

LITERATURE REVIEW

Current and emerging capabilities of AI-powered genomics, and associated ethical, legal and political debates

By

Dr Arianna Manzini

and

Timothy Lee, University of Edinburgh

April 2023

Commissioned by the Ada Lovelace Institute and the Nuffield Council on Bioethics for their joint research project 'AI and genomics futures'.

AI and genomics futures

A literature review of current and emerging capabilities of AI powered genomics, and associated ethical, legal and political debates

**Dr Arianna Manzini
Timothy Lee - University of Edinburgh**

Introduction	3
Part 1: Current and emerging genomic capabilities enabled by AI	4
1.1 Drug Discovery	5
1.2 Pharmacogenomics.....	7
1.3 Cancer	8
1.4 Diagnostics	9
1.5 Neurology/ Ageing	10
1.6 Image to genome mapping	11
1.7 Radiology.....	11
1.8 MicroRNAs	12
1.9 Sequencing.....	12
1.10 Single Cell Genomics	13
1.11 Data Sharing.....	13
1.12 Non-healthcare related applications	14
1.13 Emerging approaches and predictions	14
Part 2: Ethical, societal and legal debates concerning genomic science.....	16
2.1 Data, code and software engineering: Privacy and discrimination	17
2.2 Data and user bias: Misclassification and fairness.....	21
2.3 Explainability and interpretability: Trust and the right to an explanation	24
2.4 Changes in genomic expertise: Workforce diversity	25
2.5 Predictions of life outcomes: Genetic determinism and the spectrum of eugenics	26
2.6 Cost and opportunity cost.....	27
2.7 Macro level issues: Health disparities and concentrations of power	29
Conclusion	30
Annex 1: Methodology	33
Identifying the research questions	33
Defining the search string and selecting the sources	33
Defining criteria for literature selection and extracting data	36
Summarizing and reporting results.....	36
Annex 2: Glossary.....	38
Bibliography	39

Introduction

Claims about the potential of artificial intelligence to transform the development, application and practice of genomic science have become increasingly common in recent years. In particular, advances in techniques such as machine learning and deep learning are cited as responsible for substantial improvements in the collection, analysis and useful deployment of genomic data – and are predicted to yield further improvements in the future.

For practitioners and decisionmakers concerned with how we cultivate, manage and regulate genomics, **the emergence of AI powered genomic science raises a host of difficult, important questions:** Where, how and to what extent is AI currently changing the capabilities, viable applications and practice of genomic science? What future changes are anticipated – and how confident can we be of their emergence? And what might be the political, economic and societal impacts of these changes, should they come about?

This literature review was commissioned as part of *AI and Genomics Futures*, a joint project between the Ada Lovelace Institute and the Nuffield Council on Bioethics aimed at exploring and moving towards answers to these questions. Alongside a *horizon scanning exercise*¹ and a *scientometric analysis*², it constitutes one of a set of overlapping research activities that, when combined, are **intended to provide a clear, up to date picture of AI's current and anticipated applications to genomics.**

The review surveys current academic and grey literature with a focus on two areas: The first part of the review is devoted to collating and synthesizing academic literature on **how advances in AI are changing the capabilities, applications and practice of genomic science.** The second part considers existing literature focused on **the ethical, societal, legal and political ramifications of the changes AI is bringing about in the field of genomics.**

¹ which methodically elicits expert predictions about future capabilities and applications of AI powered genomics

² which uses data science techniques to identify macro level trends at the intersection of AI and genomics

Part 1: Current and emerging genomic capabilities enabled by AI

Genomics is the study of the function, structure, and interactions of the underlying genetic sequence of organisms. An individual's genome is a unique DNA sequence that contains their entire set of genetic instructions for development. The field was pioneered in the 1970s when the ability to extract DNA and obtain its exact sequence was established. Since then, advances in sequencing have led to large amounts of genetic information being generated.³

It is hoped that genomic analysis – the process of identifying, and comparing genomic features such as DNA sequences, structural variations, gene expression, or regulatory and functional element annotations at a genomic scale - will lead to rapid improvements in health care, especially in precision medicine for personalized therapies targeted to an individual's own genome. However, the human genome is over 3 billion nucleotide DNA bases long, and contains immense amounts of non-linear associations and interactions, making statistical analysis difficult and complex.

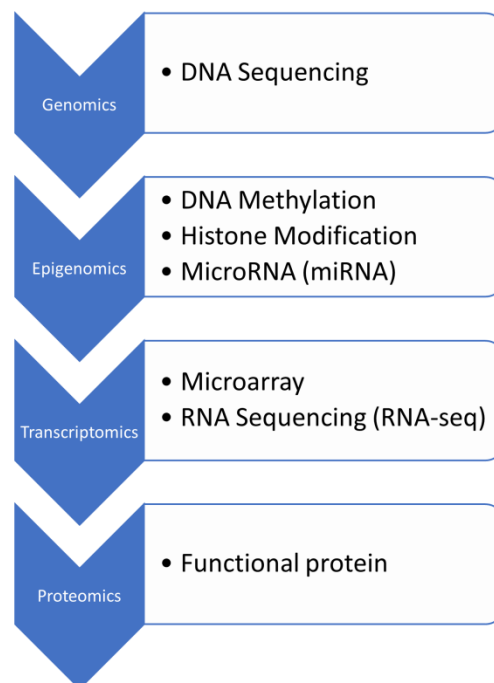


Figure 1: Sequential organization of genotype-to-phenotype process. Genomics is concerned with the underlying genome sequence of an individual. The genome sequence consists of DNA, which is transcribed to create RNA which is converted to functional protein. This functional protein is what dictates phenotype. Between the genome and transcription, there is an intermediate level of regulation which controls what transcripts are produced to be eventually translated into protein.

In response to this complexity, the field of genomics has expanded to using Artificial Intelligence (AI) and Machine Learning (ML) algorithms in order to analyse data sets to find underlying trends and associations within the genome. ML methods have the potential to resolve many of the problems that genome researchers have been facing in regard to the large amounts of complex data generated. These issues include the high dimensionality and

³ The rise of Next Generation Sequencing (NGS) is one such example of technologies responsible for the surge in genomic data generation. NGS, which is also sometimes called high-throughput sequencing, is a term used to describe various DNA sequencing platforms such as Illumina sequencing, Roche 454 sequencing, Ion Torrent, and SOLiD sequencing (F. Celesti et al. 2018).

heterogeneity of genomic data, the need for large datasets for training models, and the lack of transferability between datasets.⁴ ML approaches are employed in a variety of research fields that make use of genomic data such as cancer research, diagnostics, drug discovery, radiology, ageing, and image to genome mapping.

The noncoding portion of the genome, which hosts regulatory elements such as enhancers, promoters, splicing sites as well as prediction of microRNA (miRNA) precursors and targets for control of gene expression, has also been receiving much attention as of late with a myriad of Deep Learning (DL) tools being developed for use in those fields.

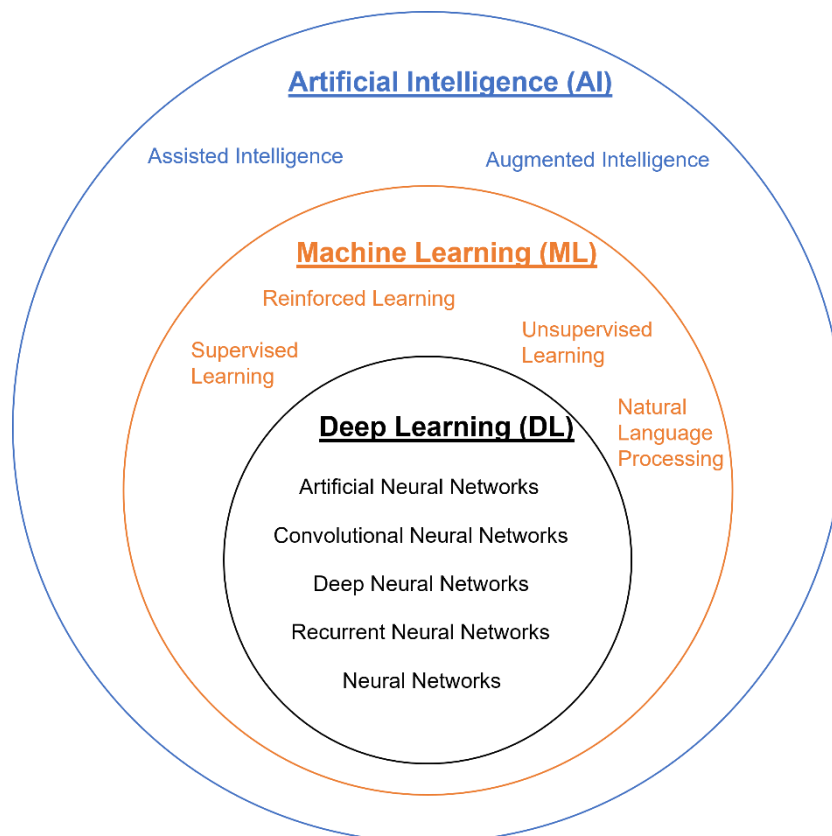


Figure 2: Hierarchy of Machine Learning and Deep Learning subfields within Artificial Intelligence. Machine learning is a subfield within AI, and deep learning is a subfield within machine learning. Various machine learning approaches are displayed, but use of neural networks is exclusively within the purview of deep learning.

1.1 Drug Discovery

The use of AI for drug discovery and repositioning of known drugs is of great interest to and has received a large amount of attention from various pharmaceutical giants. For example, since 2016, Pfizer has partnered with IBM's AI platform Watson for immune-oncology drug discovery (Batool, Ahmad, and Choi 2019). Sanofi has also been collaborating with Exscientia Ltd., an AI-driven drug design company for metabolic disorders. The prospect of using computer-based target identification for prioritization, as well as repositioning of known

⁴ Dimensionality is a function of the length of the human genome. A single sample from one individual may have millions of genetic variants. Data sets have a low number of samples when compared to the high number of possible genetic variations, resulting in high dimensionality of the data.

compounds for added efficiency is especially appealing to pharmaceutical companies for various reasons. Firstly, the cost-efficiency of highlighting and ranking potential compounds as candidates is expected to reduce the failure rate of clinical development. Secondly, repositioning of already approved drugs for new purposes avoids the need to redo trials for safety data, saving time and money while maximizing profit.

For the purposes of drug discovery, AI is used on genomic data to make inferences and identify risk genes for specific mutations that lead to disease. These risk genes are designated as excellent potential therapeutic targets. For example, a drug's effect can be modeled as a gene deleted *in-silico*, and metabolic pathways resulting from the deletion of the gene in question can be predicted as a means to assess a drug's effect (Zampieri et al. 2019). Supervised ML tools have been used also to identify cancer driver genes, which can also serve as potential new drug targets. Unsupervised ML methods can be used to augment and complement pharmacological experts' knowledge and identify genes that could serve as candidates to expand the biological effects of already known drugs (A. P. Athreya et al. 2018).

The ability of AI to enable the analysis of large amounts of data makes these approaches well-suited to drug discovery by merging genomics with pharmacokinetics. AI allows for big data analytics by employing ML to extract a plethora of useful information such as features and patterns from enormous biomedical genomic datasets (metadata), which would otherwise be an arduous and time-consuming process. For example, Gupta et al. report that with AI and ML approaches, virtual screening of compounds from chemical libraries that consist of more than 106 million compounds can be done easily and in a time-effective manner (Gupta et al. 2021). Their review highlights that another example of the use of big data analytics in genomics is the interrogation of genome wide association study (GWAS) repositories, such as GWAS central and NHGRI-EBI GWAS Catalog. Analysis of GWASs allows for the identification of locations associated with disease-associated loci. Genes associated with those loci can then be designated as potential therapeutic targets (Gupta et al. 2021). In addition to those GWAS repositories, another repository is the library of integrated network-based cellular signature (LINCS) L1000, which has data on gene expression changes on various human cell lines when exposed to various compounds.

Developments in ML techniques allow for the extraction of useful features and patterns from the various repositories above for inclusion and consideration in drug development. AI models can also deal with toxicity problems by accounting for off-target interactions with other regions of the genome using the information extracted, with ML approaches shown to more correctly identify inhibitory drug side effects at a higher accuracy than baseline methods (Zampieri et al. 2019).

GAWS databases are rich in information, and the capabilities of ML approaches are limited when dealing with such varied, heterogenous and high-dimensional data. However, DL algorithms have made it possible to integrate data at multiple levels in a nonlinear fashion while maintaining accuracy and precision (Gupta et al. 2021). This has allowed for possibilities such as the combining of large-scale genomic data alongside chemical and biological activity data with hundreds of already approved drugs (Zeng et al. 2019; Shah et al. 2019), as well as the linking of disease phenotype and symptoms to chemical drug structures, genome sequences, and gene expression (M. Bahi and M. Batouche 2018). Ideally, this should allow for the identification of candidate compounds with the highest likelihood of satisfying efficacy and toxicity parameters for a given disease or target.

Another problem that DL techniques have begun to tackle is the data-hungry nature of learning algorithms. These algorithms require large amounts of data for training and optimization (Shah et al. 2019). Stanford University has developed an algorithm that can handle the problem of data scarcity (Altae-Tran et al. 2017). Termed “one-shot learning”, it is meant to be capable of predicting a drug’s chemical properties based on heterogenous and sparse data. Such work is still in development and further model testing is needed to ensure the model is valid, especially across different types of datasets.

In silico drug repositioning is another excellent application of AI to drug discovery, with tools like deepDR systematically inferring new drug-disease relationships and other tools integrating heterogenous data of drug-target interactions, which allows the prediction of novel drug-target relationships (Zeng et al. 2019). In these cases, the targets may be areas of the patient’s genome, and predicting these interactions is key for understanding both the efficacy and toxicity of a drug.

1.2 Pharmacogenomics

Neural Networks can use multimodal learning to jointly learn from various existing datasets independently, without the common features explicitly defined (Kalinin et al 2018). This means they can combine information about a variety of biological parameters (for example, gene expression information such as RNA, miRNA, and methylation) across linked datasets. Since pharmacogenomic data is varied and heterogenous, deep networks are favored as they can deal with sparse, high-dimensional, multimodal data well (Shah et al. 2019; Koumakis 2020). Data limitations which previously caused issues in model training are lessened as deep networks can use transfer learning and apply the learned layers to other datasets.

Current trends of pharmacogenomic research points towards the noncoding portions of the human genome as novel therapeutic regions. These consist of parts of the genome that do not themselves code for proteins, but instead control the expression of the genes that do. In fact, 90% of the pharmacogenomic single nucleotide polymorphisms (SNPs) (single nucleotide base changes in a genome that influence drug efficacy and can lead to adverse effects to drugs) identified in GWAS are located in the noncoding region – particularly in the enhancers, promoters, and introns. This has been termed as the pharmacoepigenome (from epigenome, which explicitly includes regulatory elements) (Kalinin et al 2018). ML approaches have been proposed to untangle this area of high interaction between the different regulatory (or-non coding) regions of the genome – such as enhancer-promoter and drug-target inference – and predict the impact of impact SNPs on eventual phenotype. In particular, DL is implicated as suitable for this application. One such DL tool, DeepSEA, is a concurrent neural network that predicts chromatin effects on sequence variants, and can also highlight the variants that affect regulatory function (Kalinin et al. 2018; MacEachern and Forkert 2021). Many of these are also used for RNA and DNA motif mining, which is important for predicting protein binding sites (He et al. 2021). A DL framework has been developed to predict the effect of long non-coding RNAs on various types of disease (M. Zeng et al. 2021). Various applications of DL tools are currently being developed to address a range of biological questions in this area of the pharmacoepigenome such as histone modification influence of gene expression, methylation state prediction and enhancer-promoter interactions from genomic sequence, and impact of noncoding variants.

