



Published in final edited form as:

Trends Genet. 2018 July ; 34(7): 504–517. doi:10.1016/j.tig.2018.03.005.

Replicability and Prediction: lessons and challenges from GWAS

Urko M. Marigorta^{1,&}, Juan Antonio Rodríguez^{2,3,4,&,#}, Greg Gibson¹, and Arcadi Navarro^{2,5,6,7,*}

¹Center for Integrative Genomics, Georgia Institute of Technology, Atlanta, Georgia, USA

²CNAG-CRG, Centre for Genomic Regulation (CRG), Barcelona Institute of Science and Technology (BIST), Barcelona, Catalonia, Spain

³Gene Regulation, Stem Cells and Cancer Program. Centre for Genomic Regulation (CRG), Barcelona, Catalonia, Spain

⁴Universitat Pompeu Fabra (UPF), Barcelona, Catalonia, Spain

⁵Institute of Evolutionary Biology (UPF-CSIC), Department of Experimental and Health Sciences, Universitat Pompeu Fabra, Barcelona, Catalonia, Spain

⁶National Institute for Bioinformatics (INB), Barcelona, Catalonia, Spain

⁷Institució Catalana de Recerca i Estudis Avançats (ICREA), PRBB, Barcelona, Catalonia, Spain

Abstract

Since the publication of the Wellcome Trust Case Control Consortium (WTCCC) landmark study a decade ago, genome-wide association studies (GWAS) have led to the discovery of thousands of risk variants involved in disease aetiology. This successful story has two angles that are often overlooked. First, GWAS findings are highly replicable. This is an unprecedented phenomenon in complex trait genetics, and indeed in many areas of science, which in past decades had been plagued by false positives. At a time of increasing concerns about the lack of reproducibility, we examine the biological and methodological reasons that account for the replicability of GWAS and identify the challenges ahead. In contrast to the exemplary success at disease gene discovery, at present GWAS findings are not useful to predict phenotypes. We close with an overview of the prospects for individualized prediction of disease risk and its foreseeable impact in clinical practice.

Keywords

GWAS; Replicability; Genetic Architecture; Genetic Risk Score; Prediction

*Correspondence: arcadi.navarro@upf.edu (A. Navarro).

&Equal contribution

#Twitter account: @jrotwitguez

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

The high replicability of GWAS findings is an unprecedented phenomenon

Genome-wide association studies (GWAS, see **Glossary**) are the tool of choice to identify genetic variants associated with complex disease and other phenotypes of interest [1]. They have revolutionized human genetics with the discovery of thousands of alleles that influence disease risk [1]. Success, however, has been accompanied by boom and bust mood swings in the field. The early crisis of the ‘missing heritability’ is a paradigmatic example [2]. The initial hype after the publication of the Wellcome Trust Case Control Consortium (WTCCC) [3] was dampened by a generalized disappointment about the amount of **genetic variance** explained by GWAS results, which was perceived as disappointingly low. In spite of the exponential yield in terms of discovered loci, the actual relevance of GWAS findings often remains under the spotlight [4].

There is however a persistent feature of GWAS results that is often overlooked: they stand out for their high **replicability**. The last decade has been characterized by a plethora of findings that have stood the test of time almost in their entirety. This degree of success is an unprecedented phenomenon in the field of complex trait genetics and, even more generally, in many scientific fields [5], which some argue are plagued by false positive findings that generate a sense of “Replicability Crisis” that may be fuel for anti-science movements [6].

Here we examine the biological and methodological reasons that have led to the success of GWAS in replicating genetic discoveries. We start with an analysis of the replicability of disease-associated variants discovered by GWAS, with special attention to temporal and spatial patterns of replication and the inferences that can be drawn about the **genetic architecture of disease**. We next discuss the reasons behind the outstanding reliability of GWAS findings and identify extant challenges to ensure the replicability of future findings. We end with an analysis of the current state of phenotype prediction based on GWAS findings, which remains as one of the key challenges for a successful implementation precision medicine.

Widespread replicability of GWAS findings

Replication of newly reported results is the most reliable validation of scientific discoveries, as it confirms their true-positive status. In short, ‘there is consensus in science that the final arbiter is replication’ [7]. In complex trait genetics, the most used definition is *exact replication*, whereby the same genetic marker is consistently associated with the same phenotype in independent cohorts [8]. This is usually attempted in GWAS, which usually have built-in replication stages for markers with strongest evidence of association. However, these follow-up attempts do not necessarily adhere to the exact replication criterion and less stringent significance criteria are often used. Oftentimes as well, consistency of the direction of effect can lead to an initially marginal association attaining **genome-wide significance**, which is widely regarded as replication.

