

Opportunities and challenges for interpreting rare variation in clinically important genes

Gregory McInnes,^{1,16} Andrew G. Sharo,^{2,16} Megan L. Koleske,^{3,16} Julia E. H. Brown,^{4,5,16} Matthew Norstad,^{4,5} Aashish N. Adhikari,⁶ Sheng Wang,⁷ Steven Brenner,^{6,8,9,10} Jodi Halpern,¹⁰ Barbara A. Koenig,^{4,5,6,11,12} David C. Magnus,¹³ Renata C. Gallagher,^{6,14} Kathleen M. Giacomini,³ Russ B. Altman^{15*}

Abstract

Genome sequencing is enabling precision medicine—tailoring treatment to the unique constellation of variants in an individual’s genome. The impact of recurrent pathogenic variants is often understood, leaving a long tail of rare genetic variants that are uncharacterized. The problem of uncharacterized rare variation is especially acute when it occurs in genes of known clinical importance with functionally consequent frequent variants and associated mechanisms. Variants of unknown significance (VUS) in these genes are discovered at a rate that outpaces current ability to classify them using databases of previous cases, experimental evaluation, and computational predictors. Clinicians are thus left without guidance about the significance of variants that may have actionable consequences. Computational prediction of the impact of rare genetic variation is increasingly becoming an important capability. In this paper, we review the technical and ethical challenges of interpreting the function of rare variants in two settings: inborn errors of metabolism in newborns, and pharmacogenomics. We propose a framework for a genomic learning healthcare system with an initial focus on early-onset treatable disease in newborns and actionable pharmacogenomics. We argue that (1) a genomic learning healthcare system must allow for continuous collection and assessment of rare variants, (2) emerging machine learning methods will enable algorithms to predict the clinical impact of rare variants on protein function, and (3) ethical considerations must inform the construction and deployment of all rare-variation triage strategies, particularly with respect to health disparities arising from unbalanced ancestry representation.

Keywords: Precision medicine, Inborn errors of metabolism, Pharmacogenomics, Ethics, Genomics, Learning healthcare, Machine learning, Computational biology

¹Biomedical Informatics Training Program, Stanford University, Stanford, CA 94305, USA; ²Biophysics Graduate Group, University of California, Berkeley, CA 94720, USA; ³Department of Bioengineering and Therapeutics, University of California, San Francisco, CA 94143, USA; ⁴Program in Bioethics, University of California San Francisco, San Francisco, CA 94143, USA; ⁵Institute for Health & Aging, University of California San Francisco, San Francisco, CA 94143, USA; ⁷Department of Bioengineering, Stanford University, Stanford, CA 94305, USA; ⁸Department of Plant and Microbial Biology, University of California Berkeley, Berkeley, CA 94720, USA; ⁹Center for Computational Biology, University of California Berkeley, Berkeley, CA 94720, USA; ¹⁰UCSF-UCB Joint Medical Program, School of Public Health, University of California, Berkeley, CA 94720, USA; ¹¹Department of Social & Behavioral Sciences, University of California San Francisco, San Francisco, CA 94143, USA; ¹²Department of Humanities & Social Sciences, University of California San Francisco, San Francisco, CA 94143, USA; ¹³Stanford Center for Biomedical Ethics, Stanford University School of Medicine, Stanford, CA 94305, USA; ¹⁴Department of Pediatrics, University of California San Francisco, San Francisco, CA, USA; ¹⁵Departments of Bioengineering & Genetics, Stanford University, Stanford, CA 94305, USA

¹⁶These authors contributed equally to this work

*Corresponding author: russ.altman@stanford.edu

1. Introduction

We are approaching an era in which genome sequencing at birth will be a widespread practice with the potential to revolutionize healthcare. Interpretation of the genetic variants identified by sequencing, however, is a significant challenge and limits the use of DNA sequencing as a primary diagnostic screen¹. Current algorithms used to interpret the significance of genetic mutations are not reliable enough to be used without additional clinical data². Yet, accumulating biomedical data enables machine learning algorithms to predict the consequence of genetic variants with increasing accuracy. The pairing of modern algorithms and widespread genome sequencing is beginning to deliver precision medicine in limited settings³, but the broad interpretation of rare genetic variation requires algorithmic advances and improved access to data. The identification of rare variation responsible for unusual clinical phenotypes is a particularly difficult challenge because both the responsible gene and the associated variation must be identified. A slightly more tractable problem is the identification of clinically important variants in genes that are already known to be clinically significant and have known mechanisms for influencing phenotype.

This paper focuses on two fields that have known clinically important genes and in the near term should benefit greatly from improved rare variant interpretation: pharmacogenomics (PGx) and inborn errors of metabolism (IEM). IEM and PGx are subfields of genetics characterized by monogenic phenotypes for which therapeutic action can be taken in response to clinically-important variants in known genes. Both fields have been revolutionized by low-cost sequencing and the curation of large databases cataloguing the effects of specific genetic variants. Furthermore, both fields struggle with interpretation of the phenotypic effects of rare variants that have not been clinically evaluated.

As an interdisciplinary team supported by the Chan Zuckerberg Biohub, we approach these two challenges by addressing both computational and ethical issues, in order to develop a framework

for genome-informed medical care that benefits all. Here we review the current practices and limitations of variant interpretation in PGx and IEM and highlight recent computational advances that will allow researchers to improve precision medicine. Ethical considerations in these activities primarily address health disparities, since existing genetic and genomic databases are not inclusive of individuals of diverse ancestries. As the recent strategic vision from the US National Human Genomic Research Institute (NHGRI) attests, there are significant societal implications of a genomic learning healthcare system that we cannot afford to oversimplify⁴. Our focus on genes of known consequence should generalize ultimately to the more difficult cases where the gene, function and mechanism are not well-understood.

2. PGx and IEM in current clinical practice

For both PGx and IEM, our detailed understanding of the biological processes at play (the genes that are critical and how they interact) has reached a point at which routine genetic screens can inform clinical decision making. In the United States, PGx testing is mandated by the Food and Drug Administration for a number of drugs due to safety concerns, and recommended for many others. Testing for IEMs is routine practice for nearly all newborns in the United States, but the role of genetic testing is largely limited to second-tier screens and carrier testing. These two fields are linked in more ways than it may superficially appear. The clinical implications for most known PGx and IEM driven phenotypes are often caused by variants in a single gene. As monogenic traits, there is a critical importance in understanding the impact of variants in the underlying genes, but also a narrowing of the problem space that makes for a tractable solution. Additionally, the mechanisms of disease and treatment response are generally understood.

PGx describes how an individual's response to medication is influenced by genetic variation in pharmacogenes: genes encoding proteins involved in the pharmacokinetics and

pharmacodynamics of a drug⁵. Many pharmacogenes have common genetic variants with known clinical significance. These variants can affect the metabolism, transport and action of drugs throughout the body and may influence efficacy or lead to adverse events. Studies have shown as many as 99.8% of individuals carry at least one genetic variant that could lead to adverse outcomes for at least one drug^{6–8}. In the past, clinical practice overlooked the influence of genetics on drug response, and—except for several extreme cases⁹—used a standardized dose of any particular drug for most patients, with some trial-and-adjustment to determine the ideal drug and dosage. This error-prone process can lead to decreased efficacy and increased incidence of adverse events that could be otherwise avoided¹⁰. Clinical practice may be moving towards genetic testing prior to drug dosing, though at present current practice is still limited to physician-guided treatment, with genotyping or sequencing ordered by a physician and carried out clinically (Fig 1A). To date there are 60 drugs with clinical dosing guidelines published by the Clinical Pharmacogenomics Implementation Consortium (CPIC) and 94 drugs with guidelines from the Dutch Pharmacogenomics Working Group (DPWG)¹¹. As the inexpensive interrogation of genetic information gains a foothold in clinical medicine, pharmacogenetic information will increasingly be used in standard care. Importantly, when genetic information is used to guide dosing, the current focus is on common polymorphisms in individuals of European ancestry. Common polymorphisms in other ancestral groups and rare variants are generally not included in current clinical dosing guidelines. This can lead to health disparities based on ancestry, and is problematic for all individuals since rare variants are estimated to contribute to as much as 50% of interindividual variation in drug response¹².

IEM encompass more than 1,000 genetic disorders, including organic acidemias, urea cycle defects, lysosomal storage disorders, and disorders of amino acid metabolism¹³. IEM are characterized by monogenic mutations that can affect protein function and result in altered metabolite levels. The majority are autosomal recessive disorders. Many IEM are severe, early-

onset conditions amenable to therapeutic intervention, with early treatment leading to significantly improved clinical outcomes. Since the consequences of unrecognized IEM in pre-symptomatic newborns can be catastrophic, detection before symptom manifestation is essential. Newborn screening (NBS), a pervasive public health effort, detects over 40 of the most common, treatable IEM using biochemical tests performed in blood samples taken shortly after birth.