Another way that AI/ML is used in pharmacogenomics is in patient stratification. The Food and Drug Administration of the United States in 2013 called for pharmacogenomic patient stratification for clinical trials, but this has only recently been adopted by approaches utilizing AI (Kalinin et al 2018). Proprietary ML algorithms use genomic data obtained from a client's electronic health records and combine it with HapMap and 1000 Genomes Project to perform ethnicity stratification at a population, cohort, and other specified levels. The various interactions uncovered by pharmacogenomics can then be shortlisted for drug discovery and development.

1.3 Cancer

Cancer has been described as a disease of the genome and is most suitable for the application of AI in understanding the underlying interaction of genetic changes that lead to the disease. Cancer arises from unchecked mutations to the genome either due to standard replication error during cell division that escapes repair mechanisms, or due to mutagenesis. The major application of AI towards cancer genomics involves identifying driver genes of cancer and identifying mutations in those genes, as seen in breast and colorectal cancer (Vatansever et al. 2021). This is especially useful not just in a pan-cancer setting, but also when able to use a multitask information learning sharing strategy for specific cancer types and extend that towards all cancers. Cancer management, including survival and treatment plans, is especially contingent on identifying clinically significant biomarkers (Dlamini et al. 2020). ML algorithms can combine genomic information along with gene expression data obtained from microarrays to provide a better assessment of lung cancer (K. Paranjape, M. Schinkel, and P. Nanayakkara 2020).

Another large portion of the research around AI usage and genomics is directed towards personalized oncology, especially when using AI algorithms to determine mutated genes based off pathology images. For example, deep learning networks have been used on The Cancer Gene Atlas (TCGA) database histopathology images to predict the ten most mutated genes in lung adenocarcinomas. Linking image-to-genome mapping was also done in prostate cancer with TCGA obtained data to predict the presence of a SPOP gene mutation in prostate cancers (Jiang et al. 2020). Another study illustrating the potential of personalized oncology demonstrated that it is possible to use ML on genomic data for patients suffering from Acute Myeloid Leukemia to increase survivability by tailoring treatments based on remission and relapse patterns (Radakovich, Cortese, and Nazha 2020). While promising, this is an entirely novel approach and needs much more evaluation to demonstrate safety and efficacy. In addition, patients with specific genetic changes can be identified by their histology alone through an AI network, allowing for increased specificity in precision treatment (Jiang et al. 2020). ML can even predict responses on the individual patient level to certain checkpoint inhibitors or radiation, allowing for personalized and optimal treatment (K. Paranjape, M. Schinkel, and P. Nanayakkara 2020). This has potential for both predicting patient survival and prognostics but also for streamlining or optimising clinical care for patients. Convolutional neural networks have been shown to produce genetic analysis data alongside diagnostic images for predicting disease as well as treatment effectiveness, and this is particularly relevant for cancer and precision medicine (Dlamini et al. 2020).

A particular challenge in cancer treatment is the evaluation of microsatellite stability and instability. Microsatellite instability, along with other genetic traits such as the tumor mutation burden, are strong biomarkers of cancer in regards to immunotherapy (Tran et al. 2021). AI models were shown to be able to evaluate microsatellite stability based on just hematoxylin and eosin-stained gastrointestinal cancers, without the need to conduct laborious microsatellite instability assays. This model was found to be replicable not just on gastrointestinal cancer images but also on endometrial cancers as well (Jiang et al. 2020).

Other potential applications of DL in oncology include molecular subtyping of tumors, survival estimates, prognostication, and cancers with unknown primary (Tran et al. 2021). AI imaging when applied to cancers are also a great benefit for clinical diagnostics. Incorporation of multiple omics such as genomics, transcriptomics, and phenomics in DL models can categorise tumors by their molecular features into important subtypes (Tran et al. 2021). Cancers of unknown primary are a major issue, as metastasized tumors can be treated but will reoccur as the primary site is unknown. AI has been used on the genomic, transcriptomic, and methylated profiles of metastasized tumors to determine their likely origin tissue. This strategy has proven effective at predicting the origin of even rare, treatment-resistant tumors such as that of metastatic adenoid cystic carcinoma, though clinical validation is still necessary (Tran et al. 2021). Genomic data has been used also to more-effectively stratify patients based on risk for Acute Myeloid Leukemia, using a ML approach called gradient boosting trees (Radakovich et al. 2020). A framework for patient stratification using breast cancer biomarkers has also been developed, though has yet to be validated and translated for clinical use (S. Rajpal et al. 2021).

One other application is understanding the 3D tumor microenvironment. Due to the large number of available datasets for cancer research, the availability of genomic, transcriptomic and histopathology images allow for a potentially robust framework for using neural networks to create profiles of the heterogenous tumor microenvironment for subtypes of cancers. Understanding the microenvironment helps aid chemotherapeutic decisions. To further aid this endeavor, Chen et al. have proposed a transformation algorithm to transform 3D images into 2D space, allowing for convolutional neural network to be trained on small amounts of information, which is a major barrier to training ML models (X. Chen et al. 2021).

1.4 Diagnostics

Natural Language Processing-based diagnostics

Today's AI is capable of using phenotypic data (usually in the form of electronic health records) in conjunction with genomic data to provide a molecular diagnosis for use in the clinic (Dias and Torkamani 2019). Many features of AI allow for the examination of genomic data, such as the usage of Natural Language Processing (NLP) tasks. A highly cited review points out that recurrent neural networks from NLP can apply similar text learning methods from input text, such as that found in an electronic health records assessment, and use those methods for input DNA sequences (Dias and Torkamani 2019). Thus, NLP algorithms have the potential to automate genomic analysis with appropriate incorporation of supporting text data. This could be especially useful in GWAS studies with data repositories containing genomic information that are linked to electronic health records. NLP methods can link different electronic health records-based cohorts while reducing the labour that would

usually be required to do so manually (Z. Zeng et al. 2019). There is much potential for NLP to aid clinicians make diagnoses based on a combination of genomic information and phenotypic information.

Wearable Medical Device diagnostics

A study conducted by Gladding et al. showed AI can combine health data collected using wearable devices (such as heart rate, blood pressure) with multiomics and existing clinical data to identify biomarkers and phenotypes for early disease stratification and prognostication (Gladding et al. 2021). Applying an ML tool to an electrocardiogram (ECG) clinical data in combination with other monitoring apparatus allowed Gladding et al to identify patients at risk from heart failure with reduced ejection fraction.

In cardiology, AI is used in combination with genomic analysis but is applied to analyze the measurement data obtained from ECGs or wearables, or in examining clinical data. While this study demonstrates potential clinical utility of ML tools in analyzing data from wearables for more specific care, the study was small and statistically underpowered, requiring further validation before clinical adoption. There is also some skepticism towards these wearable health monitoring devices (Shah et al. 2019).

1.5 Neurology/ Ageing

Two modestly cited papers in this field have pointed out new advances in aging and aging-related disorders such as Alzheimer's disease and Parkinson's. Vatansever et al. point out that deep neural networks have been used to tackle the dimensionality problem that accompanies RNA transcript data and predict genes for Parkinson's disease (Vatansever et al. 2021). While Parkinson's has single genes linked to it, the known monogenic causes account for less than 5% of Parkinson's cases, suggesting that identification of new genes opens up new therapeutic possibilities. Other DL models have also been used to successfully identify genes linked to autism, such as a model that is trained to predict RNA splicing from the genomic DNA sequence alone (Vatansever et al. 2021). Using unlabeled learning algorithms, over 3,000 candidate age-related genes were predicted and prioritized according to their influence in aging based on gene information from 11 human biology databases (Vatansever et al. 2021).

ML has also been used in direct-to-consumer genetic testing such as 23andMe as well as Ancestry.com (Graham et al. 2020). In direct-to-consumer genetic testing, ML is employed to annotate many of the regulatory elements of the genome (splice sites, promoters, enhancers, etc), as well as identify disease biomarkers and differentiate between disease phenotypes. However, for a polygenic disease such as Alzheimer's, one weakness of using ML with GWAS datasets is the sheer mass of data needed to train such a model. This limits the depth of phenotypic data and would reduce the accuracy of these models. In a separate study, L. Li et al. used a dataset from the Alzheimer's Disease Neuroimaging Initiative (ADNI) along with participant genomes, demographic information and diagnosis information to conduct a comparison between a DL genomics (DLG) method and traditional GWAS analysis to identify significant SNP loci as contributors to Alzheimer's disease (L. Li et al. 2021). While the GWAS method identified only two significant SNPs, the DLG method identified over a thousand. The authors therefore concluded that DLG outperformed traditional GWAS analysis. While this shows the potential of AI models for finding hidden interactions of SNPs, experimental

validation of those uncovered SNPs is required to make any conclusion about the increased sensitivity of DLG methods over standard GWAS analysis.

1.6 Image to genome mapping

In areas such as macular degeneration and tumor image analysis, AI is able to more accurately predict prognostication better than human practitioners. Efforts and applications to exploit this ability of AI have led to image-to-genome mapping and other techniques that make use of AI's image-based capabilities (Shah et al. 2019). For example, identification of recurring motifs in a DNA sequence as part of the regulatory elements of a genome are similar to image pixel pattern detection by CNNs (Dias and Torkamani 2019). Imaging gives a clear identification of phenotypic differences, and ML models can be trained to link these phenotypic features to underlying genetic causes. As an example, Mobadersany et al. have developed a single prediction framework termed the genomic survival convolutional neural network that unifies both histological and genomic data that exceeds the accuracy of human experts in predicting survival (Mobadersany et al. 2018). Their study has been cited over 500 times, and is a likely indicator of the future of image-to-genome mapping. These computer vision abilities to obtain molecular information from phenotypic images are equal to or have exceeded skilled human experts and the current existing WHO paradigm of grading and classifying different gliomas (Dias and Torkamani 2019; Mobadersany et al. 2018). Facial image software and algorithms have been also able to identify certain genetic disorders and link them to molecular diagnoses (Dias and Torkamani 2019).

As anticipated in section 1.3, image-to-genome mapping is also of considerable use in cancer research. Computer vision techniques in combination with DL have been able to discern the somatic mutations that cause particular histopathologies of cancer cell images (Dias and Torkamani 2019). Histopathological images are exhibitions of underlying genomic features and AI may allow for obtaining clinically useful information without the need of having to sequence the tumor itself (Tran et al. 2021). This would reduce the need for potentially invasive biopsies.

The Human Phenotype Ontology (HPO) program also aims to extend digitally encoded phenotypes with a variety of disorders, such as immunological diseases (Köhler et al. 2019). These can potentially be used to enable the linking of genomic variants to their phenotype, expediting the identification of patients. Currently, the All of Us program launched by the National Institutes of Health in the US aims to combine the mined electronic health records data and genomic data of over 1 million people alongside NLP and HPO to create a large dataset for AI training and use for diagnosis, patient treatment and care, as well as general research (Rider, Srinivasan, and Khoury 2020).

1.7 Radiology

Trends in radiology are broadly similar to those outlined in the above section on the applications of computer vision. Given the close ties to cancer, usage of AI imaging software allows for structural and functional characterization of lesions and tumors. (N. Sedaghat et al. 2018) A small review highlights the potential to use AI image-based capabilities to construct a radiogenomic map of lesions, which can provide an insight into the organization of the intralesional heterogeneity present. Linking together digital medical images to underlying

molecular biomarkers is meant to aid in clinical treatment and diagnosis (Peng et al. 2021). Alongside the possibility of using AI and ML techniques to identify potential therapeutics specific for the type of lesion or tumour (Wong and Chaudhry 2020), this approach could complement (but not replace) existing clinician care.

1.8 MicroRNAs

In recent years, there has been an increase in the use of AI in the gene regulatory space. This is due to a combination of increased demand for this deployment of AI on the pharmaceutical industry, as well as the maturation of DL techniques. This has benefited our understanding of Micro RNAs (miRNAs).⁵ Previous, statistically-driven correlation approaches were not suitable for miRNA prediction due to the high abundance of non-linear associations (N. Sedaghat et al. 2018). Thus, the use of AI and ML applications has the potential to resolve these challenges. Prediction of these miRNAs are of great value in drug discovery and pharmacogenomics, and effort has been made to predict both mature miRNA sequences as well as the sequences of precursors (pre-miRNA) in order to design compounds targeting them. Additionally, miRNAs may contribute to cancer treatments, as aberrant gene expression of cellular replication genes can lead to cancerous cells. Thus, miRNA research as an area of gene expression regulation seems to be a prime area for employing AI and ML methods.

Traditional ML approaches in miRNA research that did not use Deep Learning include the prediction of Dicer cleavage site selection in pre-miRNAs using a SVM, and RF-based miRNA predictors, which distinguish real and pseudo-pre-miRNA from genome sequences (L. Chen et al. 2019). Some have used ML-based methods to identify relationships between miRNA and disease (Vatansever et al. 2021). Using neural networks, miRNA target discovery and disease associations have implicated new miRNAs in lung and breast neoplasms. These miRNA-disease associations are determined by a combination of miRNA sequence, disease semantics, and miRNA function through a neural network (Vatansever et al. 2021). Chen et al. have reviewed the many applications of AI and ML on miRNA prediction (for a full list of the various software tools and their intended usage, see (L. Chen et al. 2019) (L. Chen et al. 2019). This highly cited review shows that most applications of ML in miRNA involve either miRNA identification, miRNA analysis or the prediction of miRNA targets. In particular, Deep Learning has been of use in target prediction, and has been instrumental in *de novo* prediction of miRNA, predicting mature miRNA in novel precursors, and identifying miRNA in small RNA datasets (L. Chen et al. 2019). Deep Learning has even been used to predict novel miRNA in species that have not had an established sequenced genome. It has been proposed to use a Deep neural network for distinguishing pre-miRNA from pseudo hair-pins (S. Mohanty, S. Mahapatra, and T. Swarnkar 2021).

1.9 Sequencing

Sequencing technologies have come a long way since Sanger sequencing in the 1970s. Now, next-generation sequencing (NGS) is employed to conduct most sequencing work, with costs

⁵ miRNAs are short, non-coding RNAs that influence gene expression. Unlike normal messenger RNA (mRNA) transcripts, which are created from reading the genome and translated into functional proteins, miRNAs bind to mature mRNA, preventing them from being translated into protein. miRNA follows a similar life cycle to normal mRNA, with their sequences existing within the genome.

and efficiency far improved from previous sequencing methods. A review by Dlamini points out that the large-scale, complex genomic data generated by NGS makes it the ideal candidate for analysis by AI-enabled toolsets to identify patterns and correlations (Dlamini et al. 2020). NGS utilization of ML can help tackle the challenge of detecting structural variants (insertions and deletions) that occasionally occur with long read DNA sequencing.

The ease and speed of running NGS presents another big-data related issue: the complexity of errors generated from the sheer mass of information produced by multiple reads. As Celesti et al. puts it, classical statistical and ML approaches have been unable to address SNP and INDEL calling, but DL approaches are suitable to fit this problem as NGS big data provides large enough data sets with high-quality input data and training labels (F. Celesti et al. 2018). In Celesti's review of DL applications in NGS, DL has been employed for gene regulation studies involving epigenetic modification and regulatory element interactions, genome analysis and SNP prediction, as well as cancer biomarker detection.