Although external replication is always advantageous, clear-cut evidence is often lacking because very diverse definitions of replication are used in the literature. For instance, the statistical evidence warranted to declare replication may range from a soft threshold,

whereby nominal association (e.g. $P < 0.05$) is deemed enough by subsequent studies, to a hard one in which only associations repeatedly achieving genome-wide significance (and hence “re-discovered”) in several studies are considered as formally replicated. To complicate matters, positive association at variants in strong **linkage disequilibrium** with the original marker SNP is often used as evidence of replication, even if formal exact replication is not achieved. In this review, we use the ‘**re-discovery**’ label for whenever a previously discovered SNP-trait association achieves strong statistical significance in a subsequent study, and always explain the criteria used to consider replication when discussing the inferences from studies analysing GWAS replicability.

Due to these hardly reconcilable criteria, the replicability of GWAS is an important dimension that often goes without further scrutiny. A simplistic but straightforward way to assess replication is to use the NHGRI-EBI GWAS Catalog (<https://www.ebi.ac.uk/gwas/>) for evaluation of the degree to which GWAS hits are rediscovered in subsequent studies. In Figure 1 (Key Figure), we present a temporal analysis of the Catalog, a widely used resource based on manual curation of SNP-trait findings from published GWAS. Since the Catalog only includes associations with suggestive statistical evidence of association ($P < 10^{-5}$), any record in the Catalog that reports a SNP-trait pair that had been reported in a previous study can be labelled as a positive replication with strong statistical evidence of the previously known association. The figure includes all findings for 60 diseases with >15 discovered SNPs. To allow for array heterogeneity, we label as replication any instance where any SNP in linkage disequilibrium with the discovery SNP; $r^2 > 0.8$ is associated with the same disease in a subsequent study. The figure illustrates that most of the top hits included in the GWAS Catalog had already been reported in previous publications, and hence correspond to replications of known SNP-trait associations. This was the case even in the first years of GWAS, which were often met with scepticism because the yield of new associations was perceived as unsatisfactory [9].

Perhaps a more underappreciated aspect of that early period of the GWAS era is that, for the first time in complex trait genetics, findings are undoubtedly robust. For instance, all the 24 risk associations discovered at genome-wide significance by the WTCCC in 2007 were replicated within two years. This pattern has not changed over the years, and indeed the bulk of findings since 2012 are replications of associations that had been described at least twice.

Robustness is widespread and cannot be explained only because large effect associations replicate extensively. An inspection of the replicability of six disease traits illustrates this point, which holds both for re-discovery in the GWAS Catalog at $P < 10^{-5}$ as well as for nominal replications at $P < 0.05$ in large meta-analysis published after the discovery study. Specifically, we compared the average re-discovery rates of variants of moderate (Odds Ratio, OR between 1.2 and 1.5) and low **effect size** (OR < 1.2). As shown in Figure 2, re-discovery rates in the GWAS Catalog are similar regardless of effect size, namely an average of 36.5% across the six diseases for variants of moderate risk and 43% for those of weakest effect ($P = 0.93$, Kolmogorov-Smirnov test).

Re-discovery rates of ~40% would seem to imply that more than half of GWAS hits remain unreplicated. However, these percentages are severe underestimates because the GWAS

Catalog only records SNP-trait associations that achieve $P < 10^{-5}$. A more detailed analysis of the same traits focused on recently published meta-analysis renders much higher replication rates. In total, we inquired the status of 339 SNP-trait associations discovered at genome-wide significance ($P < 5 \times 10^{-8}$) by GWAS published prior to the corresponding meta-analysis study. An impressive 94% of results do indeed replicate at $P < 0.05$ (Figure 2 and Table S1) with high correlation between the discovery and replication attempt ORs (Pearson's $r = 0.54$, $P < 10^{-16}$). These magnitudes resemble the ~100% replication rates observed in the literature after accounting for statistical power and the winner's curse [10, 11]), confirming the reliability of GWAS. The relatively low attention that replication receives together with the confusion in the literature as to what can be formally considered as replication suggest that initiatives such as the GWAS Catalog should formally track replication of their SNP-trait entries.

Cross-ancestry replicability

Although more challenging, comparisons across ethnic groups reinforce the same idea and, furthermore, show that risk variants are shared across ancestries. Early analyses of the **PAGE multi-ethnic case-control cohorts** reported a perfect concordance in risk allele directions at risk SNPs for prostate cancer and type 2 diabetes [12, 13]. This consistency holds for a wide range of diseases [14] and large nominal ($P < 0.05$) replication rates are observed at the level of SNPs [10], loci pinpointed by these SNPs [15], and genome-wide effect sizes [16]. A caveat, however, is that GWAS based on individuals of non-European ancestry usually study fewer individuals [17], rendering our picture of genetic susceptibility to disease across populations incomplete [18]. Moreover, caution is required since incomplete replication can also be informative: there are sound cases of lack of inter-population replicability indicating that some risk variants are population-specific (e.g. NOD2 and Crohn's disease in East Asian populations [19, 20]). Indeed, correlations in effect size across ethnicities are often significantly less than one, and accordingly a fraction of SNPs do not replicate. For example in the comparison of European and Asian associations with major depression, failure to replicate is most commonly because of divergent allele frequencies which reduce power in one population since the proportion of attributable risk declines with minor allele frequency [21].