Population-level NBS has been a routine part of care in the United States since the 1960s. Presently, NBS detects IEM by identifying elevated metabolites in blood, which is performed using tandem mass spectrometry (MS/MS), an inexpensive and rapid test. However, disorders may be missed, some analytes are non-specific, and follow-up testing may be time-consuming and complex^{1,14}. DNA sequencing has the potential to more accurately identify disorders for which MS/MS detection is not optimal, and also identify disorders for which there is no appropriate metabolite screen.

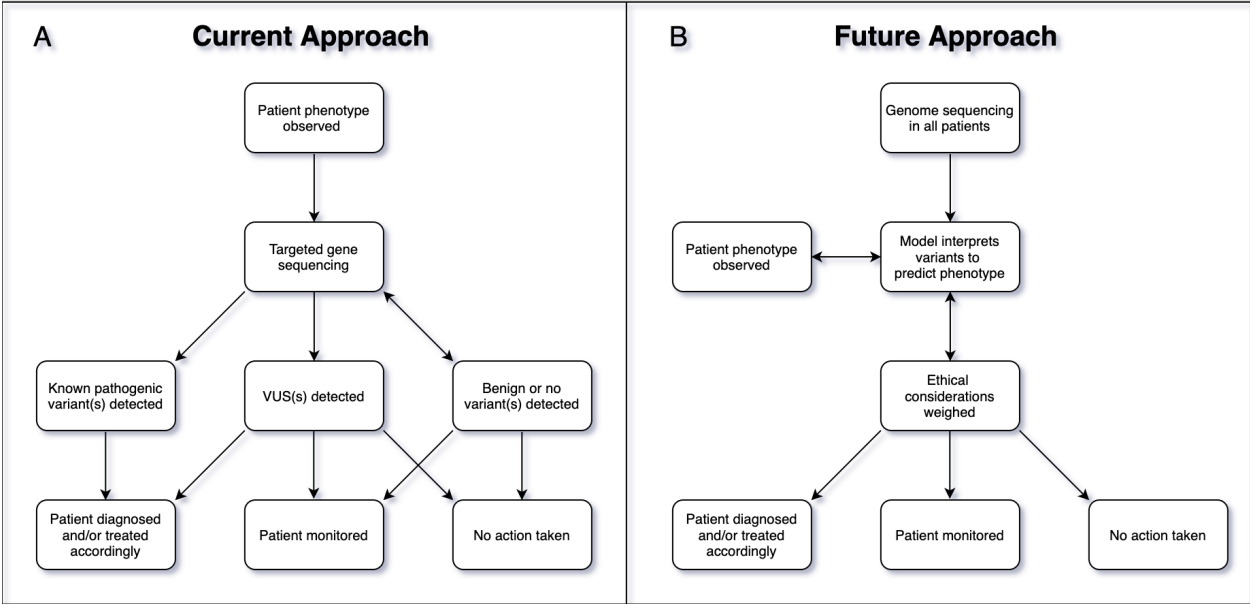


Figure 1. Simplified overview of the current approach to identification and treatment of patients with IEM or PGx related disorders contrasted with our proposed framework, which incorporates early sequencing and analysis of rare

variants with machine learning and ethical considerations. (A) Current practice for IEM and PGx begins with an observable phenotype. In IEM, this may be an altered metabolite detected by newborn screening; in PGx, perhaps an adverse event. Phenotype can also include physical examination, medical history, family history, and relevant labs or studies. Genetic sequencing is then performed, which could involve targeted single gene sequencing with copy number variant detection, gene panel, whole exome sequencing, or in some cases trio sequencing of the affected individual and parents to assess phasing and examine the possibility of de novo variation. If annotated pathogenic variants are identified in the target gene, a patient is diagnosed with an IEM or given a different drug/dose. Identification of VUS may result in either a diagnosis or lack thereof, depending on the other variants identified. In this simplified figure, VUS refers to a variant in the targeted gene of interest as opposed to an incidental finding not relevant to diagnosis. Patient diagnosis can occur without DNA sequencing, as is the case with some IEM. (B) Hypothetical future approach to patient care in the fields of PGx and IEM. All individuals undergo whole genome sequencing at birth. Machine learning models use detected variants to predict phenotype (disease risk or differential drug response). Ethics are considered and clinical action is taken accordingly.

Carrier testing provides an opportunity to detect rare variants in IEM and other disease-associated genes¹⁵ before conception. However, interpretation of genetic screening results still faces significant challenges¹⁶, especially in cases identifying variants of uncertain significance (VUS) where risk for inherited disease cannot be definitively assessed and actionability is questionable. The falling cost of next generation sequencing will continue to expand the identification of genomic variants that may cause IEM or alter drug response. While many genetic variants have catalogued associations with disease phenotypes or drug response, the majority are of unknown clinical consequence. Generating experimental data to validate the pathogenicity of individual variants is tedious and expensive, though recent advances have facilitated more large-scale generation of data¹⁷. Several databases attempt to catalog variants in disease-causing genes, but there is no central catalog for associated functional data. Thus, alternative methods for determining or predicting functional effects of genetic variants are urgently needed.

3. Ethical considerations in rare variant interpretation

Genome-informed precision medicine must include analysis of ethical, legal and social implications (ELSI) in order to improve upon rather than exacerbate existing health disparities⁴. We have identified six chief concerns with enhancing computational predictors for the phenotypic effects of rare variation at the scale proposed here. First, the uncertainty of results and, second, the return of clinical results can either improve or compromise clinical care. Although enhanced computational predictors for IEM and PGx can minimize harm from the trial-and-error of current clinical practices, consistency in clinical education and approaches to ambiguous and incidental findings will be critical to determining societal benefit. Third, there are differences between research and clinical stakeholders in approaching the classifications of VUS that need to be reconciled. Fourth, the underrepresentation of minority groups in current datasets and the underlying research that informs them needs particular attention in order to create a larger and more diverse reference genome so that biases can be reduced. Fifth, an effective genomic learning healthcare system must account for privacy risks. Sixth, there needs to be transparent data sharing expectations across all levels of human input into the learning system. Building on previous ethical frameworks^{18,19} and the need for a nuanced approach²⁰, we suggest that trade-offs between ensuring individual control over data and the social obligations of individuals have yet to be engaged with at the level of ethical governance provisions. Discussion of these concerns is guided by three central ethical questions, summarized in Table 1 and outlined within the [Spotlight Boxes](#).

Table 1: Ethical considerations of the adoption of new technologies into clinical genetics

Area of IEM & PGx and ethical issues	Key question
Whole genome sequencing for newborns 1) Uncertainty of results 2) Return of clinical results, including results from late-onset disorders	<i>Can genome sequencing improve the uncertainty of results & return of clinical results?</i>
Interpreting VUS 3) Research and clinical divide 4) Social/racial inequity	<i>Can we view the classification of VUS as a social justice opportunity, to close social and genetic ancestry gaps?</i>
Genomic learning healthcare systems 5) Privacy risks 6) Data sharing	<i>Can blood spot DNA be viewed as a public good in genomic learning healthcare systems?</i>

Ethics Spotlight 1: Can genome sequencing improve the uncertainty of results & return of clinical results?

For the use of predictive algorithms as the primary methods of analysis for IEM and PGx to be ethically justified, these methods must provide equal or greater certainty than current methods. Improving screening and predictive analysis for IEM and PGx at the testing level is contingent upon the accuracy of results, the provisions around returning results, and the impact on clinical care. Even pathogenic results can have variable penetrance and VUSs, with the possibility of reclassification, can cause significant consternation on both the part of the clinician and patient²¹. Perhaps most thoroughly documented in cancer genetics²², the clinical return of genetic results is rarely straightforward. The prohibition against the return of uncertain results, outlined by the American College of Medical Genetics (ACMG), is such that even if there is a suspicion that an uncertain variant is pathogenic, it should conservatively be classified as a VUS since this information is used in medical decisions².

The follow-up of uncertain results is complicated by clinician/researcher and patient/individual expectations and understandings of actionability. Genomic literacy across different healthcare professional roles is limited^{23,24}. The disclosing of sequencing results should be contingent upon what has been previously explained to the patient/parent about incidental findings and potential treatments²⁵. As healthcare delivery is already biased with regard to decisions about referrals or withdrawals of care, it will be challenging for algorithms to correct for existing biases in the handling of results²⁶. Uncertain and incidental (or secondary) results in clinical care should be considered in the context of existing slippages of fiduciary obligations - such as clinician biases and/or patient mistrust - that emerging tests may or may not be able to compensate for²⁷. The NHGRI has called for greater diversity amongst the genomic scientist workforce⁴.