1.10 Single Cell Genomics

Precision medicine makes use of single cell genomics to identify and classify cell populations to predict disease progression. Previously, single-DNA has been used to shed light on the evolution of tumors and somatic mosaicism in healthy tissue (W. Yang et al. 2017). The focus on single cell genomics is to delve deeper on both temporal and spatial elements of cellular change, which provides a layer of analysis that traditional methods of protein expression measurements and morphology cannot obtain. As genomic information accumulates, a significant challenge facing researchers is handling the issue of technical noise, such as transcripts that are expressed but undetected, and data scale. Previously, single cell analysis was accomplished with PCA and t-SNE to reduce the dimensionality, though these statistical models are influenced by the experimental noise (false positives, false negatives, data sparsity) from the data. W. Yang et al. used convolutional neural networks to resolve the noise issue, and deep-autoencoder to address the high dimensionality (W. Yang et al. 2017). Additionally, convolutional neural networks, recurrent neural networks, and Long Short-Term Memory (LSTM) methods were employed for data mining the text and image material of existing literature for extraction. Their future work involves combining the static screenshot of a cell's gene expression with temporal features to group the cells more accurately for brain disorder studies or drug target discovery.

1.11 Data Sharing

One very unique application of AI and ML methods to genomic data is using DL to create a minimization algorithm that allows the convenient transfer of big genomic databases. A single human genome consists of 3 billion nucleotide base pairs, and just one full, whole genome sequence generates 250 gigabytes of data (M. Aledhari et al. 2021). Sharing large amounts of data online across labs or institutions becomes a security and speed issue. Aledhari et al. have used a novel data minimization approach relying on a CNN algorithm during the data transfer process to package whole genomic data sets into smaller bits. Compared to standard encoding, this process was 98 times faster using both real and simulated datasets. While this is an innovative approach to the problem of storage and the expediency of sharing data,

security- and privacy-related issues have yet to be addressed (see section 2.1 below for a discussion of these issues).

1.12 Non-healthcare related applications

The search for non-healthcare related applications returned a total of 2404 results from Pubmed and 1194 results from IEEE Xplore. However, upon examination of the titles of the first 50 results from each database, it was found that none of the returned results dealt with any of the non-healthcare related terms. This is likely due to the chosen databases being exclusively scientifically or technically focused.

1.13 Emerging approaches and predictions

Many researchers have proposed novel uses and applications for AI in the genomics research setting. The major driver of much of these current and emerging uses is the maturation of DL approaches, especially that of convolutional neural networks and recurrent neural networks, which prove to be very suitable for genomic analysis, due to their capability to make inferences from exceedingly complex data with high dimensionality (Dias and Torkamani 2019; Koumakis 2020). Another area where DL has strong potential is the field of precision oncology, where pharmacogenomics are employed to explore drug interactions and drug mechanisms with a patient's specific genome (Tran et al. 2021). Overall, DL seems to have become the predominant approach, applied to challenges that previous ML methods were unable to resolve or improving those applications.

DL is also highly prominent in the miRNA prediction field. In fact, as seen in the literature, there is increasing focus and research in the transcriptome and epigenome, with work focusing on splicing prediction, gene expression regulation from miRNA, and the role noncoding elements of the genome play, especially in regards to pharmacogenomics and DNA binding proteins (Vatansever et al. 2021; F. Celesti et al. 2018). Additionally, DL has also been used for the prediction of DNA sequences that may serve as DNA binding sites for transcription factors (F. Ahsan et al. 2020; Q. Zhang et al. 2021). This increased attention to the regulation of the underlying genome generates more dimensional data, necessitating the usage of ML. Characterization of the 4D nucleosome, where time is added as another dimension, is a major topic in pharmacogenomic studies and requires ML usage to tease out the complex associations with the added dimension of time, such as gene expression time series (R. Yilmaz and A. E. Pusane 2020; Kalinin et al. 2018). This is an area of promising development in the future for both pharmaceutical development as well as scientific understanding.

Patient subtyping at the population and cohort levels based on individual genome features is another area that is expected to employ AI for patient categorization and stratification (Radakovich, Cortese, and Nazha 2020; Vatansever et al. 2021; Kalinin et al. 2018). This would involve assigning certain patients more or less of a compound, based on the expected efficacy of that compound on the individual, or for tailoring treatment regimens. Radakovich et al. predict that the acceptance of ML in the clinic depends largely on its ability to show transparent rationale for its output, which, given the black box problem, may be a barrier to adoption, as we discuss in section 2.3 below (Radakovich, Cortese, and Nazha 2020). They are,

however, hopeful of the utility of clinical AI in the future, with the belief that sufficiently optimized ML models would be able to estimate prognostic outcomes, along with the uncertainties attached to those outcomes, as well as clinically-accessible explanations for their predictions. The authors conclude, however, that this is a very aspirational outlook. Koumakis asserts that due to validation and standardization requirements, translating genomic research knowledge into clinically useful tools will be a slow process (Koumakis 2020).

Other omic areas are being examined and offer rich, unexplored data. Areas such as the gut microbiome, when combined with genomics, transcriptomics, and proteomics are hoped to create a more integrated and holistic profile of an individual patient (Tran et al. 2021). With the ability of neural networks to map associations and connections, it is believed that these areas will reveal new therapeutic potentials, especially for cancer. Creation of these more holistic biological models, with integrated information across the omics fields in addition to electronic health records information, is believed to significantly cut down on costs when compared to empirical trials, especially for fields such as metabolic engineering (Zampieri et al. 2019). One challenge is the reliance of AI training on large datasets (Shah et al. 2019), and much work has gone into developing and creating these datasets, such as the UK Biobank and All of Us research program. While genomic data is being generated at rapid pace, scalable phenotypic data is lacking due to the slow and expensive nature of collecting phenotypic data along with the inherent variable nature of phenotypic interpretation (Dias and Torkamani 2019). The increasing accessibility and availability of public omics data in various databases has accelerated deep learning development for cancer pharmacogenomics (Tran et al. 2021).

From a technical perspective, cloud computing using DL algorithms tends to be lacking. F. Celesti et al note that while DL is adopted in many NGS standalone applications, future cloud computing services would benefit from DL algorithms for scalability of computational resources, as well as data sharing of DNA fragments (F. Celesti et al. 2018). They predict the future will see the development of advanced DL network software tools for a variety of genomic applications (F. Celesti et al. 2018).

These emerging capabilities of AI tools in genomics research are exciting and inspiring. However, the predictive capabilities of ML/DL approaches require experimental testing for their findings to be considered valid. The hype surrounding AI approaches often far exceeds the actual science. For example, it is not expected that AI will overtake capabilities of a trained clinician (Rider, Srinivasan, and Khoury 2020; MacEachern and Forkert 2021), although the tools that AI offers will mean that a clinician that ignores these developing tools will likely be replaced by a competent physician that does make use of them. A judicious degree of responsibility is needed when considering AI's increasing influence in the clinic, especially given a previous case where IBM Watson gave incorrect recommendations for cancer patients (Rider, Srinivasan, and Khoury 2020). Therefore, in addition to investigating the capabilities and applications enabled or enhanced by AI in genomics, it is important to understand their ethical, societal, legal, and political ramifications, to which we now turn.

Part 2: Ethical, societal and legal debates concerning genomic science

In the review of the academic literature focused on the ethical, societal, legal, and political ramifications of the capabilities AI is engendering in genomics, we identified 102 pieces of literature from Web of Science, of which 44 were relevant after screening title, abstract, and keywords, and 186 pieces from Scopus, of which 63 resulted relevant. After screening the full text and deleting duplicates, 59 papers from the two databases resulted relevant, of which 30% were included for analysis. A number of additional pieces of literature were added through Google searches and manual searches in the bibliography of the identified papers, and from websites of relevant organisations.

Most literature on the ethical, societal, legal, and political impacts discussed healthcare applications of the emerging genomic capabilities made possible by AI. In this setting, a widely debated application is the use of AI methods to integrate large volumes of heterogeneous types of patients' data, including clinical, genomic, metabolomics, proteomics, and lifestyle data, and publicly available annotated data such as genes, variants, drugs, and biomarkers (Ahmed 2020). As the first part of this review has illustrated, this large amount of data requires advanced computing capabilities to be analysed (Alaidarous 2020). It is hoped that these approaches will allow better understanding of fundamental biological processes to lead to enhanced understanding of human diseases; improvements in diagnostic predictions and patient classification, for example by identifying subtypes of patients suffering from a certain condition; and identification of drug targets and interactions, or prediction of patients' different drug responses (Ching et al. 2018). In turn, this could inform decisions around allocation of healthcare resources (Kimberly 2019). Altogether these advances aim at providing precision medicine, particularly in areas like oncology, and suggest that we may be moving towards a learning healthcare system, which "applies the best known evidence, encourages continuous learning, and allows for knowledge generation as a natural by-product of patient care delivery" (Flahault et al. 2017).

Among the identified literature, three papers discussed applications in the context of law enforcement (Arnold 2020; Ramjee and Ringrose 2020; Ivanov, Matsepuro, and Trubnikova 2022). In this context, AI could be used for forensic genetic genealogy to identify alleged law offenders, by comparing suspects' DNA to profiles stored in genealogy databases and piecing together family trees. Another application is forensic DNA phenotyping, in which algorithms are trained to identify patterns between DNA samples and face photographs, and can then be used to predict individual characteristics of an unknown person (e.g., potential suspect) from DNA found at the crime scene (Arnold 2020). These techniques can also be used in efforts to identify a suspected homicide victim from unidentified remains.

A couple of papers focused on healthcare education (Khamisy-Farah et al. 2021; Green et al. 2020). In particular, authors argue that medical education should equip the next generation of healthcare professionals with the skills required to work in healthcare systems that are changing due to recent scientific (genomic and post-genomic) and computational (data science, Artificial Intelligence) advances, and their interactions in disciplines like bioinformatics.

An emerging strand of literature focused on the ethical implications arising from the convergence between, on the one hand, research on the genetic underpinnings of human learning and, on the other hand, advances in AI. The convergence between these two

research areas is aimed at predicting children's educational, socio-economic, and life outcomes (Adam 2019; Williamson 2020; Comfort 2018).

Finally, only a minority of the identified literature engaged with issues around the costs of implementing the new genomic capabilities enabled by AI (Ching et al. 2018; Nuffield Council on Bioethics 2018); and with issues at the macro level, particularly how advances in the analysis of large amount of different types of data may impact the distribution of power across world regions (Hummel and Braun 2020).

Most papers made considerations that are not tied to any specific country or region. However, some literature focused on the following countries: North America (US and Canada), Europe (France, Germany, Italy and other European countries), Australia, China, Malaysia, Saudi Arabia, and Qatar.

Below is a summary of the main ethical, societal, legal, and political impacts discussed in the literature.

2.1 Data, code and software engineering: Privacy and discrimination

Data sharing in healthcare

In healthcare settings, data sharing among laboratories is essential for the computational analysis of genomic data aimed at, for example, classifying patients in groups with similar genotypes or for drug responses (Ching et al. 2018), and allows a range of advantages, provided that high quality data are shared. In particular, it allows the creation of large datasets for algorithm training and accuracy (Cudai et al. 2021). Data sharing also allows to increase cohort sizes, and so avoid statistical difficulties that derive from having a much larger number of variables (genomic measurements or features) than the number of patients on whom these measurements are collected (samples) (Azencott 2018). Some also argue that it reduces the need for animal experimentation, and may save costs that laboratories would have to sustain to produce new data (Cudai et al. 2021).

To make data sharing possible, collaborations need to be established across institutions and stakeholders, particularly government, academia, industry (spanning from those working on life sciences to those working on artificial intelligence), and civil society (A Jamal 2021; Green et al. 2020; Nellaker et al. 2019). This is currently difficult due to inconsistent or evolving standards for data sharing in healthcare. For example, while healthcare records would be a great source of data for making progress in precision medicine, as they contain information about patients' characteristics, biometrics, health conditions, health screenings, insurance status, and medications (Breen et al. 2019), ambiguous regulatory language hinders the sharing of such records across institutions (Ching et al. 2018). Moreover, medical data usually cannot leave the data owner's closed environment (e.g., hospitals) without ethics approval and patients' informed consent (Eggel, Schaer, and Muller 2018). In academic publishing too practice varies, with some journals having no policy about data sharing, while others require data sharing with no restrictions or on request (Car et al. 2019). One paper in this review also argues that data sharing practices are hindered by concerns around intellectual property, as they raise questions about who owns the shared data and to whom findings that result from the use of such data belong (Nellaker et al. 2019).

Furthermore, some have argued that these collaborations run the risk of disadvantaging institutions with fewer resources, whose data could be used to “further the careers of their peers in richer countries” (Car et al. 2019). However, data sharing would make it possible for those laboratories that, due to limited budget, cannot produce their own data, to access data collected elsewhere, and so could, on the contrary, contribute to advance their research (Caudai et al. 2021). Yet, the issue remains that sharing large datasets requires receiving institutions to have high storage capacity and strong computing infrastructures (see section 1.11 above), which could contribute to the digital divide between research institutions or tech companies in higher and lower resource countries (Eggel, Schaer, and Muller 2018). Although not much discussed in the identified literature, this consideration raises the question of whether recent scientific and technical advances at the intersection between genomics and AI may increase the concentration of economic and political power in the hands of a few organizations, usually residing in wealthier countries, which will therefore reap most of the benefits from research and development in this area (see also section 2.7 on *Concerns at the macro level*) (Stahl 2021).

Sharing of healthcare data also raises a range of ethical and legal considerations. One important concern among members of the public is the impact that the availability of their healthcare data could have on individual privacy and confidentiality (Alrefaei et al. 2022). For the purposes of this section, privacy concerns relate primarily to cases in which data are shared (e.g., with employers, insurers, or tech companies) for uses that data subjects are unaware of, and which may not align with their interests or preferences (Hern 2017). However, data storage and sharing may also be subject to unintentional breaches, hence causing concerns about the possible subsequent misuse of the leaked data. In both cases, privacy violations may lead to public mistrust towards institutions owing or sharing patient data, and public backlash if technological and scientific advances are not perceived to be pursued in the interest of society as a whole (Nuffield Council on Bioethics 2018).

Genetic data, in particular, are of sensitive nature, as they contain information about an individual’s susceptibility to certain health conditions. Thus, if genomic data are unintentionally leaked or intentionally shared beyond healthcare institutions, individuals may be subject to genetic discrimination, a term used to refer to the risk that people may be treated differently because they have (or are perceived to have) certain genetic variants (A Jamal 2021; Azencott 2018). In addition, genomic data are not only relevant to the individual, but may reveal information about a patient’s family members and are even relevant to their broader ethnic group (Caudai et al. 2021; Azencott 2018). This causes significant challenges to traditional methods of consent to data storage and management, and suggests that such methods may be inadequate in genomics, as the data collected from individuals who give their consent can be used to draw inferences about individuals who never consented (Hummel and Braun 2020; Nellaker et al. 2019).