Replicability is a function of the genetic architecture of disease

Besides confirming the reliability of GWAS as a tool to dissect genetic risk, patterns of temporal and cross-ancestry replicability can illuminate the genetic architecture of complex traits, informing about the reliability of effect estimations and their variability across human ancestries. Age-related macular degeneration (AMD) is a case in point. In one of the first ever GWAS, Klein et al. [22] detected a strong risk variant located in *CFH* (OR=4.6). Given their exceedingly low sample size (96 patients and 50 controls), this report could have been a fluke, or worse, a severe over-estimation of the real effect size. With time, however, this and a few other high-risk variants have been extensively replicated, indicating that a large proportion of AMD heritability is accounted by a handful of variants with extremely large effect sizes [23].

More generally, the power to discover an association is proportional to the amount of phenotypic variance explained by the polymorphism, namely the product of the effect size-squared and the heterozygosity ($VE = 2pq\beta^2$). Here the effect size in standard deviation units is the average effect of substituting one allele for the other. For a trait regulated by thousands of genes, even for common alleles, β will generally be smaller than one twentieth of the standard deviation (for example, the equivalent of less than one mm of height). Very large sample sizes are required to detect such effects and there is usually a winner's curse of over-estimating the effect size, but suitably large replication studies can confirm their influence. Replicability thus also informs about the shared architecture of other diseases, vastly more complex than AMD. For instance, the convergence in associated loci confirms that a large proportion of the genetic risk for immune diseases is shared across conditions [24, 25]. Similarly, the relative lack of findings for mental disorders suggests that psychiatric traits have an infinitesimal architecture controlled by many loci of tiny effects that can only be discovered through large meta-analyses [4, 26].

A corollary is that patterns of replicability at specific associations can afford valuable lessons about disease heterogeneity. For instance, in 2007 three contemporaneous scans for T2D by the WTCCC, DGI and FUSION consortia collaborated in one of the first coordinated data-sharing efforts to enhance evidence about the detected loci [3, 27, 28]. From the incomplete inter-study replication for the signal at the *FTO* locus, with clear heterogeneity in effect, it was learned that *FTO* exerts its effects through BMI, playing only an indirect role in T2D risk [29]. More recently, the development of new methodologies to systematically explore inter-cohort and inter-disease heterogeneity suggests that studying replicability will be key to improve our understanding of the genetic aetiology of disease [30, 31].

Collaborative science and large cohorts were incentives for reproducibility

The reliability of GWAS findings is a novelty that departs sharply from historical experience in the field. For over two decades, the literature in medical genetics has been tormented by promising but later discredited discoveries [32, 33]. This was specially the case for **candidate-gene studies**, in which research groups would perform association analysis focused on genes that had been selected previously based on their biological plausibility. These studies in particular have been characterized by high rates of false positives [34]. For instance, after thousands of candidate-gene studies published during the 90s and early 2000s, literally only a handful of the most studied associations have shown consistent replication [32]. Moreover, the small fraction of findings from the candidate-gene era that have been confirmed by GWAS [35] is composed mainly by risk variants of low effect sizes discovered by the studies with the largest sample sizes [36]. Many factors have contributed to the poor track record of candidate-gene reproducibility, including insufficient availability of genetic markers, inadequate handling of population structure, lack of statistical power due to low sample size, improper control of multiple testing and extensive **publication bias** [32].

The surplus of promising but eventually failed associations seriously undermined the credibility of the whole LD-association-mapping approach, but on the bright side it made researchers aware that they needed to do better. Awareness about the ubiquity of false

positives and about effects being smaller than initially expected shaped the field profoundly (pre-GWAS simulation studies concurrent with the development of the **common disease-common variant** (CD/CV) hypothesis systematically assumed OR's of 1.5, 2, 4 or larger [37]). Such awareness led to (i) a push for increased sample size, (ii) the use of stringent measures to control for multiple testing and avoid false positives (**type I error**), and (iii) a culture of collaboration and data sharing [36].

While it is also true that technological advances in genomics helped to pave the way for GWAS, the generalization of proper multiple testing corrections and setting adequate p-value thresholds to declare significance might have been the most important factor favouring the large replicability of GWAS. The push for statistical rigour was embedded in a new culture of data sharing and demands for external replication before publication (often enforced by funders [38] and journals [39]), which not only translated into more powered studies but also reduced the incidence of publication bias. For instance, the strict threshold to declare genome-wide significance proposed in the original WTCCC study and accompanying feature [40] was a clear guideline against false positives adopted by subsequent studies.