In order to contain immediate risks around uncertainty of results and focus resources, is there a case for tiered approaches? For example, beginning with targeted sequencing and, upon accuracy improvements, expanding programs to include non-targeted sequencing, or, at the individual level, only sequencing specific genes as a second tier option if a positive test result arises in genome sequencing? Certainly, implementing genome sequencing at the routine screening level requires greater computational accuracy, accessibility and more nuanced ethical safeguards^{4,20}. In the US healthcare context, it is difficult to resolve the issue of healthcare insurance coverage. Can financial disparity in the follow up of results be partially alleviated with temporary coverage through Risk Sharing Agreements between payers and manufacturers of tests?²⁸ Can ethical priorities of the clinician and patient transaction may be corroborated with the needs of the genomic learning healthcare system (that must maximize scarce resources) such that genomic sequencing improves healthcare across all of society?

4. Evaluating variants of unknown significance

Variants in functionally important genes are often suspected to lead to clinical consequences. For IEM and PGx, there are hundreds of genes in which nonsense and missense variants are associated with clinical outcomes. Although additional genetic, epigenetic, and environmental factors alter disease risk and drug response, the gene sequence is the primary determinant of phenotype for these genes. Thousands of pathogenic rare variants in these genes have been characterized with clinical consequences often well understood and catalogued. Yet exome and genome sequencing continue to identify novel variants in these genes at a rapid pace. The ACMG has developed guidelines to interpret these variants, but by design, conclusive evidence is required to assert a variant is pathogenic, even in known disease genes². For example, defects in *PAH* cause phenylketonuria (PKU), an IEM that can lead to severe intellectual disability and seizures when untreated. In gnomAD²⁹, a population database of variants seen in more than 100,000 individuals, 57% of observed protein-altering variants in *PAH* have unknown pathogenicity. Individuals who are homozygous for these variants at birth will have an unknown risk of developing PKU, and carriers of these variants cannot be advised of their risk of having a child with PKU. Thus, predicting the functional consequence of rare variants in IEM and PGx is an important challenge.

To begin to address this issue, numerous publicly available databases actively catalog genetic variants and associated disease and drug response phenotypes. These databases are typically human-curated, and bring together information that would otherwise be dispersed across the literature, allowing researchers and clinicians to quickly access existing knowledge. Several databases focus on the pathogenicity of variants genome-wide, including thousands of variants in IEM and PGx genes. These include ClinVar, ClinGen, HGMD, and Online Mendelian

Inheritance in Man (OMIM)^{30–32}. Such platforms have a shared goal of linking genes with disease, although they take different approaches. ClinVar allows submissions from clinical laboratories, research groups, and specialized databases, presenting all submitted data through an online interface. Most submissions are not manually vetted and are presented as submitted. ClinGen and OMIM attempt to provide authoritative curation of known variants and their relationship to disease. Curators review literature and experimental data to determine pathogenicity of genetic variants. ClinVar and ClinGen share and collaboratively curate data. In addition to being used for standardizing the set of variants with known consequences, these databases are also used by researchers and clinicians to evaluate the evidence that an uncatalogued VUS causes disease based on its similarity to catalogued variants (e.g., if a VUS results in the same amino acid change as a catalogued pathogenic variant, this VUS now has strong evidence for being pathogenic)². Similarly, efforts have been made to catalog the relationship between genetic variation and drug response, exemplified by databases including PharmVar and PharmGKB^{33–35}. Like ClinVar, PharmVar relies on user submissions of discovered haplotypes in genes related to pharmacogenomics.

These variant databases encapsulate the combined expertise of thousands of clinical researchers across the world, but also reveal a large amount of uncertainty. The majority of possible missense variants in IEM and PGx genes are classified as VUS or are altogether missing from databases. ClinVar alone contains more than 6,000 variants classified as VUS in IEM genes, and more than 10,000 VUS in PGx genes (Figure 2a,b). Variants in ClinVar change classification as researchers submit new evidence, but very few VUS variants are resolved as fully pathogenic or benign (Figure 2c,d). Instead, many variants become associated with conflicting classifications. Indeed, 41% of IEM and PGx variants in ClinVar are of uncertain significance or have conflicting interpretations of clinical importance. For novel variants, it is often challenging to establish pathogenic certainty until they are observed by multiple clinicians who submit consistent

classifications to a variant database. For VUS without further clinical or experimental evidence, computational methods offer a possible resolution.

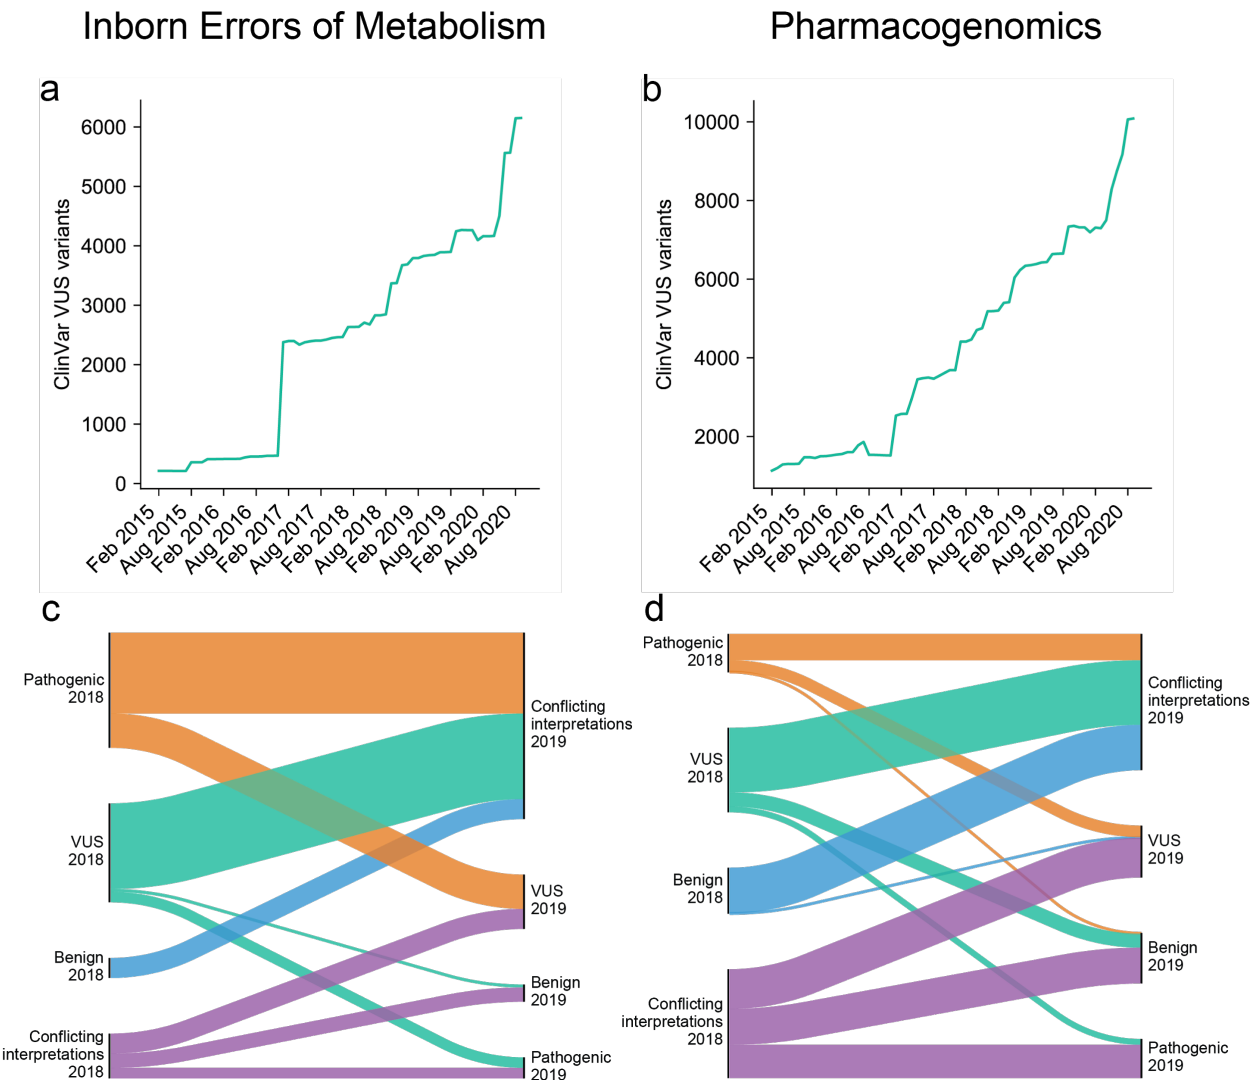


Fig 2. ClinVar variants of unknown significance in genes related to IEM and PGx. **a,b**, the number of VUS in ClinVar between 2015 and 2020 in IEM and PGx genes, respectively. **c,d** All ClinVar variants in

IEM and PGx genes that were reclassified between Feb 2018 and Feb 2019. Height of bars is proportional to number of variants reclassified. A total of 293 variant reclassifications are shown in **c**, and 434 variant reclassifications are shown in **d**.