Concerns around discrimination become even more pressing with the rapid advances in AI and genomics. Some have argued that, because AI tools are allowing new understandings of the human genome, it has become increasingly important to keep genetic and genomic data confidential, as we currently do not know what information of an individual’s genomic makeup will tell us about them in the future (Caudai et al. 2021). Moreover, in computations made possible by ML and DL, genomic information is combined with sensitive clinical information (e.g., lifestyle data like drug use), as well as data (e.g., social media use) that may indirectly provide information about an individual’s health status, whilst not being subject to

the same legal protection that healthcare data are (Nuffield Council on Bioethics 2018). The combination of these large datasets can reveal information about an individual that genomic information alone cannot, which provides further evidence for the need to manage genomic data in a way that protects patients' privacy.

In addition to data sharing, code sharing and open source software are considered essential to further accelerate progress in genomic research (Ching et al. 2018). However, these options raise issues around liability (Dias and Torkamani 2019). Standard licence agreements of proprietary software include provisions about the responsibilities of both developers and end users, and usually the former provide indemnification and warranty, as they maintain copyright of the software and its code. Caudai and colleagues explain that most bioinformatics software is instead released under the GPL (General Public Licence) or the MIT licence. These licences have a disclaimer that limits the liability of software developers, and so shifts responsibility (e.g., in case of software failure) onto the end users, who are however not involved in software development (Caudai et al. 2021).

Beyond healthcare: AI & genomics applications for law enforcement

Beyond healthcare, concerns around privacy violation due to data sharing have also been raised in relation to the use of genetic data deposited in genealogy databases for law enforcement purposes. Cases in which data have been shared by genetic genealogy companies in violation of their agreement with users (Arnold 2020) have prompted the US Department of Justice to produce the *Interim Policy on Forensic Genetic Genealogical DNA Analysis and Searching* (The US Department of Justice 2019). This document permits law enforcement agencies to use forensic genetic genealogy only for serious violent crimes (e.g., rape and murder), and only after other leads have been exhausted, including a search of the FBI's national DNA database for criminal forensic investigation (Combined DNA Index System or CODIS). The policy also prohibits arrest of suspects based solely on a genetic association detected with data stored in a genealogy database (The US Department of Justice 2019). In some cases, agreements between companies and end users do include details about how and when genealogy data can be shared with other institutions. As has been argued, in these cases it is essential that those details are made clear and transparent to users, together with information about the possible implications in the context of criminal justice, rather than being hidden in Terms of Service documents that users tend not to read (Berkman, Miller, and Grady 2018).

Applications at the intersection between genomics, AI, and law enforcement also run the risk of reinforcing prejudices and enabling or justifying discriminatory practices against ethnic minority groups. In particular, one concern that emerged from the literature is that these applications, which link genealogy to crime, do not consider the historical and societal reasons why individuals from certain groups (e.g., people of colour) are more likely to end up in criminal DNA databases (see also section 2.2. on bias) (Arnold 2020). The reviewed literature suggests that these considerations are particularly important in contexts like the US, where in 2020 regulations were introduced to allow the Department of Justice to collect DNA from immigrant detainees and upload the resulting sequences to CODIS; and China, where genetic profiling has been used to target members of the Uyghur group (Arnold 2020).

Proposed mitigations from the literature

Limitations in data sharing practices have “the practical effect of reducing cohort size, limiting statistical significance, preventing the detection of weak effects [143], and restricting the

number of parameters that can be trained in a model” (Ching et al. 2018). In healthcare, this means slowing down scientific discoveries and reducing the potential health benefits that may derive from data sharing (Eggel, Schaer, and Muller 2018). Thus, scholars have highlighted the need to identify solutions to reduce the risk of genetic discrimination when genomic data are shared, and to protect the privacy and confidentiality of data subjects when their genetic data are combined with other data for computations enabled by Machine Learning or Deep Learning methods (A Jamal 2021; Ahmed 2020).

Legal and regulatory initiatives to address concerns around genetic discrimination have been available for some time. For example, in the EU protections against this form of discrimination comes from Article 21 of the EU Charter of Fundamental Rights (European Union 2010) and from the more recent General Data Protection Regulation (GDPR) (European Union 2018). In the USA, the Genetic Information Nondiscrimination Act (GINA) prohibits the use of genetic test results by employers and health insurers (United States Government 2008).

More relevant to the present investigation at the intersection between AI and genomics, some papers discuss technical solutions to protect patients’ privacy beyond anonymization techniques. Anonymization, which removes identifiers (e.g. names) from the datasets, can indeed be a target of linkage attacks, which make re-identification possible through the use of auxiliary information, especially with increased dataset dimensionality (Azencott 2018; Cudai et al. 2021). An alternative that is often discussed in the literature is differential privacy, in which noise is added to computations so that it becomes impossible to tell if any individual’s data were included in the original datasets by looking at the results; in this way, databases can be queried without the risk that personal data are inferred from the output of the computation (Kearns and Roth 2020). However, noise injection can lead to inaccurate results, unless large datasets are used (Azencott 2018). As another example, k-anonymity, which is used to “selectively remove or generalize data until no attribute combination is shared by less than k records”, has been used to overcome some (but not all) loopholes found in traditional anonymization techniques (Cudai et al. 2021).

Recent work has also led to the development of software (e.g., DataShield) that allow computation to be carried out at the site where the data usually reside, and so where they remain under the control of the primary custodian (Ching et al. 2018). By “bringing the algorithm to the data”, data remain hidden from external researchers who want to run their algorithms on a particular dataset, as the researchers’ algorithms are executed directly on the data owner’s internal infrastructure (Eggel, Schaer, and Muller 2018). Cloud computing services that are secure and compliant with regulations have also been developed, in order to reduce concerns around storing sensitive (e.g., healthcare) data in the cloud (Ching et al. 2018).

Finally, in order to address societal concerns around the use of personal data to realise the new capabilities that AI is enabling in genomic research, and so to limit potential mistrust and backlash against AI, initiatives could be established to engage with the views, preferences, and interests of members of the public (Nuffield Council on Bioethics 2018). In one study investigating Saudi society’s awareness and attitudes towards the Saudi Human Genome Program (SHGP) and the role of AI in genetic data analysis within SHGP, participants reported to be in favour of public initiatives aimed at educating Saudi society around issues related to sharing and privacy of genetic data (Alrefaei et al. 2022).

2.2 Data and user bias: Misclassification and fairness

Types of bias

In computer science, the expression “garbage in garbage out” is used to refer to the fact that errors in the input data of an ML model will be reflected in the model outputs. Machine bias can take various forms and it is important to consider in computational advances in genomics, as it gives rise to problems of relevance, completeness, balance, and accuracy that have hindered the deployment of AI algorithms in clinical diagnostics (Alaidarous 2020).

Ching and colleagues (2018) explain that “biomedical data do not consist of precise measurements but of estimates with noise” (Ching et al. 2018). To increase reliability in automated decisions made in the context of healthcare, it is therefore important to obtain measures of how noise is propagated from input to output. Two main sources of uncertainty can be distinguished: epistemic uncertainty, which is due to insufficient training data or differences in the distributions of the training and testing datasets, and so can be mitigated by collecting more data; and aleatoric uncertainty which, as we discuss below, is due to noisy or missing data (Ching et al. 2018).

In the healthcare contexts considered in this review, noisy data are usually a consequence of biological variability, and hence of the physiological differences among individuals or within the same individual over time; as well as of technical artifacts. These are errors introduced by the equipment and techniques used (e.g. sequencing methods), which lead to results that do not represent the “true” biological material (Cudai et al. 2021). These sources of noise need to be cleaned “so that biologically relevant differences can be reconstructed and analyzed accurately” (Li, Brouwer, and Luo 2022).

Another source of bias is incomplete or missing data. Issues around completeness may be the result of unavailable values or measurement failure, inconsistencies in the way in which electronic health records are compiled (Ching et al. 2018; Cudai et al. 2021), or data fragmentation, due to the fact that patients often receive care from different institutions (Breen et al. 2019). Moreover, advances at the intersection between AI/ML and genomics require the integration of databases with disparate data types. Not only can some data point be missing in some databases; the integration of databases with different schemas can also lead to error and data loss (Cudai et al. 2021).

Datasets can also be unbalanced due to representation biases, which are a result of sampling strategies that lead to data classes being represented unequally. Thus, “Non-representative samples lack the diversity of the population, with missing subgroups and other anomalies” (Mehrabi et al. 2021). Alaidarous gives the example of potential bias introduced when computation analysis for the diagnosis of a genetic condition is based only on the functional expression of a pathogenic gene (Alaidarous 2020). However, a more widely reported issue in this context is the lack of diversity in the populations from which genomic data are collected, which tend to have an over representation of male and sick individuals from European ancestry (Azencott 2018; Ching et al. 2018). Although this allows, for example, the identification of ancestry-specific biomarkers, it also makes it difficult to know whether findings will be relevant to individuals from other groups (Cudai et al. 2021). To give an example, Breen et al. explain that research on the BRCA gene (linked to the development of hereditary breast–ovarian cancer) has showed that there are substantial differences in variants within this gene between individuals of Chinese and non-Chinese ethnicities (Breen et al. 2019). This concern is exacerbated by the fact that, to create large datasets for

computations, biased genomic datasets are combined with other types of datasets that suffer from the same representation issues. For example, sensor data from devices like smart watches are skewed towards those who own such devices, who usually are more technologically literate and come from higher socio-economic backgrounds (Breen et al. 2019).

Finally, classification problems require that the information encoded in the training data is accurate. As Ching and colleagues have highlighted, bias can be introduced in ML models when training data are not accurately labelled (Ching et al. 2018).

However, bias can occur in ML models even without errors being introduced in the way in which data are collected, measured, sampled, or labelled. This is due to so-called historic bias, which Suresh and colleagues define as “a normative concern with the world as it is; it is a fundamental, structural issue with the first step of the data generation process and can exist even given perfect sampling and feature selection” (Suresh and Gutttag 2021). A common example of historic bias is when automated decision-making used to select candidates for tech jobs make recommendations that are not gender neutral, because they are trained on datasets that, due to male dominance and historical prejudices against women in the tech sector, include a majority of male applicants (Dastin 2018). Although less discussed in healthcare contexts, historic bias is an issue that can affect healthcare datasets too. For example, certain medical specialities like psychiatry rely heavily on clinical notes. Because these notes are written by humans, they reflect the prejudices that, as evidence shows, affect the way in which healthcare professionals interact with, and make treatment recommendations for, ethnic minority groups and female patients (I. Y. Chen, Szolovits, and Ghassemi 2019). When data from healthcare records are combined with genomic and other data, these historic biases become embedded in the models used for personalised medicine (Dias and Torkamani 2019).

Consequences of biased datasets on healthcare applications

Machine bias has important consequences in functional genomics and healthcare applications and can limit the practical utility of AI systems applied to genomics in clinical diagnostic (Dias and Torkamani 2019). In particular, biases can lead to poor model accuracy. Here accuracy should not be understood as accuracy of input data (see section on *Types of bias*), but rather as accuracy of outputs, i.e. the correctness of the model decisions or predictions about a specific person and across a wider population relative to the task at hand (Rodolfa, Lamba, and Ghani 2021).

Bias in a classification algorithms can indeed raise problems of misclassification, particularly of either false positives (non-member elements are assigned to a class) or false negatives (some elements are mistakenly not assigned to the class they belong) (Cudai et al. 2021). If the “class” is, for example, cancer diagnosis, false positives may cause psychological distress and lead to increased costs due to unnecessary follow-up tests; and false negatives may cause delayed diagnosis and, in turn, risk jeopardizing patients’ health (Alaidarous 2020).

Furthermore, if algorithms are trained on unbalanced datasets, they risk underperforming when applied to underrepresented groups (Dias and Torkamani 2019). For example, an ML model for genetic risk predictions that is trained only on data collected from patients from European ancestry may identify associations that are only relevant to that ethnic group, and so it may underperform when used in other ethnic groups. Underperformance with certain population groups raises concerns around fairness, as it may lead to inequalities in the way in

which people are treated within the healthcare system, and to potential poorer health outcomes for certain groups, particularly those with protected characteristics (Ching et al. 2018). This is why, in its 2020 strategic vision, the National Human Genome Research Institute identified attention to diversity, and particularly the systematic inclusion of ancestrally diverse and underrepresented individuals in genomic studies, as one of the guiding principles and values for human genomics (Green et al. 2020).

Proposed mitigations from the literature

It is because of these and similar concerns that efforts have been made to develop debiasing techniques, and so reduce machine bias (Alaidarous 2020). As already anticipated in section 1.10, various statistical methods can be used for data denoising (Li, Brouwer, and Luo 2022), i.e. to either correct artefacts or remove incorrect and irrelevant measures in datasets (Cudai et al. 2021). Further, solutions to incomplete datasets consists in either removal of variables or samples with missing values, which however may lead to more bias; or data imputation, i.e. forecasting and imputing the most appropriate value for each missing measurement, for which several technical strategies exist (Cudai et al. 2021). In relation to incomplete and fragmented data from electronic healthcare records, Breen et al. also describe studies that have standardized data collection and measures across different healthcare systems, and initiatives that have developed methods to combine and compare electronic health records across health service providers (Breen et al. 2019). In addition, to target representation bias in datasets, and so minimize inequalities, scholars have advocated for initiatives for large-scale data collection, as well as for regulation introducing fairness standards (Dias and Torkamani 2019). Finally, misclassification issues could be addressed by either using AI methods trained to recognize classification errors, or by adopting guidelines that require re-examination of complex cases by experts (Cudai et al. 2021). More generally, deep networks that use transfer learning to learn sets of rules based on one dataset, and then apply those rules to another dataset (see section 1.2 above), could help in overcoming issues related to models trained on a single dataset, which will necessarily replicate the biases intrinsic in that dataset.

However, as explained above, biases are not all due to the way in which data are collected, measured, sampled, or labelled, and so not all types of solutions can concern sample and feature selection. This is because many ML models are trained using user-generated data and parameters, which reflect any inherent (e.g., historical) biases in users (Mehrabi et al. 2021). In this way, even though they allow “correct” predictions relative to the task at hand, AI applications in genomics may end up reinforcing and exacerbating pre-existing societal prejudices, hence causing tensions between accuracy and fairness. This last concern applies both to healthcare, for example due to bias against female patients or ethnic minority groups in electronic healthcare records (Dias and Torkamani 2019), and to applications in criminal justice. This is why, outside of healthcare, some scholars have criticized advances at the intersection between AI, genomics, and law enforcement, as linking genealogy and crime may reinforce prejudices against people of colour, who are disproportionately stopped by police and overrepresented in criminal DNA databases, whilst distracting from efforts to address structural and societal inequalities that contribute to lawbreaking (Arnold 2020). The development of interpretable models could help understand the human values that underpin a machine decision-making (Dias and Torkamani 2019).

2.3 Explainability and interpretability: Trust and the right to an explanation

Definitions

The efficiency and performance of AI models used to make complex computational analyses, for example in bioinformatics, have rapidly advanced in the past few years, especially since increasing research has been conducted in Deep Learning. However, AI algorithms are now commonly referred to as “black boxes”, because to increased architecture complexity corresponds increased opaqueness and decreased transparency (Dias and Torkamani 2019). In other words, these models cause issues around interpretability and explainability.

