEALTH CARE (2019) OECD/WHO Centre for Health Development in Kobe

ADDRESSING PROBLEMATIC OPIOIDS USE IN OECD COUNTRIES (2019)

OECD REVIEW OF PUBLIC HEALTH: JAPAN (2019)

OECD REVIEW OF PUBLIC HEALTH: CHILE (2019)

For a full list, consult the OECD health web page at <http://www.oecd.org/health/>