Most computational approaches predict the functional impact of single nucleotide polymorphisms (SNPs) and small insertions and deletions (INDELs) by training predictive models on variants across the entire genome. The popular tool CADD uses a logistic regression model and more than 60 genomic features to learn the features that distinguish randomly generated variants from recently fixed variants in humans³⁶. The resulting predictor has been used to predict the pathogenicity of clinical variants, and is currently used in clinical analysis pipelines³⁷. REVEL, a meta-predictor, uses the ensemble of scores from several prediction algorithms like CADD, each with different strengths and weaknesses, and is trained to differentiate rare unlabeled variants from HGMD pathogenic variants³⁸. Both CADD and REVEL are capable of predicting the effects of variants in any gene, which is typical of predictors used in clinical research. However, predictors that are gene-, gene family-, or locus-specific generally perform better for both IEM and PGx in comparison to predictors that rely on data from the entire genome^{39–46}. Despite their promise, bespoke methods are constrained by the limited data available for most genes, such as the number of known pathogenic variants and associated functional data. Because these methods are designed to predict the functional impact of a variant, their predictions can be some layers removed from the clinical consequence. Additionally, pharmacogenes are not under the same evolutionary constraint as genes involved in disease, limiting the effectiveness of most predictive algorithms^{47,41}

To combine the best features of variant databases and computational predictors, automated systems that use both in tandem are already being tested to predict the pathogenicity of rare

variants. Consider one recent study evaluating IEM detection by sequencing dried blood spots drawn from newborns¹. This study compared the performance of MS/MS to exome sequencing as a primary screen for IEM on a set of 805 newborns with confirmed IEM. Variants identified by sequencing were automatically assessed on rarity, protein consequence, predicted pathogenicity (including CADD), and matched with catalogued pathogenic variants in ClinVar and HGMD to predict disease status. Overall, this combination was neither sufficiently sensitive nor specific compared to MS/MS, and exome sequencing notably missed a number of cases in which a pair of rare, protein-altering variants were absent from the causal gene. However, performance varied between IEM, and in some cases, provided more specific diagnoses than analyte testing. 32% of pathogenic variants were absent from HGMD and ClinVar. Critically, sequencing led to several false positives in which an individual harbored a pair of rare, protein-altering variants in an IEM gene, but did not have the associated disorder. These false positives significantly limit the ability to use DNA sequencing for screening, and could be mitigated by more accurate computational methods that distinguish pathogenic from benign protein-altering variants.

Ethics spotlight 2: Can we view the classification of VUS as a social justice opportunity?

Whether the classification of VUS and IEMs can offer a fairer distribution of the benefits of sequencing technologies across all population groups is a significant question. Most large datasets in the US contain homogeneous ancestry that is unrepresentative of disadvantaged groups^{48,49}. In addition to the need to improve predictive methods for IEMs, screened individuals need to be considered as part of a social group in relationship to a wider and unequal social system. The moral obligations embedded within the ethics of clinical research and practice need to be better integrated¹⁸. For individuals seeking healthcare, polygenic risk scores are more accurate for patients of European ancestry because the data on which algorithms are trained comes largely from individuals of European ancestry^{50,51}. Similarly, variant impact predictors tend to train catalogued variants from databases, which are not representative of all ancestries. ClinVar was recently found to be missing a large number of hearing impairment variants that primarily affect individuals of African ancestry⁵², likely indicative of a broader pattern. For variant predictors, this bias will lead to greater reliance on European ancestry variants and European genetic context, producing less accurate classification of IEM and PGX variants in other ancestries (e.g., African ancestry), which would only compound existing injustice in healthcare access for under-represented populations^{53,54}. Disparity in ancestry representation is especially stark in data sources for Genome Wide Association Studies (GWAS), where European ancestry disproportionately represents 81% of the data set population⁴⁸.

Can we alleviate healthcare disparity by closing current ancestry gaps in GWAS? Given evidence that polygenic risk scores can be improved upon by incorporating datasets for a broader range of genetic ancestries⁵⁵, it is imperative that GWAS strive for fairer training data also. As the field matures to consider the role of genetic modifiers⁵⁶, as well as social and environmental interactions⁵⁷, results from GWAS of diverse individuals are needed to consider the effects of genetic modifiers and the environment on variants. Newborn screening programs, with their mandatory collection and the near universal application of testing, provide a diverse and truly representative set of individuals⁵⁸. That said, racial discrimination in healthcare and healthcare research is not simply resolvable through technical fixes. Redressing data under-representation and health equity in machine learning precision medicine must be viewed in the context of broader social change, which we pick up on in the next Box regarding questions of social obligation.

5. Opportunities in rare variant evaluation

In predicting the effect of a variant on gene function, we can predict its effects on the system, such as a metabolic pathway, and then on the physiology/pathophysiology. Cataloging observed likely clinically impactful variants in databases such as ClinVar and PharmVar³² can be effective for determining the pathogenicity of more frequent rare variants, (allele frequency between 0.01% and 1%). These variants are common enough that they have been identified in multiple individuals and therefore the effect on phenotype can be verified. However, ultra-rare variants, defined as having an allele frequency less than 0.01%, are responsible for a large portion of rare genetic disorders. Publicly available databases of PKU patients indicate that 60% of cases involve at least one ultra-rare SNV, and in 28% of cases the individual carries an ultra-rare variant on both copies of *PAH*. Some of these ultra-rare variants may be *de novo* mutations, and the individual may be the only person known to harbor that exact variant⁵⁹. The vast majority of ultra-rare variants are absent from clinical databases, indicating that the current approach of cataloguing observed genetic variants fails when allele frequencies are especially low. For *PAH*, which is one of the most studied metabolic genes, only 9% of possible SNVs have functional impact classified in ClinVar.

Emerging computational algorithms may serve as a means for evaluating the impact of rare variants in IEM and PGx genes. As previously stated, existing algorithms have limited ability to accurately predict the impact of variants in these genes, especially among rare variants. Methods have been developed to specifically evaluate variants in pharmacogenes, but these are largely based on existing methods and may have some of the same inherent biases⁴¹. Machine learning has revolutionized computer vision and natural language processing by effectively analyzing spatial and sequential data^{60–62}. With the rapid growth of biological data, deep learning has also

been extensively used in bioinformatics^{63–70}, including transcription factor binding site prediction⁷¹, genome functional annotation⁷², and assessment of variant function^{73,74}. Several methods have been developed specifically for the evaluation of alleles in pharmacogenes, namely *CYP2D6*^{75,76}. These purpose-built models outperform existing methods and are capable of assessing the impact of any combination of variants observed in a haplotype, rather than single variants. One major drawback of deep learning is that it requires an immense amount of data in order to estimate the large number of parameters required for good performance^{77,78}.

Transfer learning offers an opportunity to leverage the power of deep learning in situations where data is limited. It is difficult to obtain sufficient data to develop phenotype prediction algorithms from genomic data using deep learning, especially when we only have 10s or 100s of individuals with both genome sequencing data and well-characterized clinical or molecular phenotypes. Transfer learning is an emerging approach for overcoming the limited data challenges. The idea is to build models that perform a task (X) that is similar to the goal task (Y) but for which there are large amounts of relevant real or simulated data. Once the model for solving task X is performing well, it can be refined with data relevant to task Y. In the case of predicting variants, we might build a model using data from a well-studied gene (X) and then refine the model with data from a poorly-studied gene (Y). The resulting model may perform very well on Y, since the “lessons” learned in modeling X transfer well to Y^{79–84}. There are several flavors of transfer learning that have been applied to applications in genetics and proteomics. Convolutional neural network (CNN) based approaches pre-train weights of convolutional layers on large datasets that can be fine-tuned on smaller datasets⁷⁵. Transformer based approaches, frequently used in natural language processing, have been applied to functional predictions of variants in proteins^{85,86}. Graph-CNNs have been used to make drug binding predictions using protein structure data after being pre-trained using an unsupervised learning step⁸⁷. These transfer learning methods could in theory be used to create structure-based predictions of the effect of amino acid changes on

drug binding. These methods combined with chemoinformatics approaches for representing drug molecules could be used to create substrate-specific predictions of drug-protein interactions and how genetic variants may influence that behavior.