The two terms are often used interchangeably, but it may be useful to separate them for explanatory purposes. Interpretability refers to *how* an algorithm that, for example, makes treatment recommendation based on a range of patient data in input arrives at those decisions and recommendations (Caudai et al. 2021). Those concerned with interpretability issues argue that the *how* needs to be human interpretable. Developing human interpretable algorithm can be a way to verify that the model accurately represents the problem at stake. Here, we introduce two further understandings of the term “accuracy”. An algorithm may be accurate when tested because it identifies unknown hidden relationships among training and testing data, or finds shortcuts during training, but it may underperform when applied to real-world bioinformatic tasks (Caudai et al. 2021) – see the case of the skin-cancer detecting algorithm that learnt to detect rulers in the pictures on which it was trained, as rulers are often photographed next to skin lesions that are cause of concern (Patel 2017). Moreover, an algorithm may be accurate in the sense that it performs well within the parameters set by its developers, but the problem under study may have been poorly formulated, if we do not know (or cannot include) all the variables and features (e.g., socio-economic constructs) that should be inputted into it. In this sense, the algorithm does not accurately represent the problem under study, and if the setting is clinical diagnostics, this implies that the model is unlikely to be useful to address specific individual needs. An interpretable model may help unveil the human decisions that went into defining or constraining how to study the problem in the first place.

Interpretability can be distinguished from explainability, which focuses on explaining the reasons *why* the model arrives to certain conclusions (Caudai et al. 2021). A model can be interpretable by its developers, but these need also to be able explain (in an understandable way) its decisions to end users and other stakeholders.

In the computer science literature, some also use these terms to differentiate between two different areas of Machine Learning. In this context, *interpretable ML* focuses on models that are inherently interpretable or human readable (e.g. decision trees), while *explainable ML* aims to provide post-hoc explanations for proprietary models or black box models that are incomprehensible to humans (see the section below about *Proposed mitigations from the literature*) (Rudin 2019).

Why are interpretability and explainability important in healthcare?

Problems around interpretability and explainability are not unique to healthcare, but they raise important concerns in this setting. In healthcare we are often confronted with high-risk and sensitive situations, in which answering the “why” question is not just something that patients value, but is often essential to make subsequent clinical decisions (Dias and Torkamani 2019). Without the ability to trace and explain processes of feature selection for a

cancer diagnosis, health care professionals' and patients' trust may decrease, and so adoption of AI systems may be compromised (Cudai et al. 2021). Trust-related concerns are exacerbated by the fact that questions about who is and should be considered accountable for automated decisions (Marcu, Boyd, and Bezak 2019), as well as how people harmed by such decisions can seek redress, remain open (Nuffield Council on Bioethics 2018).

The recognized importance of explainability in scenarios like recommending medical treatments or getting insurance or a loan has led to discussions around the "right to an explanation" (Wachter, Mittelstadt, and Floridi 2017; Wachter, Mittelstadt, and Russell 2017), which refers to the right of an individual to be given explanations for why an algorithm has taken a decision that affects them (Cudai et al. 2021). This is seen as a way to increase accountability and transparency of automated decision-making (Wachter, Mittelstadt, and Floridi 2017), and hence reduce the risk that AI models make decisions in ways that are discriminatory or unlawful (Ching et al. 2018), which can be more difficult with risk scoring algorithms with proprietary or poorly documented software (Breen et al. 2019). Discussions around the right to an explanation are contributing to shape the way in which the GDPR is being implemented in European countries, and have contributed to amendments to the Data Protection Bill in the UK (The Alan Turing Institute 2018).

Proposed mitigations from the literature

Concerns around interpretability and explainability have given rise to the research area of explainable AI (XAI). Cudai and colleagues describe two overarching strategies for addressing these concerns, i.e. transparency and post-hoc analysis (Cudai et al. 2021). The former works on simpler models that are not black boxes, such as decision trees, that are directly or inherently interpretable, in the sense that the model explanation is the model itself. However, the capabilities of these models are limited, and so they cannot be used for more complex computations. In these cases, post hoc analysis, which can be model-agnostic or model-dependent, needs to be performed (Cudai et al. 2021). In these cases, the black box model is analysed by using another model that approximates it, or is visualised to understand which parts of the input the model focuses on to produce its predictions. Ching and colleagues discuss a range of technical solutions to address explainability concerns of Deep Learning models, including backpropagation-based approaches, in which the signal from a target output neuron is propagated backwards to the input layer; and perturbation-based approaches, where parts of the input are changed to see the impact on the output (Ching et al. 2018). The latter would allow a doctor to query an AI model to see how its predictions about a patient's health outcome would change by changing (e.g., genomic or lifestyle) data in input (Dias and Torkamani 2019).

2.4 Changes in genomic expertise: Workforce diversity

The increasing relevance of AI to genomics has given rise to questions about the changes needed in education and training of healthcare professionals to foster a new generation of workers that can leverage on big data in their day-to-day work. Khamisy-Farah and colleagues (2021) conducted a review of the literature on the impact of recent genomic and computational/digital advances on medical education in North America, Europe, and China. This review showed that, among young doctors, there is low level of awareness of both the relevance of genomics to disciplines like medical oncology, and of digital health and big data (Khamisy-Farah et al. 2021).

These and similar studies have highlighted the importance of diversity among genomic experts. Workforce diversity refers, first, to the need for a range of different skillsets to conduct genomic science. Some authors support the integration of multidisciplinary teaching in biomedical education, through subjects like biomathematics, biostatistics, and bioinformatics; as well as teaching in machine learning, mathematical modelling, computation and simulation, to equip biomedical students with quantitative skills (Khamisy-Farah et al. 2021; Marcu, Boyd, and Bezak 2019). These considerations are echoed by the National Human Genome Research Institute (NHGRI)'s 2020 strategic vision, which encourages further knowledge exchange between genomic researchers and data scientists, as data science skills have become a prerequisite for conducting research in genomics (Green et al. 2020). Given the increased recognised importance for technologists (and scientists more generally) to consider the ethical and societal implications of their work, discussions should also be facilitated about key ethical and societal issues arising from genomic advances (Green et al. 2020), and teaching in computer science should be reformed so that ethics is integrated within the core curricula, rather than being taught as part of stand-alone modules (Zou and Schiebinger 2021). Zou and Schiebinger, in particular, recommend that ethicists become part of AI development teams, which in turn is likely to require changes in the way in which researchers are trained within universities, so to equip them with the skills to conduct interdisciplinary work (Zou and Schiebinger 2021).

Workforce diversity also refers to having team members with a range of various demographic characteristics and backgrounds, who can help broaden the types of questions we ask of genomic data, contribute to medical research with unique insights, and reflect the concerns of the populations that will be affected by technological advances (Green et al. 2020; Avni and El Kaliouby 2020). The importance of this type of diversity has been recognised both in genomics and in AI research, where initiatives have started to emerge to reduce barriers to career opportunities to women, underrepresented ethnic groups, disadvantaged populations, and disabled people. For example, in early 2022 the UK government has invested £23 million to fund scholarships for conversion courses, so to encourage young people from underrepresented groups to gain the skills necessary to get jobs in the tech sector (UK Government 2022).

2.5 Predictions of life outcomes: Genetic determinism and the spectrum of eugenics

The convergence of increasing interest in the cognitive, neural, and genetic substrates of human learning and advances in Machine Learning have led to a new strand of research that aims to use biodata to make predictions about individuals' academic achievement and inclination to other behaviours (e.g., aggression, investment behaviour) that influence life outcomes, including people's socio-economic status (Abdellaoui et al. 2019; Comfort 2018). This has been termed "social genomics" to indicate an area of research that uses large amount of data and computing power to identify genetic contributions to complex social traits (Adam 2019). Some consider studies in this area praiseworthy, because they could, for example, inform changes in education policies and contribute to deliver personalised learning, i.e. to target education to each student's DNA, and in turn improve their life chances. Findings from sociogenomics are said to be particularly useful to identify children with extra educational

needs, and to allow parents of children with conditions such as ADHD or autism to lobby for extra support in schools (Adam 2019). However, many ethical concerns have been raised.

Scientists warn that, taken individually, single-nucleotide polymorphisms (SNPs), which are common but tiny variations in the DNA, provide a small contribution to the behavioural traits with which they are found to be associated (Adam 2019). Moreover, given that any complex trait results from the interactions of tens of thousands of SNPs, as well as their interaction with environmental factors, genomic studies on complex traits only provide indication of correlations between genotype and phenotype, rather than mapping traits to their genetic causes (Comfort 2018). This also suggests that polygenic risk scores, which sum individual SNP effect into single scores, may be relevant at the level of the population, but provide little information for any specific individual (Williamson 2020); in fact, we already have better predictors of students' future educational performance, such as information about their previous academic attainment (Adam 2019). Furthermore, findings from sociogenomic studies replicate the same concerns discussed above about representation biases (section 2.2.), given that they use data primarily collected from older males of European descent, which raises questions about their relevance for people who are underrepresented in existing datasets (Adam 2019; Williamson 2020).

Ethicists are also concerned about how findings from social genomics research could be used beyond the lab, such as in schools or by employers and insurers (Comfort 2018). A major concern is that sociogenomics could encourage genetic determinism, and so lead to discriminatory practices towards some children (Williamson 2020). Genetic determinism is the view according to which our genes determine our characteristics, such as our educational performance, and so are sufficient to define our future (Wachbroit 2002). The belief that inequalities in life outcomes are genetically determined could shift away attention from efforts that target social and structural inequalities, reinforce the belief that such efforts are futile, and even reanimate old eugenic aspirations (Comfort 2018).

This is why some authors in this review argue that scientists working in social genomics have an ethical responsibility to make all the possible steps to ensure that their research findings are used to help and support, rather than leaving behind, children who are struggling (Adam 2019). One way to do so would be to use the results from genomic studies as a control variable to inform understanding of the environmental, rather than genetic, contributions to complex social traits (Adam 2019). Moreover, scientists could accompany their publications with blog posts or a list of frequently asked questions that clarify what could and could not be legitimately inferred from their papers (Comfort 2018; Adam 2019).

2.6 Cost and opportunity cost

The problem

Because healthcare resources are scarce, the integration of the new AI-enabled capabilities in genomics raises questions around opportunity cost, which is measured by “the health benefits [...] that could have been achieved had the money been spent on the next best alternative intervention or healthcare programme” (Palmer and Raftery 1999). In this review, some argue that, although implementing precision medicine may lead to increased costs due to an increase in genetic tests required, it will overall help to cut on costs associated with misdiagnoses, unsuccessful treatments, and patients' poor health (A Jamal 2021).

However, the benefits derived from decreased costs can only be fully achieved if access to genomics in healthcare is an option for all members of society, and so society benefits equitably from advances in precision medicine (Green et al. 2020; A Jamal 2021). For example, if data-intensive processing enabled by AI allows us to identify similar genomic risk profiles, but individuals with similar profiles differ in entitlements to access healthcare services, questions around the fair allocation of healthcare resources, and so about distributed justice, arise (Hummel and Braun 2020). Moreover, algorithms that recommend the best course of treatment based on a patient's data may be trained to meet certain cost efficiency standards; thus, they may recommend that treatment is not undertaken for a particular patient if others are more likely to benefit from that intervention (Cohen et al. 2014). While this may contribute to save resources or cut down costs, it may simultaneously lead to recommendations that are in conflict with healthcare professionals' responsibilities and obligations, which are usually directed towards the best interests of individual patients (Nuffield Council on Bioethics 2018).

Although some of these concerns may become more relevant as AI methods to support clinical decision-making become widespread within healthcare systems, Ching and colleagues discuss a range of costs that are already implicated by AI advances in genomics, particularly in relation to training algorithms. Producing high-performing (especially Deep Learning) models requires much computing power, which comes with high costs in terms of time, memory, and energy (Ching et al. 2018), and so can also have environmental implications. Moreover, generating ground-truth labels for classification tasks can be an expensive process, because it requires that some clinicians annotate patients' data (e.g., healthcare records) and that other clinicians validate and interpret the features constructed by the algorithms (Ching et al. 2018). To reduce costs, the stage of expert review could be skipped, but any technology that lacks this stage is likely to encounter resistance from healthcare professionals due to trust-related concerns (Ching et al. 2018).

Proposed mitigations from the literature

Recently, the more widespread use of Graphics Processing Units (GPUs) has contributed to reduce efficiency issues. This is because, differently from Central Processing Units (CPUs), GPUs make it possible to process multiple computations in parallel. However, authors in this review argue that GPUs too have limited memory, which will make it difficult to implement increasingly complex models in a single GPU unless improvements keep up with the pace with which growing biological datasets are created (Ching et al. 2018). In addition to the aforementioned cloud computing, which is already largely used in genomics, another method to address increasingly demanding computational needs is distributed computing, which allows to run multiple software components, spread across different computers, as a single system (Ching et al. 2018).

However, Jevons' paradox warns against the risk that increased efficiency in training AI models could be followed by increase in training of larger models, which may in turn counteract efforts to reduce energy costs (Weidinger et al. 2021). Moreover, even if these and other technical solutions were identified and adopted to limit costs related to training algorithms, the new capabilities that AI is enabling in genomics would not lead to concrete health benefits without a radical change in funding schemes. Studies have indeed showed that in genomics, and in science more generally, limited resources are available for implementation science, i.e. for the development of methods to promote uptake of research findings in routine healthcare practice, as much more fundings are invested for discovery-

focused research than for translational efforts (Car et al. 2019). As a consequence, “Translation from bench science to real-world practice generally averages 17 years” (Breen et al. 2019). To address issues related with the limited uptake of genomics in clinical care, the National Human Genome Research Institute (NHGRI) has advocated for the need to develop highly specialised services for genomic referrals (similarly to existing transplantation or cancer centres), whilst simultaneously aiming at increasing the availability of services more broadly, by providing specialized genomic care for example through telemedicine (Green et al. 2020)

2.7 Macro level issues: Health disparities and concentrations of power

A minority of the identified literature focuses on macro ethical issues, particularly on how the use of AI for the analysis of large amount of different types of data may impact the distribution of power across countries, and the relationship between the public and private sector in the provision of healthcare services. Stahl argues that the large computing resources and amounts of data required of AI approaches imply that only the organisations with access to such resources will be able to benefit from AI (Stahl 2021). This risks giving large amounts of power and dominance to some actors, particularly private tech companies and research institutions in high-income countries. In turn, this raises concerns about the scope for public scrutiny over AI advances in genomics (Lévesque 2019), [and questions around](#) which research will get funded. For example, although discovery and translational research on infectious diseases could have great global benefits, this research area may not get prioritized due to the fact that the greatest burden of infection disease is in low- and middle-income countries, while the biggest investments in health research come from high-income countries (Hummel and Braun 2020).

Considerations about which actors are driving advances at the intersection between AI and genomics are relevant to research on health disparities. In recent years, much research has been conducted on the difference in rates of premature death or illness not just between groups within the same country, but also between populations in low-income countries and high-income countries. Such research aims to improve our understanding of how various determinants (e.g., racism) contribute to poorer health outcomes in certain groups or populations. Despite these efforts, health disparities still exist (Breen et al. 2019). This is in part because to understand the natural, social, and economic environments that influence behaviour and health outcomes, we need more granular information than we have traditionally been able to collect, and hence bigger amount of data from different sources, including surveys and statistics, electronic medical records, genomics research, social media, and sensing from personal devices (Breen et al. 2019). Moreover, while taken individually each data type contributes to elucidate disease causes, we still have limited understanding of how these determinants interact with each other, and so how intersecting characteristics (e.g., gender, ethnicity, age, etc.) affect inequalities (Breen et al. 2019; Zou and Schiebinger 2021). Combining data science and health disparity studies thanks to data mining approaches made possible by advances in ML and DL can help address these two issues, and so make progress in translating research findings into real-world interventions. In analogy to precision medicine, Flahault and colleagues have described these efforts as “precision global health”, to refer to the use of information-driven approaches and digital technologies to improve equality in global health (Flahault et al. 2017).

However, if not done carefully, the implementation of AI to healthcare can actually lead to increased health disparities (Zou and Schiebinger 2021). This is because of representation and other bias-related issues in datasets that we described in section 2.2, which limit our capacity to understand causes of health disparities (Breen et al. 2019). An observation that is particularly pertinent to the macro scale on which this section focuses is that fewer data come from low-income countries, due to the fact that in those settings health information systems are not as powerful as in richer parts of the world. For example, high-speed internet access is still the exception in rural areas in low-income countries (Flahault et al. 2017). Combined with the need for high storage capacity and strong computing infrastructures for data sharing (see section 2.1 above), this runs the risk of creating a vicious circle, in which more data are collected from and processed in high-income countries for the benefit of the local population, hence exacerbating the health divide between richer and poorer regions of the world.

To ensure that biomedical AI benefits diverse populations, a range of solutions have been proposed. Breen and colleague advocate for the need of community-based participatory research, in which groups experiencing health disparities collaborate with researchers to identify research priorities, and so define the research questions and scope; help shape research methodologies; and contribute to data collection and interpretation, as well as dissemination of results (Breen et al. 2019). This is important because members of those groups that experience health disparities have better understanding of their health needs, and so can help inform the identification of relevant data sources that will improve our capacity to translate findings into real-world applications. Interviews with these groups could be used “to understand sensitive situations and complex life contexts experienced by vulnerable groups, and this knowledge can be used to develop quantitative instruments that are more sensitive to the meanings and interpretations of respondent reports” (Breen et al. 2019). In this way, qualitative data could be used to enhance or inform quantitative research based on big data analytics.

Zou and Schiebinger also mention policy and regulatory solutions, for example to ensure the review of training data for bias, and highlight the role that funding agencies could play in demanding the implementation of these solutions (Zou and Schiebinger 2021). Some founding agencies are already doing so. For example, the German Research Foundation has introduced guidelines for sex, gender, and diversity analysis in research proposals across disciplines (Deutsche Forschungsgemeinschaft 2021). Similarly, some conferences and peer reviewed journals have started to develop policies on diversity and inclusion, and expect that authors include considerations of the broader ethical and societal implications of their work in their publications. Zou and Schiebinger encourage these and similar initiatives, whilst also cautioning against the risk of “ethics washing”, the practice by which researchers may make unsubstantiated or misleading claims about the ethical values and benefits of their work to appear more ethical than they actually are (Floridi 2019).

Conclusion

In this literature review, we presented the current and emerging capabilities and applications enabled or enhanced by AI in genomics, as well as their ethical, societal, legal, and political ramifications. An overall consideration that is worth noting is that most of the identified literature, both in the sciences and social sciences, focused on healthcare applications of the advances at the intersection between genomics and AI. The range of scientific applications

for AI in genomic research were vast, and two major trends were evident from the literature. Firstly, Deep Learning approaches have surpassed previous machine capabilities in addressing some of the scientific questions of genomic research such as in areas of metagenomics and transcriptomics. Secondly, an increasing number of applications are directed towards the regulatory control of gene expression compared to other fields. Other noteworthy fields are that of drug discovery applications and cancer. The entirety of the scientific literature that was found was concerning these healthcare related domains, and was lacking in non-healthcare related concepts such as personality or behavioural dispositions. The increased number of publications in more recent years indicates that AI and ML research in genomics is expected to continue to rapidly increase. However, one major challenge is that, though many frameworks and tools have been developed to predict various attributes based on genomic data, these require experimental validation in order to be considered for clinical translation.

With regard to healthcare applications, the ethics and social science literature was concerned primarily with the risk of violating patients' privacy, when different types of datasets are shared for precision medicine purposes; the potential for misclassifying patients or treating certain groups of patients unfairly, when models are trained on biased datasets; and the risk of feeding into public mistrust towards AI applications, when proprietary and black-box models are used to support clinical decision-making. Only a minority of the literature on societal implications discussed applications that are relevant to law enforcement or prediction of life outcomes. Future research should investigate whether this partiality of the literature towards healthcare reflects where major research investments at the intersection between genomics and AI currently are, or where experts believe more credible progress can be made, or whether this partiality is because non-healthcare-related research, and the ethical complexities they may give rise to, have so far been neglected in the academic and public debate.

Interestingly, while the review of the scientific literature identified a range of specific areas of application of AI methods to genomics, the ethics and social science literature seemed to make more generic and well-rehearsed considerations about the issues raised by big data analytics. This suggests that this literature has not caught up with the rapid advances made in science, which makes it urgent to develop follow-up investigations on the ethical, societal, legal, and political ramifications of each research area identified at the intersection between genomics and AI. The divide between what has been found in the technical literature and in the ethics/social science literature may also be a result of disciplinary silos, which needs to be further explored to identify ways in which closer cross-disciplinary collaborations could be established to anticipate and address concerns in genomics and AI.

Moreover, the ethics discussion seems to come primarily from industrialised countries, which are likely to be the actors that are driving progress in the new genomic capabilities enabled by AI. Interestingly, the literature covers the importance of workforce diversity (see section 2.4), but it focuses on the micro-scale (i.e., the need for teams of geneticists or AI researchers that come from different ethnic backgrounds, have a range of genders, disabilities, etc.), rather than making in-depth considerations about the macro-scale, i.e. the geographical distribution of research institutions, organisations, and tech companies. This may be why we have not identified important discussions about, for example, the need for workforce with understanding and awareness of the socio-environmental characteristics of the local realities where AI-enabled genomics applications are implemented, despite evidence shows that

socio-economic factors do influence how AI models are implemented locally, and may have implications on system performance (Beede et al. 2020).

This is a symptom of a broader limitation that we identified in the review of the ethical, legal, and policy literature. While most authors discussed traditional ethical and bioethical topics (e.g., privacy, genetic discrimination and determinism, bias, trust), only a minority of the literature focused on macro ethical and political issues. In particular, a couple of papers discussed how research at the intersection between genomics and AI is likely to increase the concentration of economic power in the hands of private or public institutions in high-income countries, which have access to large computing resources and can sustain the costs involved in implementing AI-enabled practices in genomics. While the review of the science and technical literature identified interest in AI for drug discovery from pharmaceutical companies, potential implications of such an interest were not discussed in the ethics and social science literature. Similarly, no paper discussed how the increasing availability of individual data beyond healthcare (e.g., through wearable medical devices) is changing the relationship and involvement of the state and private sector in providing health services (Ada Lovelace Institute 2020); or whether private companies training their proprietary models on publicly funded datasets have any obligation towards the publics that have contributed to those databases.

Thus, we could only partially answer questions (see Appendix 1) concerned with how AI-based emerging capabilities in the analysis of the human genome will affect the distributions of economic and political power, and the delivery of healthcare; the boundary between healthcare and other domains of life; and the relationship between, and respective responsibilities of, citizens, the state and the private sector in healthcare and broader public service provision. Increasingly, investigations of macro-ethical issues, sometimes termed “decolonial AI”, are being conducted in the broader field of Artificial Intelligence, and it is essential that similar efforts are translated to the specific research area at the intersection between AI and genomics, given the relevance that advances at this intersection have on people’s health.

This literature review was the first building block towards informing an understanding of how AI is transforming the capabilities and practice of genomic science, and what such a transformation could mean for people and society. The next steps in the AI and Genomics Futures project will aim to further investigate the trends identified here, as well as address the questions that have remained unanswered, through bibliometric analyses, horizon scanning exercises, futures mapping techniques, and engagement with various experts and stakeholders from the public and private sector. The insights derived from these initiatives will help inform decision makers as they prepare for future scenarios of genomics and AI technologies.

Annex 1: Methodology

To conduct this literature review, we followed a four-stage process, adapted from (Arksey and O'Malley 2005).

Identifying the research questions

The research questions of this review were as follows:

Primary research questions:

- a. What is the current state of AI in research on the human genome, and what are the upcoming and emerging AI technologies or applications in this field?
- b. What are the ethical, societal, legal, and political impacts of the capabilities that AI is engendering in the sequence, analysis, and understanding of the human genome?

Secondary research questions:

- a. What is AI enabling (and what might it enable) genomic science to do, or to do better?
- b. What are the current, emerging, and predicted applications of the new genomic capabilities enabled by AI?
- c. How, if at all, is the increasing relevance of AI to genomics changing the way genomic science is conducted, the data and expertise needed to conduct it, and the actors involved?
- d. How will AI-based emerging capabilities in the analysis of the human genome affect:
 - Different groups within society and the distributions of risks/benefits across groups?
 - Distributions of economic and political power?
 - Societal attitudes towards the allocation of scarce resources, human agency, responsibility and desert, and how will these attitudes affect or come into tension with practices in employment, criminal justice, education and fiscal policy?
 - The delivery of healthcare, the boundary between healthcare and other domains of life, and the relationship between the citizen, state and private sector in healthcare and broader public service provision?
- e. How will emerging AI-based capabilities in the analysis of the human genome be affected by data protection norms and rules, ownership of genomic data and genomic capacity, public attitudes towards genomics and advances in gene editing technologies such as CRISPR.

Defining the search string and selecting the sources

Based on the primary and secondary research questions, we defined a range of key search terms to identify the relevant literature.

To investigate how advances in AI are changing the capabilities, applications and practice of genomic science, we used the search string in Table 1.

	Example	Alternatives
Concept 1: Emerging	Emerging	new OR novel
Concept 2: AI	Artificial Intelligence	“augmented intelligence” OR “autonomous intelligence” OR “Deep Learning” OR “Machine Learning” OR “supervised learning” OR “unsupervised learning”
Concept 3: Genomics	Genom*	Genetic*

Table 1: Key search terms to identify the literature on the capabilities that AI is engendering in the sequence, analysis, and understanding of the human genome.

We adopted an iterative process whereby we first conducted informal searches to refine our search string. We then performed the final search in two bibliographic databases (PubMed and IEEE Xplore) in April 2022 to retrieve eligible publications. We searched title, abstract, and keywords for the following terms: (emerging OR new OR novel) AND (“artificial intelligence” OR “augmented intelligence” OR “autonomous intelligence” OR “deep learning” OR “machine learning” OR “supervised learning” OR “unsupervised learning”) AND (genom* OR genetic*) NOT (“genetic algorithm” OR “antimicrobial resistance” OR “antibacterial resistance” OR “antibiotic resistance”). The NOT search terms were added to exclude results that were not relevant to the research question.

As the premier biomedical literature database, PubMed hosts comprehensive literature relating to the life sciences and healthcare and was chosen as suitable for providing literature to answer the research question. IEEE Xplore is a repository hosted by the Institute of Electrical and Electronics Engineers (IEEE) and is a database for scientific and technical content focusing on electrical engineering, computer science, and electronics. IEEE Xplore was chosen as a database to provide more technically focused specific applications of AI and ML techniques.

On the same databases, a search was also conducted to purposively identify literature on non-healthcare applications of AI in genomics (i.e. criminality, intelligence, sexual orientation etc). The search was performed with the following terms: (("genom*") AND ("artificial intelligence" OR "machine Learning" OR "deep learning") AND ("behaviour" OR "behavioural disposition*" OR "criminality" OR "personality" OR "big five personality traits" OR "intelligence" OR "IQ" OR "intelligence quotient" OR "educational attainment" OR "life outcomes" OR "socioeconomic status" OR "performance" OR "endurance" OR "physical traits" OR "sexual orientation"))).

To identify relevant academic and grey literature focused on the ethical, societal, legal, and political ramifications of the capabilities AI is engendering in genomics, the following search terms were used.

	Example	Alternatives
Concept 1: Artificial Intelligence	“Artificial intelligence”	AI OR “Machine learning” OR ML OR “Deep learning”
Concept 2: Genomics	Genetic*	genomic* OR “genetic medicine” OR “genomic medicine” OR “human genetic*” OR “medical genetic*” OR “DNA sequencing” OR “genom* sequencing”

Concept 3: Features (the ethical, societal, legal, and political features impacting the capabilities that AI is engendering in the analysis of the human genome)	Ethic*	bioethic* OR law OR legal OR polic* OR social
Concept 4: Specific ethical and legal issues	Bias	divers* OR fair* OR justice OR privacy OR confidentiality OR trust OR responsibility OR accountability OR "resource allocation" OR explainability OR interpretability OR "data protection" OR "data ownership" OR "data sharing"

Table 2: Key search terms to identify the literature on the ethical, societal, legal, and political impacts of the capabilities that AI is engendering in the sequence, analysis, and understanding of the human genome.

We adopted an iterative process whereby we first conducted informal searches, informed by the results from the first part of the literature review, to refine our search string. We then performed the final search in two bibliographic databases (Web of Science and Scopus) in April 2022 to retrieve eligible publications. We searched title, abstract, and keywords for the following terms: ("artificial intelligence" OR AI OR "Machine learning" OR ML OR "Deep learning") AND (genetic* OR genomic* OR "genetic medicine" OR "genomic medicine" OR "human genetic*" OR "medical genetic*" OR "DNA sequencing" OR "genom* sequencing") AND (ethic* OR bioethic* OR law OR legal OR polic* OR social) AND (bias OR divers* OR fair* OR justice OR privacy OR confidentiality OR trust OR responsibility OR accountability OR "resource allocation" OR explainability OR interpretability OR "data protection" OR "data ownership" OR "data sharing") NOT ("genetic programming" OR "genetic algorithm*" OR "genetic fuzzy-systems"). The NOT search terms were added to exclude results that were not relevant to the research questions.

The rationale for selecting Web of Science and Scopus to search for relevant literature was that these two databases have a cross-disciplinary scope that covers science, technology, and social sciences research. Such a wide scope seemed to be a good fit for a literature review at the intersection between AI, genomics, and societal implications.

Because advances at the intersection between AI and genomics are relatively recent, for the part of the literature review focused on the ethical, societal, legal, and political ramifications, we also looked into grey literature, such as reports from relevant organizations, regulatory documents, and conference proceedings, which present ongoing research that has not yet been formally published or is not published in standard academic fora. In particular, grey literature was searched in the following sources:

- Google.com
- Websites of relevant organizations and institutions, particularly website sections like "Documents", "Reports", "Policy", and "Library". Examples are provided below:
 - [The European Society of Human Genetics](#)
 - [National Human Genome Research Institute](#)
 - [Other genetic research organisations](#) in the world
 - [DeepMind](#)
 - [The Alan Turing Institute](#)

- [Nuffield Council on Bioethics](#)
- [Ada Lovelace Institute](#)
- Conference papers/proceedings were retrieved from Web of Science and Scopus.

Defining criteria for literature selection and extracting data

Literature retrieved from the databases was downloaded on a citation management software (i.e., Zotero) and duplicates were deleted. The literature was screened for relevance by using the selection (inclusion/exclusion) criteria outlined below. In the first round of screening, we looked into the titles, abstracts, and keywords of the identified literature to exclude irrelevant items, while in the second round, we looked into the full text. Discussions were held among the research team members to resolve cases in which we were uncertain whether to include or exclude certain items. The selection criteria were also used to determine which grey literature to include.