The underlying homology between genes existing in gene families and across species may allow for an increased ability to perform transfer learning. We may be able to use knowledge learned in some domains to inform others. Not surprisingly, some rare diseases have received more attention than others, often due to the frequency of disease, serendipitous factors, and scientific opportunities. These well-studied diseases typically have significantly more variant impact data available than others. PKU has an incidence of 1 in 10,000 newborns, and there are hundreds of disease-associated catalogued variants. In comparison, tyrosine hydroxylase deficiency (THD) affects fewer than 1 in 100,000 newborns and has been associated with fewer than 20 variants. Sequencing benefits individuals with THD less, simply because the disease is rarer and few known pathogenic variants exist. The chemical similarity of phenylalanine and tyrosine leads to a high degree of homology between phenylalanine hydroxylase (PAH) and tyrosine hydroxylase (TH), which presents an opportunity to transfer knowledge about PKU variants to better understand THD—for example, in understanding which parts of the protein may be more or less tolerant of non-synonymous mutations.

Ultimately, the goal of any variant interpretation method is to improve clinical care. Integration of genetics into the clinic is already quite challenging, and integration of computational methods for predicting variant function is rife with further challenges. Learning health systems have long been proposed as models for improving healthcare^{88–90}, but integration of genetic data into such a system would allow for the accumulation of data to train more sophisticated predictive models as well as an opportunity to iteratively improve upon such algorithms.

A genomic learning healthcare system would allow for rapid collection and phenotyping of rare variants. Learning health systems have been proposed in healthcare since 2007, but none have integrated genetics to inform patient treatment⁸⁹. In these implemented systems, the algorithms are constantly improving based on a feedback loop of data that is collected over the course of patient treatment. A genomic learning healthcare system would operate in much the same way, with the addition that clinical decision support is provided based on genetic data as well as clinical data²⁸. In this proposed system, collection, sequencing, and analysis of patient data would be required as a first step, and would need to be available as part of the patient's clinical record in the electronic health system. This would enable clinical decision support for IEM and PGx related conditions, providing doctors with diagnosis and treatment guidance. Then, the algorithms underlying the clinical decision support can be evaluated regularly and updated based on newly available patient data. In addition to evaluating the algorithms, sequencing and analyzing important genes for every individual treated will allow for more rapid collection and phenotyping of ultra-rare variants.

The ultimate goal of a genomic learning healthcare system is to improve treatment for all patients by leveraging their genetic data. This includes triaging rare variants that may be previously unseen in patients and potentially making clinical decisions based on their predicted impact. As a conservative first step, a genomic learning healthcare system could implement existing clinical guidance models for IEM and PGx, such as the pharmacogenomics dosing recommendations from the Clinical Pharmacogenomics Implementation Consortium (which was recently supported by Medicare). Once genetic data are collected for each patient, predictive models for rare variants can be developed and implemented in clinical practice at such a time when there is sufficient confidence in the predictions of the model. Careful analysis will be needed in selecting and evaluating predictive models for both IEM and PGx, and it is likely that gene-specific models will

be needed. We illustrate this framework in Figure 3, before turning to the ethical questions to be taken into account.

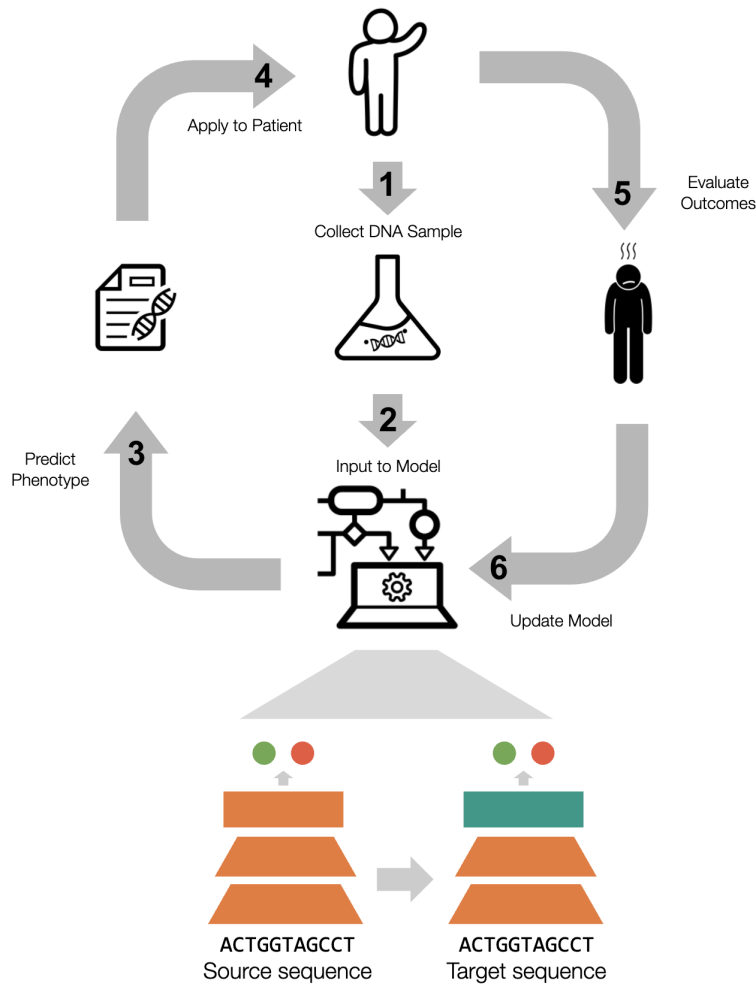


Figure 3. Proposed workflow for a genomic learning healthcare system. Patients DNA samples are collected and sequenced, with genomic data input to computational models. The model outputs a predicted phenotype for the patient with results reviewed by clinicians and applied to the patient. Outcomes are evaluated and the model continues to learn from a feedback loop to improve outcomes for future patients. Icons are from The Noun Project^{91–94}.

Ethics Spotlight 3: Can blood spot DNA be viewed as a public good in a genomic learning healthcare system?

Data sources and determination of data ownership and consent for secondary data use will significantly shape whether or not genomic learning healthcare systems can improve accuracy and reduce biases. Learning health care systems present unique ethical challenges that traditional clinical and research ethics - focusing on individual harms and a sharp research/clinical divide - will find difficult to address¹⁸. Data collection and input (step 1 and step 2 of Fig 3) differs between clinical and public health repositories in terms of provisions around secondary use. The use of AI in healthcare systems is complicated by issues arising from the possible encoding and routinization of human bias, even with the use of seemingly neutral data sources⁴⁹. Further, AI is becoming the repository of the collective medical mind²⁶. More than simply doing no harm, a genomic learning healthcare system should actively support greater health equity^{4,95}. Can clinical data be viewed as a 'public good' insofar as all stakeholders (both healthcare and private industry) hold a moral obligation to use and share clinical data in ways that benefit society over and above individual or commercial interests¹⁹?

If viewing clinical data as a public good, of greatest difficulty would be how to deal with computational predictors and healthcare outcomes that accurately capture differences not so much resulting from human input biases but rather serving unfair social conditions. For public health data use, it is important to identify and address social and political inconsistencies in the ethical oversight from Institutional Review Boards and government bodies, particularly in regard to informed consent and anonymization of data⁹⁶. This requires careful consideration for the nuances of beneficence regarding collections and distribution of genomic information⁹⁷. The current justification for the mandatory nature of newborn screening rests on the potential harms to the child were they not screened for these treatable conditions (see Johnston et al. 2018 for a full historical justification)²⁰. Safeguards are needed to protect the storage and research use of genetic data, which could become more identifiable.⁹⁸ With such protections, could the practice of informed consent with individuals be seen as less important than another process to ensure respect for autonomy at a group level, in order to meet social obligations to contribute to both greater knowledge and efforts to reduce social inequity in health¹⁸? Because biobanks of newborn blood-spots provide a rich and unique dataset for research and improving newborn screening (and other genetic testing) -- with enormous potential for contribution to a genomic learning healthcare system -- the loss of such potential (if secondary use of blood spots is only permissible on an individual consent basis) needs to be carefully weighed up against ethical concerns about respect for individual control. How do we ensure respect for individuals in a genomic learning healthcare system that relies on the collective contributions of entire populations in order for everyone to potentially benefit? These are issues that our research must continue to engage with directly.