Inclusion criteria

- Peer reviewed journal articles, book chapters, grey literature
- Published since 2017
- Published in English
- Discussing the *ethical, societal, legal, and political impacts* of using AI for the analysis of the human genome (for the part of the literature review focused on societal implications)
- Focused on *human genome*
- Focused on genomic *analysis* rather than genomic *editing*
- Focused on healthcare and non-healthcare settings
- Worldwide
- Accessible (i.e. reviewers can retrieve full document)

Exclusion criteria

- Books
- Older than 5 years
- Inaccessible in English
- Lacking relevance to AI or genomic science
- Lacking reference to the *ethical, societal, legal, and political impacts* of using AI for the analysis of the human genome (for the part of the literature review focused on societal implications)
- Not focused on *humans*
- Focused on genomic *editing* rather than genomic *analysis*
- Inaccessible (i.e. reviewers cannot retrieve full document)

The literature that matched the inclusion and exclusion criteria of this review was saved on a citation management software (i.e., Zotero) for data extraction.

Summarizing and reporting results

The information extracted from the various pieces of literature were grouped thematically (Braun and Clarke 2006). This means that, across different pieces of literature, we identified similar overarching themes that outline the current and emerging capabilities that AI is

engendering in the sequence, analysis, and understanding of the human genome, and their key ethical, societal, legal, and political impacts. Themes that outlined the ethical, societal, legal, and political factors affecting (rather than resulting from) emerging AI-based capabilities in the analysis of the human genome were also noted. The research team discussed these themes at regular meetings to develop a coherent narrative that could answer the primary and secondary research questions of the literature review.

Annex 2: Glossary

Big Data = the rapidly escalating volume, velocity (torrent of information) and variety (heterogeneity) of data produced and stored (Degeling et al. 2020; Breen et al. 2019).

Big Data Analytics = the process of collecting, organising and analysing large datasets, to discover patterns and generate useful, actionable information (Garattini et al. 2019)

Precision medicine = An approach to individualize the prevention, diagnosis and treatment of various diseases to each individual patient. By collecting and combining a large amount of different types of data (genetic, lifestyle and environmental factors), it aims to provide the right therapy to the right patient at the right time, while also increasing treatment response and reducing side effects (Erdmann, Rehmann-Sutter, and Bozzaro 2021).

Single nucleotide polymorphism = single nucleotide base change in a genome

Enhancers = short region of DNA that has heightened likelihood of protein binding, which in turn increases the likelihood of initiating transcription

Promoters = short region of DNA that initiates the process of transcription of a nearby gene

Introns = the sequence of DNA that corresponds to the RNA that is removed from the RNA transcript before translation. Intron removal is part of the process of RNA splicing, from which mature messenger RNA is created, ready for translation into protein

DNA = deoxyribonucleic acid. DNA in the genome is comprised of a double helix made of four nucleotide bases and provides the genetic instructions for the entirety of an organism's living processes

RNA = ribonucleic acid. Generated from DNA during the process of transcription. Mature RNA is then translated to be converted into protein.

Mutagenesis = changes to an individual's genome. This can be due to a variety of causes, such as radiation or ongoing errors in DNA replication.

Radiogenomic = the intersection between radiological images and underlying gene expression

INDELS = insertion or deletions of DNA bases

Bibliography

- A Jamal, A.R. 2021. "Precision Medicine: Making It Happen for Malaysia." *Malaysian Journal of Medical Sciences* 28 (3): 1–4. <https://doi.org/10.21315/mjms2021.28.3.1>.
- A. P. Athreya, A. J. Gaglio, J. Cairns, K. R. Kalari, R. M. Weinshilboum, L. Wang, Z. T. Kalbarczyk, and R. K. Iyer. 2018. "Machine Learning Helps Identify New Drug Mechanisms in Triple-Negative Breast Cancer." *IEEE Transactions on NanoBioscience* 17 (3): 251–59. <https://doi.org/10.1109/TNB.2018.2851997>.
- Abdellaoui, Abdel, David Hugh-Jones, Loic Yengo, Kathryn E. Kemper, Michel G. Nivard, Laura Veul, Yan Holtz, et al. 2019. "Genetic Correlates of Social Stratification in Great Britain." *Nature Human Behaviour* 3 (12): 1332–42. <https://doi.org/10.1038/s41562-019-0757-5>.
- Ada Lovelace Institute. 2020. "The Data Will See You Now. Datafication and the Boundaries of Health." <https://www.adalovelaceinstitute.org/report/the-data-will-see-you-now/>.
- Adam, David. 2019. "The Promise and Peril of the New Science of Social Genomics." *Nature* 574 (7779): 618–21.
- Ahmed, Zeeshan. 2020. "Practicing Precision Medicine with Intelligently Integrative Clinical and Multi-Omics Data Analysis." *Human Genomics* 14 (1): 35. <https://doi.org/10.1186/s40246-020-00287-z>.
- Alaidarous, MA. 2020. "The Emergence of New Trends in Clinical Laboratory Diagnosis." *SAUDI MEDICAL JOURNAL* 41 (11): 1175–80. <https://doi.org/10.15537/smj.2020.11.25455>.
- Alrefaei, A.F., Y.M. Hawsawi, D. Almaleki, T. Alafif, F.A. Alzahrani, and M.A. Bakhrebah. 2022. "Genetic Data Sharing and Artificial Intelligence in the Era of Personalized Medicine Based on a Cross-sectional Analysis of the Saudi Human Genome Program." *Scientific Reports* 12 (1). <https://doi.org/10.1038/s41598-022-05296-7>.
- Altae-Tran, Han, Bharath Ramsundar, Aneesh S. Pappu, and Vijay Pande. 2017. "Low Data Drug Discovery with One-Shot Learning." *ACS Central Science* 3 (4): 283–93. <https://doi.org/10.1021/acscentsci.6b00367>.
- Arksey, Hilary, and Lisa O'Malley. 2005. "Scoping Studies: Towards a Methodological Framework." *International Journal of Social Research Methodology* 8 (1): 19–32. <https://doi.org/10.1080/1364557032000119616>.
- Arnold, C. 2020. "The Controversial Company Using DNA to Sketch the Faces of Criminals." *Nature* 585 (7824): 178–81. <https://doi.org/10.1038/d41586-020-02545-5>.
- Avni, Ronit, and Rana El Kaliouby. 2020. "Here's Why AI Needs a More Diverse Workforce." *World Economic Forum* (blog). September 21, 2020. <https://www.weforum.org/agenda/2020/09/ai-needs-diverse-workforce/>.
- Azencott, C.-A. 2018. "Machine Learning and Genomics: Precision Medicine versus Patient Privacy." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376 (2128): 20170350. <https://doi.org/10.1098/rsta.2017.0350>.
- Batool, Maria, Bilal Ahmad, and Sangdun Choi. 2019. "A Structure-Based Drug Discovery Paradigm." *International Journal of Molecular Sciences* 20 (11). <https://doi.org/10.3390/ijms20112783>.
- Beede, Emma, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M Vardoulakis. 2020. "A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy." In , 1–12.

- Berkman, Benjamin E, Wynter K Miller, and Christine Grady. 2018. "Is It Ethical to Use Genealogy Data to Solve Crimes?" *Annals of Internal Medicine* 169 (5): 333–34. <https://doi.org/10.7326/M18-1348>.
- Braun, Virginia, and Victoria Clarke. 2006. "Using Thematic Analysis in Psychology." *Qualitative Research in Psychology* 3 (2): 77–101. <https://doi.org/10.1191/1478088706qp063oa>.
- Breen, N, D Berrigan, JS Jackson, DWS Wong, FB Wood, JC Denny, XZ Zhang, and PE Bourne. 2019. "Translational Health Disparities Research in a Data-Rich World." *HEALTH EQUITY* 3 (1): 588–600. <https://doi.org/10.1089/heq.2019.0042>.
- Car, J, A Sheikh, P Wicks, and MS Williams. 2019. "Beyond the Hype of Big Data and Artificial Intelligence: Building Foundations for Knowledge and Wisdom." *BMC MEDICINE* 17 (July). <https://doi.org/10.1186/s12916-019-1382-x>.
- Caudai, Claudia, Antonella Galizia, Filippo Geraci, Loredana Le Pera, Veronica Morea, Emanuele Salerno, Allegra Via, and Teresa Colombo. 2021. "AI Applications in Functional Genomics." *Computational and Structural Biotechnology Journal* 19: 5762–90. <https://doi.org/10.1016/j.csbj.2021.10.009>.
- Chen, Irene Y., Peter Szolovits, and Marzyeh Ghassemi. 2019. "Can AI Help Reduce Disparities in General Medical and Mental Health Care?" *AMA Journal of Ethics* 21 (2): E167-179. <https://doi.org/10.1001/amajethics.2019.167>.
- Chen, Liang, Liisa Heikkinen, Changliang Wang, Yang Yang, Huiyan Sun, and Garry Wong. 2019. "Trends in the Development of MiRNA Bioinformatics Tools." *Briefings in Bioinformatics* 20 (5): 1836–52. <https://doi.org/10.1093/bib/bby054>.
- Ching, T, DS Himmelstein, BK Beaulieu-Jones, AA Kalinin, BT Do, GP Way, E Ferrero, et al. 2018. "Opportunities and Obstacles for Deep Learning in Biology and Medicine." *JOURNAL OF THE ROYAL SOCIETY INTERFACE* 15 (141). <https://doi.org/10.1098/rsif.2017.0387>.
- Cohen, I. Glenn, Ruben Amarasingham, Anand Shah, Bin Xie, and Bernard Lo. 2014. "The Legal And Ethical Concerns That Arise From Using Complex Predictive Analytics In Health Care." *Health Affairs* 33 (7): 1139–47. <https://doi.org/10.1377/hlthaff.2014.0048>.
- Comfort, Nathaniel. 2018. "Sociogenomics Is Opening a New Door to Eugenics." *MIT Review Technology*, 2018. <https://www.technologyreview.com/2018/10/23/139420/sociogenomics-is-opening-a-new-door-to-eugenics/>.
- Dastin, Jeffrey. 2018. "Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women." *Reuters*, October 11, 2018. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>.
- Deutsche Forschungsgemeinschaft. 2021. "Relevance of Sex, Gender and Diversity in Research." 2021. https://www.dfg.de/en/research_funding/principles_dfg_funding/diversity_dimensions/index.html.
- Dias, Raquel, and Ali Torkamani. 2019. "Artificial Intelligence in Clinical and Genomic Diagnostics." *Genome Medicine* 11 (1): 70. <https://doi.org/10.1186/s13073-019-0689-8>.
- Dlamini, Zodwa, Flavia Zita Francies, Rodney Hull, and Rahaba Marima. 2020. "Artificial Intelligence (AI) and Big Data in Cancer and Precision Oncology." *Computational and Structural Biotechnology Journal* 18: 2300–2311. <https://doi.org/10.1016/j.csbj.2020.08.019>.
- Eggel, Ivan, Roger Schaer, and Henning Muller. 2018. "Distributed Container-Based Evaluation Platform for Private/Large Datasets." In *2018 17th International Symposium on Parallel and Distributed Computing (ISPDC)*, 93–100. Geneva: IEEE. <https://doi.org/10.1109/ISPDC2018.2018.00022>.

European Union. 2010. *Charter of Fundamental Rights of the European Union. Official Journal of the European Union C83*. Vol. 53. Brussels: European Union.

— — —. 2018. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation)*, OJ 2016 L 119/1. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>.

F. Ahsan, A. Drouin, F. Laviolette, D. Precup, and M. Blanchette. 2020. "Phylogenetic Manifold Regularization: A Semi-Supervised Approach to Predict Transcription Factor Binding Sites." In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 62–66. <https://doi.org/10.1109/BIBM49941.2020.9313437>.

F. Celesti, A. Celesti, J. Wan, and M. Villari. 2018. "Why Deep Learning Is Changing the Way to Approach NGS Data Processing: A Review." *IEEE Reviews in Biomedical Engineering* 11: 68–76. <https://doi.org/10.1109/RBME.2018.2825987>.

Flahault, A, A Geissbuhler, I Guessous, PJ Guerin, I Bolon, M Salathe, and G Escher. 2017. "Precision Global Health in the Digital Age." *SWISS MEDICAL WEEKLY* 147 (April). <https://doi.org/10.4414/smw.2017.14423>.

Floridi, Luciano. 2019. "Translating Principles into Practices of Digital Ethics: Five Risks of Being Unethical." *Philosophy & Technology* 32 (2): 185–93. <https://doi.org/10.1007/s13347-019-00354-x>.

Gladding, Patrick A., Suzanne Loader, Kevin Smith, Erica Zarate, Saras Green, Silas Villas-Boas, Phillip Shepherd, et al. 2021. "Multiomics, Virtual Reality and Artificial Intelligence in Heart Failure." *Future Cardiology* 17 (8): 1335–47. <https://doi.org/10.2217/fca-2020-0225>.

Graham, Sarah A., Ellen E. Lee, Dilip V. Jeste, Ryan Van Patten, Elizabeth W. Twamley, Camille Nebeker, Yasunori Yamada, Ho-Cheol Kim, and Colin A. Depp. 2020. "Artificial Intelligence Approaches to Predicting and Detecting Cognitive Decline in Older Adults: A Conceptual Review." *Psychiatry Research* 284 (February): 112732. <https://doi.org/10.1016/j.psychres.2019.112732>.

Green, Eric D., Chris Gunter, Leslie G. Biesecker, Valentina Di Francesco, Carla L. Easter, Elise A. Feingold, Adam L. Felsenfeld, et al. 2020. "Strategic Vision for Improving Human Health at The Forefront of Genomics." *Nature* 586 (7831): 683–92. <https://doi.org/10.1038/s41586-020-2817-4>.

Gupta, Rohan, Devesh Srivastava, Mehar Sahu, Swati Tiwari, Rashmi K. Ambasta, and Pravir Kumar. 2021. "Artificial Intelligence to Deep Learning: Machine Intelligence Approach for Drug Discovery." *Molecular Diversity* 25 (3): 1315–60. <https://doi.org/10.1007/s11030-021-10217-3>.

He, Ying, Zhen Shen, Qinhu Zhang, Siguo Wang, and De-Shuang Huang. 2021. "A Survey on Deep Learning in DNA/RNA Motif Mining." *Briefings in Bioinformatics* 22 (4). <https://doi.org/10.1093/bib/bbaa229>.

Hern, Alex. 2017. "Royal Free Breached UK Data Law in 1.6m Patient Deal with Google's DeepMind." *The Guardian*, 2017. <https://www.theguardian.com/technology/2017/jul/03/google-deepmind-16m-patient-royal-free-deal-data-protection-act>.

Hummel, P, and M Braun. 2020. "Just Data? Solidarity and Justice in Data-Driven Medicine." *LIFE SCIENCES SOCIETY AND POLICY* 16 (1). <https://doi.org/10.1186/s40504-020-00101-7>.

Ivanov, V.V., D.M. Matsepuro, and T.V. Trubnikova. 2022. *National Judicial Bodies in Search of a Balance of Public and Private Interests. Lect. Notes Networks Syst.* Vol. 345 LNNS. Springer Science and Business Media Deutschland GmbH. https://doi.org/10.1007/978-3-030-89708-6_14.