Conclusion

The defining problem of the genomic age is the interpretation of human variation. In reviewing computational advancements and ethical concerns, we look to develop gene-specific variant interpretation algorithms with a genomic learning healthcare system that builds from a focus on

early-onset treatable disease in newborns and actionable pharmacogenomics recommendations. We seek diagnosis of IEM and treatment for PGx that is tailored to each individual, and treatment outcomes that are shared to improve treatment for future patients across all of society. The existing system is the first step towards this goal, as evidenced by confirmatory sequencing of patients and variant cataloging in databases such as ClinVar. Yet the existing system falls short, because it is reactive rather than predictive, and accurate treatment depends on whether the variant has been previously seen and cataloged. Importantly, it remains to be determined whether computational methods can alleviate health inequity that is reinforced by these limited variant databases. Pervasive sequencing may indeed present a social justice opportunity: to actively promote a more fair and consistent distribution of treatment across all population groups. Yet, there are many barriers in the way, including unrepresentative sequencing databases, secondary data use permissions, barriers to healthcare access, and existing biases at the human interface of research and caregiving.

There are technical challenges, including accurate variant classification, data limitations, and growing numbers of variants of uncertain significance. A combination of integrated learning and transfer learning can overcome existing data limitations in order to improve the computational prediction of variants. An increased understanding of each patient's variants will enable more precise diagnosis and treatment. Most importantly, as more patients provide information into the system, lessons learned from one patient may inform the care of and benefit all patients. A dynamic and fair genomic learning healthcare system will create the greatest patient benefit from the captured genomic and phenotypic information, but this will fundamentally depend on careful consideration of societal implications.

Acknowledgements

G.M. is supported by the Big Data to Knowledge (BD2K) from the National Institutes of Health (T32 LM012409).

A.G.S. is supported by a National Science Foundation Graduate Research Fellowship under Grant No. DGE 1752814.

M.N. and B.K. are supported by U19HD077627.

D.C.M. is supported by Stanford Clinical & Translational Science Award (NIH, NCATS: UL1TR003142)

R.B.A. is supported by NIH GM102365 and HG010615.

Additionally, G.M., M.L.K., J.E.H.B, M.N., S.W., K.M.G., and R.B.A. are funded by the Chan Zuckerberg Biohub.

Competing Interests

Aashish N. Adhikari is currently an employee of Illumina, Inc.

References

1. Adhikari, A.N., Gallagher, R.C., Wang, Y., Currier, R.J., Amatuni, G., Bassaganyas, L., Chen, F., Kundu, K., Kvale, M., Mooney, S.D., et al. (2020). The role of exome sequencing in newborn screening for inborn errors of metabolism. *Nat. Med.*
2. Richards, S., Aziz, N., Bale, S., Bick, D., Das, S., Gastier-Foster, J., Grody, W.W., Hegde, M., Lyon, E., Spector, E., et al. (2015). Standards and guidelines for the interpretation of sequence variants: a joint consensus recommendation of the American College of Medical Genetics and Genomics and the Association for Molecular Pathology. *Genet. Med.* 17, 405–424.
3. Kather, J.N., Pearson, A.T., Halama, N., Jäger, D., Krause, J., Loosen, S.H., Marx, A., Boor, P., Tacke, F., Neumann, U.P., et al. (2019). Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. *Nat. Med.* 25, 1054–1056.
4. Green, E.D., Gunter, C., Biesecker, L.G., Di Francesco, V., Easter, C.L., Feingold, E.A., Felsenfeld, A.L., Kaufman, D.J., Ostrander, E.A., Pavan, W.J., et al. (2020). Strategic vision for improving human health at The Forefront of Genomics. *Nature* 586, 683–692.

5. Lavertu, A., McInnes, G., Daneshjou, R., Whirl-Carrillo, M., Klein, T.E., and Altman, R.B. (2018). Pharmacogenomics and big genomic data: from lab to clinic and back again. *Hum. Mol. Genet.* **27**, R72–R78.
6. Van Driest, S.L., Shi, Y., Bowton, E.A., Schildcrout, J.S., Peterson, J.F., Pulley, J., Denny, J.C., and Roden, D.M. (2014). Clinically actionable genotypes among 10,000 patients with preemptive pharmacogenomic testing. *Clin. Pharmacol. Ther.* **95**, 423–431.
7. Reisberg, S., Krebs, K., Lepamets, M., Kals, M., Mägi, R., Metsalu, K., Lauschke, V.M., Vilo, J., and Milani, L. (2019). Translating genotype data of 44,000 biobank participants into clinical pharmacogenetic recommendations: challenges and solutions. *Genet. Med.* **21**, 1345–1354.
8. McInnes, G., Lavertu, A., Sangkuhl, K., Klein, T.E., Whirl-Carrillo, M., and Altman, R.B. (2020). Pharmacogenetics at scale: An analysis of the UK Biobank.
9. Martin, M.A., Klein, T.E., Dong, B.J., Pirmohamed, M., Haas, D.W., Kroetz, D.L., and Clinical Pharmacogenetics Implementation Consortium (2012). Clinical pharmacogenetics implementation consortium guidelines for HLA-B genotype and abacavir dosing. *Clin. Pharmacol. Ther.* **91**, 734–738.
10. of Medicine, I. (2000). *To Err Is Human: Building a Safer Health System* (Washington, DC: The National Academies Press).
11. Bank, P.C.D., Caudle, K.E., Swen, J.J., Gammal, R.S., Whirl-Carrillo, M., Klein, T.E., Relling, M.V., and Guchelaar, H.-J. (2018). Comparison of the Guidelines of the Clinical Pharmacogenetics Implementation Consortium and the Dutch Pharmacogenetics Working Group. *Clin. Pharmacol. Ther.* **103**, 599–618.
12. Ingelman-Sundberg, M., Mkrtchian, S., Zhou, Y., and Lauschke, V.M. (2018). Integrating rare genetic variants into pharmacogenetic drug response predictions. *Hum. Genomics* **12**, 26.
13. Ferreira, C.R., van Karnebeek, C.D.M., Vockley, J., and Blau, N. (2019). A proposed nosology of inborn errors of metabolism. *Genet. Med.* **21**, 102–106.
14. Azzopardi, P.J., Upshur, R.E.G., Luca, S., Venkataramanan, V., Potter, B.K., Chakraborty, P.K., and Hayeems, R.Z. (2020). Health-care providers' perspectives on uncertainty generated by variant forms of newborn screening targets. *Genet. Med.* **22**, 566–573.
15. Azimi, M., Schmaus, K., Greger, V., Neitzel, D., Rochelle, R., and Dinh, T. (2016). Carrier screening by next-generation sequencing: health benefits and cost effectiveness. *Mol Genet Genomic Med* **4**, 292–302.
16. Kraft, S.A., Duenas, D., Wilfond, B.S., and Goddard, K.A.B. (2019). The evolving landscape of expanded carrier screening: challenges and opportunities. *Genet. Med.* **21**, 790–797.
17. Fowler, D.M., and Fields, S. (2014). Deep mutational scanning: a new style of protein science. *Nat. Methods* **11**, 801–807.
18. Faden, R.R., Kass, N.E., Goodman, S.N., Pronovost, P., Tunis, S., and Beauchamp, T.L. (2013). An ethics framework for a learning health care system: a departure from traditional research ethics and clinical ethics. *Hastings Cent. Rep. Spec No*, S16–S27.
19. Larson, D.B., Magnus, D.C., Lungren, M.P., Shah, N.H., and Langlotz, C.P. (2020). Ethics of

Using and Sharing Clinical Imaging Data for Artificial Intelligence: A Proposed Framework. *Radiology* 295, 675–682.