- Jiang, Yahui, Meng Yang, Shuhao Wang, Xiangchun Li, and Yan Sun. 2020. "Emerging Role of Deep Learning-Based Artificial Intelligence in Tumor Pathology." *Cancer Communications (London, England)* 40 (4): 154–66. <https://doi.org/10.1002/cac2.12012>.
- K. Paranjape, M. Schinkel, and P. Nanayakkara. 2020. "Short Keynote Paper: Mainstreaming Personalized Healthcare—Transforming Healthcare Through New Era of Artificial Intelligence." *IEEE Journal of Biomedical and Health Informatics* 24 (7): 1860–63. <https://doi.org/10.1109/JBHI.2020.2970807>.
- Kalinin, Alexandr A., Gerald A. Higgins, Narathip Reamaroon, Sayedmohammadreza Soroushmehr, Ari Allyn-Feuer, Ivo D. Dinov, Kayvan Najarian, and Brian D. Athey. 2018. "Deep Learning in Pharmacogenomics: From Gene Regulation to Patient Stratification." *Pharmacogenomics* 19 (7): 629–50. <https://doi.org/10.2217/pgs-2018-0008>.
- Kearns, Michael 1971-, and Aaron 1984- Roth. 2020. *The Ethical Algorithm : The Science of Socially Aware Algorithm Design*. 1 online resource (218 pages) : illustrations, map vols. New York, NY, United States of America: Oxford University Press. <https://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=2255401>.
- Khamisy-Farah, R., P. Gilbey, L.B. Furstenau, M.K. Sott, R. Farah, M. Viviani, M. Bisogni, J.D. Kong, R. Ciliberti, and N.L. Bragazzi. 2021. "Big Data for Biomedical Education with a Focus on the Covid-19 Era: An Integrative Review of the Literature." *International Journal of Environmental Research and Public Health* 18 (17). <https://doi.org/10.3390/ijerph18178989>.
- Kimberly, R.P. 2019. "Translational Research—for the Individual and the Community." *Journal of Health Care for the Poor and Underserved* 30 (4): 79–85. <https://doi.org/10.1353/hpu.2019.0118>.
- Köhler, Sebastian, Leigh Carmody, Nicole Vasilevsky, Julius O B Jacobsen, Daniel Danis, Jean-Philippe Gourdi, Michael Gargano, et al. 2019. "Expansion of the Human Phenotype Ontology (HPO) Knowledge Base and Resources." *Nucleic Acids Research* 47 (Database issue): D1018–27. <https://doi.org/10.1093/nar/gky1105>.
- Koumakis, Lefteris. 2020. "Deep Learning Models in Genomics; Are We There Yet?" *Computational and Structural Biotechnology Journal* 18: 1466–73. <https://doi.org/10.1016/j.csbj.2020.06.017>.
- L. Li, Y. Huang, Y. Han, and J. Jiang. 2021. "Use of Deep Learning Genomics to Discriminate Alzheimer's Disease and Healthy Controls." In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 5788–91. <https://doi.org/10.1109/EMBC46164.2021.9629983>.
- Lévesque, Maroussia. 2019. "Looking Back to the Future of AI." *Indigenous AI* (blog). January 13, 2019. <https://www.indigenous-ai.net/looking-back-to-the-future-of-ai>.
- Li, Hui, Cory R. Brouwer, and Weijun Luo. 2022. "A Universal Deep Neural Network for In-Depth Cleaning of Single-Cell RNA-Seq Data." *Nature Communications* 13 (1): 1901. <https://doi.org/10.1038/s41467-022-29576-y>.
- M. Aledhari, M. D. Pierro, M. Hefeida, and F. Saeed. 2021. "A Deep Learning-Based Data Minimization Algorithm for Fast and Secure Transfer of Big Genomic Datasets." *IEEE Transactions on Big Data* 7 (2): 271–84. <https://doi.org/10.1109/TBDATA.2018.2805687>.
- M. Bahi and M. Batouche. 2018. "Deep Semi-Supervised Learning for DTI Prediction Using Large Datasets and H2O-Spark Platform." In *2018 International Conference on Intelligent Systems and Computer Vision (ISCV)*, 1–7. <https://doi.org/10.1109/ISACV.2018.8354081>.
- M. Zeng, C. Lu, Z. Fei, F. -X. Wu, Y. Li, J. Wang, and M. Li. 2021. "DMFLDA: A Deep Learning Framework for Predicting lncRNA–Disease Associations." *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 18 (6): 2353–63. <https://doi.org/10.1109/TCBB.2020.2983958>.

- MacEachern, Sarah J., and Nils D. Forkert. 2021. "Machine Learning for Precision Medicine." *Genome* 64 (4): 416–25. <https://doi.org/10.1139/gen-2020-0131>.
- Marcu, L.G., C. Boyd, and E. Bezak. 2019. "Current Issues Regarding Artificial Intelligence in Cancer and Health Care. Implications for Medical Physicists and Biomedical Engineers." *Health and Technology* 9 (4): 375–81. <https://doi.org/10.1007/s12553-019-00348-x>.
- Mehrabi, Ninareh, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. "A Survey on Bias and Fairness in Machine Learning." *ACM Computing Surveys* 54 (6): 1–35. <https://doi.org/10.1145/3457607>.
- Mobadersany, Pooya, Safoora Yousefi, Mohamed Amgad, David A. Gutman, Jill S. Barnholtz-Sloan, José E. Velázquez Vega, Daniel J. Brat, and Lee A. D. Cooper. 2018. "Predicting Cancer Outcomes from Histology and Genomics Using Convolutional Networks." *Proceedings of the National Academy of Sciences of the United States of America* 115 (13): E2970–79. <https://doi.org/10.1073/pnas.1717139115>.
- N. Sedaghat, M. Fathy, M. H. Modarressi, and A. Shojaie. 2018. "Combining Supervised and Unsupervised Learning for Improved miRNA Target Prediction." *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 15 (5): 1594–1604. <https://doi.org/10.1109/TCBB.2017.2727042>.
- Nellaker, C, FS Alkuraya, G Baynam, RA Bernier, FPJ Bernier, V Boulangerw, M Brudno, et al. 2019. "Enabling Global Clinical Collaborations on Identifiable Patient Data: The Minerva Initiative." *FRONTIERS IN GENETICS* 10 (July). <https://doi.org/10.3389/fgene.2019.00611>.
- Niet, AG van der, and A Bleakley. 2021. "Where Medical Education Meets Artificial Intelligence: 'Does Technology Care?'" *MEDICAL EDUCATION* 55 (1): 30–36. <https://doi.org/10.1111/medu.14131>.
- Nuffield Council on Bioethics. 2018. "Artificial Intelligence (AI) in Healthcare and Research." <https://www.nuffieldbioethics.org/publications/ai-in-healthcare-and-research>.
- Palmer, Stephen, and James Raftery. 1999. "Opportunity Cost." *BMJ* 318 (7197): 1551. <https://doi.org/10.1136/bmj.318.7197.1551>.
- Patel, Neel V. 2017. "Why Doctors Aren't Afraid of Better, More Efficient AI Diagnosing Cancer." *Daily Beast*, December 22, 2017. <https://www.thedailybeast.com/why-doctors-arent-afraid-of-better-more-efficient-ai-diagnosing-cancer>.
- Peng, Zhouying, Yumin Wang, Yaxuan Wang, Sijie Jiang, Ruohao Fan, Hua Zhang, and Weihong Jiang. 2021. "Application of Radiomics and Machine Learning in Head and Neck Cancers." *International Journal of Biological Sciences* 17 (2): 475–86. <https://doi.org/10.7150/ijbs.55716>.
- Q. Zhang, W. Yu, K. Han, A. K. Nandi, and D. -S. Huang. 2021. "Multi-Scale Capsule Network for Predicting DNA-Protein Binding Sites." *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 18 (5): 1793–1800. <https://doi.org/10.1109/TCBB.2020.3025579>.
- R. Yilmaz and A. E. Pusane. 2020. "Deep Learning Based Automatic Modulation Classification in the Case of Carrier Phase Shift." In *2020 43rd International Conference on Telecommunications and Signal Processing (TSP)*, 354–57. <https://doi.org/10.1109/TSP49548.2020.9163509>.
- Radakovich, Nathan, Matthew Cortese, and Aziz Nazha. 2020. "Acute Myeloid Leukemia and Artificial Intelligence, Algorithms and New Scores." *Best Practice & Research. Clinical Haematology* 33 (3): 101192. <https://doi.org/10.1016/j.beha.2020.101192>.
- Ramjee, D, and K Ringrose. 2020. "THE CHALLENGES OF FORENSIC GENEALOGY: DIRTY DATA, ELECTRONIC EVIDENCE, AND PRIVACY CONCERNS." *DENVER LAW REVIEW* 98 (1): 157–95.

- Rider, Nicholas L., Renganathan Srinivasan, and Paneez Khoury. 2020. "Artificial Intelligence and the Hunt for Immunological Disorders." *Current Opinion in Allergy and Clinical Immunology* 20 (6): 565–73. <https://doi.org/10.1097/ACI.0000000000000691>.
- Rodolfa, Kit T., Hemank Lamba, and Rayid Ghani. 2021. "Empirical Observation of Negligible Fairness–Accuracy Trade-Offs in Machine Learning for Public Policy." *Nature Machine Intelligence* 3 (10): 896–904. <https://doi.org/10.1038/s42256-021-00396-x>.
- Rudin, Cynthia. 2019. "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead." *Nature Machine Intelligence* 1 (5): 206–15.
- S. Mohanty, S. Mahapatra, and T. Swarnkar. 2021. "Deep Learning Techniques for Pre-miRNA Prediction: Current Challenges and Future Directions." In *2021 9th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO)*, 1–4. <https://doi.org/10.1109/ICRITO51393.2021.9596072>.
- S. Rajpal, M. Agarwal, V. Kumar, A. Gupta, and N. Kumar. 2021. "Triphasic DeepBRCA-A Deep Learning-Based Framework for Identification of Biomarkers for Breast Cancer Stratification." *IEEE Access* 9: 103347–64. <https://doi.org/10.1109/ACCESS.2021.3093616>.
- Shah, Pratik, Francis Kendall, Sean Khozin, Ryan Goosen, Jianying Hu, Jason Laramie, Michael Ringel, and Nicholas Schork. 2019. "Artificial Intelligence and Machine Learning in Clinical Development: A Translational Perspective." *NPJ Digital Medicine* 2: 69. <https://doi.org/10.1038/s41746-019-0148-3>.
- Stahl, Bernd Carsten. 2021. "Ethical Issues of AI." *Artificial Intelligence for a Better Future: An Ecosystem Perspective on the Ethics of AI and Emerging Digital Technologies*, March, 35–53. https://doi.org/10.1007/978-3-030-69978-9_4.
- Suresh, Harini, and John V. Guttag. 2021. "A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle." In *Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–9. <https://doi.org/10.1145/3465416.3483305>.
- The Alan Turing Institute. 2018. "A Right to Explanation." 2018. <https://www.turing.ac.uk/research/impact-stories/a-right-to-explanation>.
- The US Department of Justice. 2019. "Interim Policy on Forensic Genetic Genealogical DNA Analysis and Searching." [https://www.justice.gov/olp/page/file/1204386/download#:~:text=The%20purpose%20of%20this%20interim,searching%20\('FGGS'\)](https://www.justice.gov/olp/page/file/1204386/download#:~:text=The%20purpose%20of%20this%20interim,searching%20('FGGS')).
- Tran, Khoa A., Olga Kondrashova, Andrew Bradley, Elizabeth D. Williams, John V. Pearson, and Nicola Waddell. 2021. "Deep Learning in Cancer Diagnosis, Prognosis and Treatment Selection." *Genome Medicine* 13 (1): 152. <https://doi.org/10.1186/s13073-021-00968-x>.
- UK Government. 2022. "£23 Million to Boost Skills and Diversity in AI Jobs." February 10. <https://www.gov.uk/government/news/23-million-to-boost-skills-and-diversity-in-ai-jobs>.
- United States Government. 2008. *Public Law 110 - 233 - Genetic Information Nondiscrimination Act of 2008*. <https://www.govinfo.gov/app/details/PLAW-110publ233>.
- Vatanever, Sezen, Avner Schlessinger, Daniel Wacker, H. Ümit Kaniskan, Jian Jin, Ming-Ming Zhou, and Bin Zhang. 2021. "Artificial Intelligence and Machine Learning-Aided Drug Discovery in Central Nervous System Diseases: State-of-the-Arts and Future Directions." *Medicinal Research Reviews* 41 (3): 1427–73. <https://doi.org/10.1002/med.21764>.
- W. Yang, S. M. Weissman, R. Guan, J. Zhang, A. Canaan, and M. Q. Yang. 2017. "Invited Talk: Developing Deep Multi-Source Intelligent Learning That Facilitates the Advancement of Single Cell Genomics Research." In *2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 750–53. <https://doi.org/10.1109/BIBM.2017.8217749>.

- Wachbroit, Robert. 2002. "Genetic Determinism, Genetic Reductionism, and Genetic Essentialism." *Encyclopedia of Ethical, Legal and Policy Issues in Biotechnology*.
- Wachter, Sandra, Brent Mittelstadt, and Luciano Floridi. 2017. "Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation." *International Data Privacy Law* 7 (2): 76–99.
- Wachter, Sandra, Brent Mittelstadt, and Chris Russell. 2017. "Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR." *Harv. JL & Tech.* 31: 841.
- Weidinger, Laura, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, and Atoosa Kasirzadeh. 2021. "Ethical and Social Risks of Harm from Language Models." *ArXiv Preprint ArXiv:2112.04359*.
- Williamson, B. 2020. "Bringing up the Bio-Datafied Child: Scientific and Ethical Controversies over Computational Biology in Education." *ETHICS AND EDUCATION* 15 (4): 444–63. <https://doi.org/10.1080/17449642.2020.1822631>.
- Wong, Chi Wah, and Ammar Chaudhry. 2020. "Radiogenomics of Lung Cancer." *Journal of Thoracic Disease* 12 (9): 5104–9. <https://doi.org/10.21037/jtd-2019-pitd-10>.
- X. Chen, X. Lin, Q. Shen, and X. Qian. 2021. "Combined Spiral Transformation and Model-Driven Multi-Modal Deep Learning Scheme for Automatic Prediction of TP53 Mutation in Pancreatic Cancer." *IEEE Transactions on Medical Imaging* 40 (2): 735–47. <https://doi.org/10.1109/TMI.2020.3035789>.
- Z. Zeng, Y. Deng, X. Li, T. Naumann, and Y. Luo. 2019. "Natural Language Processing for EHR-Based Computational Phenotyping." *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 16 (1): 139–53. <https://doi.org/10.1109/TCBB.2018.2849968>.
- Zampieri, Guido, Supreeta Vijayakumar, Elisabeth Yaneske, and Claudio Angione. 2019. "Machine and Deep Learning Meet Genome-Scale Metabolic Modeling." *PLoS Computational Biology* 15 (7): e1007084. <https://doi.org/10.1371/journal.pcbi.1007084>.
- Zeng, Xiangxiang, Siyi Zhu, Xiangrong Liu, Yadi Zhou, Ruth Nussinov, and Feixiong Cheng. 2019. "DeepDR: A Network-Based Deep Learning Approach to in Silico Drug Repositioning." *Bioinformatics (Oxford, England)* 35 (24): 5191–98. <https://doi.org/10.1093/bioinformatics/btz418>.
- Zou, J, and L Schiebinger. 2021. "Ensuring That Biomedical AI Benefits Diverse Populations." *EBIOMEDICINE* 67 (May). <https://doi.org/10.1016/j.ebiom.2021.103358>.