20. Johnston, J., Lantos, J.D., Goldenberg, A., Chen, F., Parens, E., Koenig, B.A., and members of the NSIGHT Ethics and Policy Advisory Board (2018). Sequencing Newborns: A Call for Nuanced Use of Genomic Technologies. *Hastings Cent. Rep.* 48 Suppl 2, S2–S6.
21. Couzin-Frankel, J. (2014). Unknown significance. *Science* 346, 1167–1170.
22. Vineis, P. (1997). Ethical issues in genetic screening for cancer. *Ann. Oncol.* 8, 945–949.
23. Hippman, C., and Nislow, C. (2019). Pharmacogenomic Testing: Clinical Evidence and Implementation Challenges. *J Pers Med* 9,.
24. Popejoy, A.B., Crooks, K.R., Fullerton, S.M., Hindorff, L.A., Hooker, G.W., Koenig, B.A., Pino, N., Ramos, E.M., Ritter, D.I., Wand, H., et al. (2020). Clinical Genetics Lacks Standard Definitions and Protocols for the Collection and Use of Diversity Measures. *Am. J. Hum. Genet.* 107, 72–82.
25. McCullough, L.B., Brothers, K.B., Chung, W.K., Joffe, S., Koenig, B.A., Wilfond, B., Yu, J.-H., and Clinical Sequencing Exploratory Research (CSER) Consortium Pediatrics Working Group (2015). Professionally Responsible Disclosure of Genomic Sequencing Results in Pediatric Practice. *Pediatrics* 136, e974–e982.
26. Char, D.S., Shah, N.H., and Magnus, D. (2018). Implementing Machine Learning in Health Care - Addressing Ethical Challenges. *N. Engl. J. Med.* 378, 981–983.
27. Martinez-Martin, N., and Magnus, D. (2019). Privacy and ethical challenges in next-generation sequencing. *Expert Review of Precision Medicine and Drug Development* 4, 95–104.
28. Lu, C.Y., Williams, M.S., Ginsburg, G.S., Toh, S., Brown, J.S., and Khoury, M.J. (2018). A proposed approach to accelerate evidence generation for genomic-based technologies in the context of a learning health system. *Genet. Med.* 20, 390–396.
29. Karczewski, K.J., Francioli, L.C., Tiao, G., Cummings, B.B., Alföldi, J., Wang, Q., Collins, R.L., Laricchia, K.M., Ganna, A., Birnbaum, D.P., et al. (2020). The mutational constraint spectrum quantified from variation in 141,456 humans. *Nature* 581, 434–443.
30. Hamosh, A., Scott, A.F., Amberger, J.S., Bocchini, C.A., and McKusick, V.A. (2005). Online Mendelian Inheritance in Man (OMIM), a knowledgebase of human genes and genetic disorders. *Nucleic Acids Res.* 33, D514–D517.
31. Rehm, H.L., Berg, J.S., Brooks, L.D., Bustamante, C.D., Evans, J.P., Landrum, M.J., Ledbetter, D.H., Maglott, D.R., Martin, C.L., Nussbaum, R.L., et al. (2015). ClinGen—the clinical genome resource. *N. Engl. J. Med.* 372, 2235–2242.
32. Landrum, M.J., Lee, J.M., Benson, M., Brown, G.R., Chao, C., Chitipiralla, S., Gu, B., Hart, J., Hoffman, D., Jang, W., et al. (2018). ClinVar: improving access to variant interpretations and supporting evidence. *Nucleic Acids Res.* 46, D1062–D1067.
33. Gaedigk, A., Ingelman-Sundberg, M., Miller, N.A., Leeder, J.S., Whirl-Carrillo, M., Klein, T.E., and PharmVar Steering Committee (2018). The Pharmacogene Variation (PharmVar) Consortium: Incorporation of the Human Cytochrome P450 (CYP) Allele Nomenclature

Database. *Clin. Pharmacol. Ther.* 103, 399–401.

34. Gaedigk, A., Sangkuhl, K., Whirl-Carrillo, M., Twist, G.P., Klein, T.E., Miller, N.A., and PharmVar Steering Committee (2019). The Evolution of PharmVar. *Clin. Pharmacol. Ther.* 105, 29–32.

35. Whirl-Carrillo, M., McDonagh, E.M., Hebert, J.M., Gong, L., Sangkuhl, K., Thorn, C.F., Altman, R.B., and Klein, T.E. (2012). Pharmacogenomics knowledge for personalized medicine. *Clin. Pharmacol. Ther.* 92, 414–417.

36. Rentzsch, P., Witten, D., Cooper, G.M., Shendure, J., and Kircher, M. (2019). CADD: predicting the deleteriousness of variants throughout the human genome. *Nucleic Acids Res.* 47, D886–D894.

37. van der Velde, K.J., Kuiper, J., Thompson, B.A., Plazzer, J.-P., van Valkenhoef, G., de Haan, M., Jongbloed, J.D.H., Wijmenga, C., de Koning, T.J., Abbott, K.M., et al. (2015). Evaluation of CADD Scores in Curated Mismatch Repair Gene Variants Yields a Model for Clinical Validation and Prioritization. *Hum. Mutat.* 36, 712–719.

38. Ioannidis, N.M., Rothstein, J.H., Pejaver, V., Middha, S., McDonnell, S.K., Baheti, S., Musolf, A., Li, Q., Holzinger, E., Karyadi, D., et al. (2016). REVEL: An Ensemble Method for Predicting the Pathogenicity of Rare Missense Variants. *Am. J. Hum. Genet.* 99, 877–885.

39. Li, Q., Liu, X., Gibbs, R.A., Boerwinkle, E., Polychronakos, C., and Qu, H.-Q. (2014). Gene-specific function prediction for non-synonymous mutations in monogenic diabetes genes. *PLoS One* 9, e104452.

40. Hamasaki-Katagiri, N., Salari, R., Wu, A., Qi, Y., Schiller, T., Filiberto, A.C., Schisterman, E.F., Komar, A.A., Przytycka, T.M., and Kimchi-Sarfaty, C. (2013). A gene-specific method for predicting hemophilia-causing point mutations. *J. Mol. Biol.* 425, 4023–4033.

41. Zhou, Y., Mkrtchian, S., Kumondai, M., Hiratsuka, M., and Lauschke, V.M. (2019). An optimized prediction framework to assess the functional impact of pharmacogenetic variants. *Pharmacogenomics J.* 19, 115–126.

42. Adhikari, A.N. (2019). Gene-specific features enhance interpretation of mutational impact on acid α -glucosidase enzyme activity. *Human Mutation* 40, 1507–1518.

43. Lal, D., May, P., Perez-Palma, E., Samocha, K.E., Kosmicki, J.A., Robinson, E.B., Møller, R.S., Krause, R., Nürnberg, P., Weckhuysen, S., et al. (2020). Gene family information facilitates variant interpretation and identification of disease-associated genes in neurodevelopmental disorders. *Genome Med.* 12, 28.

44. Heyne, H.O., Baez-Nieto, D., Iqbal, S., Palmer, D.S., Brunklaus, A., May, P., Collaborative, E., Johannesen, K.M., Lauxmann, S., Lemke, J.R., et al. (2020). Predicting functional effects of missense variants in voltage-gated sodium and calcium channels. *Sci. Transl. Med.* 12,.

45. Clerx, M., Heijman, J., Collins, P., and Volders, P.G.A. (2018). Predicting changes to INa from missense mutations in human SCN5A. *Sci. Rep.* 8, 12797.

46. Li, B., Mendenhall, J.L., Kroncke, B.M., Taylor, K.C., Huang, H., Smith, D.K., Vanoye, C.G., Blume, J.D., George, A.L., Jr, Sanders, C.R., et al. (2017). Predicting the Functional Impact of

KCNQ1 Variants of Unknown Significance. *Circ. Cardiovasc. Genet.* *10*,.

47. Jorge, L.F., Eichelbaum, M., Griese, E.U., Inaba, T., and Arias, T.D. (1999). Comparative evolutionary pharmacogenetics of CYP2D6 in Ngawbe and Embera Amerindians of Panama and Colombia: role of selection versus drift in world populations. *Pharmacogenetics* *9*, 217–228.
48. Popejoy, A.B., and Fullerton, S.M. (2016). Genomics is failing on diversity. *Nature* *538*, 161–164.
49. Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science* *366*, 447–453.
50. Duncan, L., Shen, H., Gelaye, B., Meijssen, J., Ressler, K., Feldman, M., Peterson, R., and Domingue, B. (2019). Analysis of polygenic risk score usage and performance in diverse human populations. *Nat. Commun.* *10*, 3328.
51. Martin, A.R., Gignoux, C.R., Walters, R.K., Wojcik, G.L., Neale, B.M., Gravel, S., Daly, M.J., Bustamante, C.D., and Kenny, E.E. (2017). Human Demographic History Impacts Genetic Risk Prediction across Diverse Populations. *Am. J. Hum. Genet.* *100*, 635–649.
52. Chakchouk, I., Zhang, D., Zhang, Z., Francioli, L.C., Santos-Cortez, R.L.P., Schrauwen, I., and Leal, S.M. (2019). Disparities in discovery of pathogenic variants for autosomal recessive non-syndromic hearing impairment by ancestry. *Eur. J. Hum. Genet.* *27*, 1456–1465.
53. Perera, M.A., Cavallari, L.H., and Johnson, J.A. (2014). Warfarin pharmacogenetics: an illustration of the importance of studies in minority populations. *Clin. Pharmacol. Ther.* *95*, 242–244.
54. Amendola, C.P., Silva-Jr, J.M., Carvalho, T., Sanches, L.C., Silva, U.V. de A.E., Almeida, R., Burdmann, E., Lima, E., Barbosa, F.F., Ferreira, R.S., et al. (2018). Goal-directed therapy in patients with early acute kidney injury: a multicenter randomized controlled trial. *Clinics* *73*, e327.
55. Martin, A.R., Kanai, M., Kamatani, Y., Okada, Y., Neale, B.M., and Daly, M.J. (2019). Clinical use of current polygenic risk scores may exacerbate health disparities. *Nat. Genet.* *51*, 584–591.
56. Rahit, K.M.T.H., and Tarailo-Graovac, M. (2020). Genetic Modifiers and Rare Mendelian Disease. *Genes* *11*,.
57. Hunter, D.J. (2005). Gene-environment interactions in human diseases. *Nat. Rev. Genet.* *6*, 287–298.
58. Feuchtbaum, L., Carter, J., Dowray, S., Currier, R.J., and Lorey, F. (2012). Birth prevalence of disorders detectable through newborn screening by race/ethnicity. *Genet. Med.* *14*, 937–945.
59. Blau, N., Shen, N., and Carducci, C. (2014). Molecular genetics and diagnosis of phenylketonuria: state of the art. *Expert Rev. Mol. Diagn.* *14*, 655–671.
60. Hinton, G., Deng, L., Yu, D., Dahl, G., Mohamed, A.-R., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Kingsbury, B., et al. (2012). Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal Process. Mag.* *29*,.

61. Kiros, R., Zhu, Y., Salakhutdinov, R.R., Zemel, R., Urtasun, R., Torralba, A., and Fidler, S. (2015). Skip-Thought Vectors. In *Advances in Neural Information Processing Systems 28*, C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, and R. Garnett, eds. (Curran Associates, Inc.), pp. 3294–3302.
62. Collobert, R., and Weston, J. (2008). A Unified Architecture for Natural Language Processing: Deep Neural Networks with Multitask Learning. In *Proceedings of the 25th International Conference on Machine Learning*, (New York, NY, USA: ACM), pp. 160–167.
63. Ching, T., Himmelstein, D.S., Beaulieu-Jones, B.K., Kalinin, A.A., Do, B.T., Way, G.P., Ferrero, E., Agapow, P.-M., Zietz, M., Hoffman, M.M., et al. (2018). Opportunities and obstacles for deep learning in biology and medicine. *J. R. Soc. Interface* 15,.
64. Min, S., Lee, B., and Yoon, S. (2017). Deep learning in bioinformatics. *Brief. Bioinform.* 18, 851–869.
65. Zou, J., Huss, M., Abid, A., Mohammadi, P., Torkamani, A., and Telenti, A. (2019). A primer on deep learning in genomics. *Nat. Genet.* 51, 12–18.
66. Libbrecht, M.W., and Noble, W.S. (2015). Machine learning applications in genetics and genomics. *Nat. Rev. Genet.* 16, 321–332.
67. Yue, T., and Wang, H. (2018). Deep Learning for Genomics: A Concise Overview.
68. Angermueller, C., Pärnamaa, T., Parts, L., and Stegle, O. (2016). Deep learning for computational biology. *Mol. Syst. Biol.* 12, 878.
69. Ma, J., Yu, M.K., Fong, S., Ono, K., Sage, E., Demchak, B., Sharan, R., and Ideker, T. (2018). Using deep learning to model the hierarchical structure and function of a cell. *Nat. Methods* 15, 290–298.
70. Kelley, D.R., Snoek, J., and Rinn, J.L. (2016). Basset: learning the regulatory code of the accessible genome with deep convolutional neural networks. *Genome Res.* 26, 990–999.
71. Shen, Z., Bao, W., and Huang, D.-S. (2018). Recurrent Neural Network for Predicting Transcription Factor Binding Sites. *Sci. Rep.* 8, 15270.
72. Khodabandelou, G., Mozziconacci, J., and Routhier, E. (2018). Genome functional annotation using deep convolutional neural network. *bioRxiv* 330308.
73. Quang, D., Chen, Y., and Xie, X. (2015). DANN: a deep learning approach for annotating the pathogenicity of genetic variants. *Bioinformatics* 31, 761–763.
74. Zhou, J., and Troyanskaya, O.G. (2015). Predicting effects of noncoding variants with deep learning-based sequence model. *Nat. Methods* 12, 931–934.
75. McInnes, G., Dalton, R., Sangkuhl, K., Whirl-Carrillo, M., Lee, S.-B., Tsao, P.S., Gaedigk, A., Altman, R.B., and Woodahl, E.L. (2020). Transfer learning enables prediction of CYP2D6 haplotype function.
76. van der Lee, M., Allard, W.G., Rolf H A, Baak-Pablo, R.F., Menafrá, R., Birgit A L, Deenen, M.J., Neven, P., Johansson, I., Gastaldello, S., et al. (2020). A unifying model to predict variable drug response for personalised medicine.

77. Erhan, D., Bengio, Y., Courville, A., Manzagol, P.-A., Vincent, P., and Bengio, S. (2010). Why Does Unsupervised Pre-training Help Deep Learning? *J. Mach. Learn. Res.* *11*, 625–660.
78. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* *15*, 1929–1958.
79. Pan, S.J., and Yang, Q. (2010). A Survey on Transfer Learning. *IEEE Trans. Knowl. Data Eng.* *22*, 1345–1359.
80. Shao, L., Zhu, F., and Li, X. (2015). Transfer learning for visual categorization: a survey. *IEEE Trans Neural Netw Learn Syst* *26*, 1019–1034.
81. Weiss, K., Khoshgoftaar, T.M., and Wang, D. (2016). A survey of transfer learning. *Journal of Big Data* *3*, 9.
82. Zamir, A.R., Sax, A., Shen, W., Guibas, L.J., Malik, J., and Savarese, S. (2018). Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3712–3722.
83. Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. (2014). How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, eds. (Curran Associates, Inc.), pp. 3320–3328.
84. Taroni, J.N., Grayson, P.C., Hu, Q., Eddy, S., Kretzler, M., Merkel, P.A., and Greene, C.S. (2019). MultiPLIER: A Transfer Learning Framework for Transcriptomics Reveals Systemic Features of Rare Disease. *Cell Syst* *8*, 380–394.e4.
85. Rao, R., Bhattacharya, N., Thomas, N., Duan, Y., Chen, P., Canny, J., Abbeel, P., and Song, Y. (2019). Evaluating Protein Transfer Learning with TAPE. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett, eds. (Curran Associates, Inc.), pp. 9689–9701.
86. Elnaggar, A., Heinzinger, M., Dallago, C., Rehawi, G., Wang, Y., Jones, L., Gibbs, T., Feher, T., Angerer, C., Steinegger, M., et al. (2020). ProtTrans: Towards Cracking the Language of Life's Code Through Self-Supervised Deep Learning and High Performance Computing.
87. Torng, W., and Altman, R.B. (2019). Graph Convolutional Neural Networks for Predicting Drug-Target Interactions. *J. Chem. Inf. Model.* *59*, 4131–4149.
88. Guise, J.-M., Savitz, L.A., and Friedman, C.P. (2018). Mind the Gap: Putting Evidence into Practice in the Era of Learning Health Systems. *J. Gen. Intern. Med.* *33*, 2237–2239.
89. Etheredge, L.M. (2007). A rapid-learning health system: what would a rapid-learning health system look like, and how might we get there? *Health Aff.*
90. Greene, S.M., Reid, R.J., and Larson, E.B. (2012). Implementing the learning health system: from concept to action. *Ann. Intern. Med.* *157*, 207–210.
91. Lamm, V. People Drawn Thin Collection, Person.

92. Alberto Gongora, H. Big Data Collection, Machine Learning.
93. Wray, A. Gene Testing.
94. ProSymbols, U.S. STEM Elements Line Icons Collections, DNA.
95. Rajkomar, A., Hardt, M., Howell, M.D., Corrado, G., and Chin, M.H. (2018). Ensuring Fairness in Machine Learning to Advance Health Equity. *Ann. Intern. Med.* 169, 866–872.
96. O'Doherty, K.C., Christofides, E., Yen, J., Bentzen, H.B., Burke, W., Hallowell, N., Koenig, B.A., and Willison, D.J. (2016). If you build it, they will come: unintended future uses of organised health data collections. *BMC Med. Ethics* 17, 54.
97. Green, A.R., Tan-McGrory, A., Cervantes, M.C., and Betancourt, J.R. (2010). Leveraging quality improvement to achieve equity in health care. *Jt. Comm. J. Qual. Patient Saf.* 36, 435–442.
98. Shringarpure, S.S., and Bustamante, C.D. (2015). Privacy Risks from Genomic Data-Sharing Beacons. *Am. J. Hum. Genet.* 97, 631–646.