As a final point, the expectation that genetic effects should for the most part be very small has allowed the flagging and correction of results that looked flawed, such as the original GWAS for exceptional longevity [39]. Collaboration and data sharing, once restricted to large-scale scientific infrastructures such as large particle accelerators, has more recently been spearheaded by the genomics community. We propose that similar changes would benefit other scientific areas. This is not to say that we simply favour “big science”, but rather that we believe all scientific endeavours, regardless of their scale, benefit from openness.

What next? Replicability in the full genome era

The cost-effectiveness of SNP array genotyping followed by imputation suggests that GWAS will remain for quite some time as the main instrument to detect loci associated with complex traits [1, 41]. New technological advances, however, may challenge the replicability achieved by the current generation of GWAS. For instance, the decreasing cost of high-throughput sequencing (NGS) allows for exome and whole-genome studies that will however need new thresholds to correct for multiple testing [42]. In addition, the availability of ample clinical history for some big cohorts, such as the UK Biobank, asks for careful control of spurious correlations that may arise from multiple comparisons.

New methodological approaches can also potentially harm the reliability of findings. For instance, the interest in non-additive effects such as epistasis and gene-by-environment interactions demands careful attention to latent covariates, which can introduce biases if unaccounted for [43, 44]. Similarly, the list of new tools, such as multi-marker analyses, which go beyond single SNP tests, or the inclusion of functional evidence to reweight GWAS results, is growing by the day [45, 46]. The heterogeneity in these methodologies will necessarily complicate the evaluation of replicability.

Strict guidelines for publication and the ample experience gathered suggest that GWAS will maintain large replicability rates. However, exponential increases in available genomic and phenotype information may lead to progressive specializations in the hypotheses tested, necessarily departing from a classical GWAS framework. For instance, investigators working on exome data might decide to increase statistical power by aggregate genes into pathways. Under such scenarios, progressive relaxation of the statistical standards may lead to higher rates of false positives.

Indeed, extremely large GWAS finding hundreds of signals are already neglecting the importance of ascertaining the positive replication status of each discovered variant. For instance, the 2014 meta-analysis for height by the GIANT consortium focused on the overall concordance of effects between discovery and independent cohorts [47]. Instead of discussing the replication status of specific findings, this study highlighted the overall amount of phenotypic variance explained [47]. As explained in the next section, this approach is growing in parallel with the interest in using GWAS data for phenotype prediction. However, for variant and gene discovery purposes, insisting on strict significance thresholds and on garnering replication evidence for each individual variant will still be the best recipe for ensuring high replicability rates.

The problem of low predictability

As the biological underpinnings of disease are progressively unveiled, using the findings from GWAS to predict risk of disease at the individual level has gained traction. Yet, despite its importance for the prospects of implementing precision medicine in the clinic, phenotype predictability is still very low [48]. In the following paragraphs we summarize the reasons for the current state of phenotype prediction. We provide specific examples and calculations to visualize the performance of current genetic predictors, and end with a discussion of the prospects for using predictors in the clinic.

The most used method for risk profiling involves the calculation of Genetic Risk Scores (GRSs) that sum up the number of genome-wide significant alleles per individual in a target cohort. However, even if cases tend to have larger risk scores on average, the distributions of cases and controls tend to overlap for the most part. Although reasons related to the genetic architecture of complex traits partially account for the limited power of GRSs, this is also an unfortunate consequence of the limited fraction of disease susceptibility that GWAS have uncovered.

One of the most used measures to test the performance of any risk score is the area under the ROC curve (AUC), which summarizes the true positive rate (sensitivity) and false positive rate (1 - specificity) for all possible cut-offs. A classifier performing randomly and thus with null predictive power has an AUC of 0.5, whereas a perfect one reaches an AUC of 1. Thanks to the public availability of summary statistics from GWAS, it is straightforward to explore the performance of GRS in any target dataset. In Figure 3 we use genotypes from the WTCCC to offer a glimpse of the temporal progress of GRSs and resulting improvements in AUC values for five different conditions. In solid lines, we show the evolution of classical predictors based on genome-wide significant SNPs. The plot indicates that, although the list

of risk SNPs available in the GWAS Catalog has increased steadily over the years, the power of GRS seems to have plateaued and only tiny improvements in the ability to distinguish between cases and controls can be noticed for the most recent years. In the dashed lines, we add the equivalent calculations using variants detected by the COMBI algorithm that was recently proposed by some of us [45]. This method adds a preliminary step based on support vector machine that aims to improve the detection of variants, and which as seen in the figure, has the potential to refine GRSs and provide gains in performance.

Given the low power of GWAS to detect small effects, most causal variants are known to remain as false negatives. In recent years, genetic prediction has moved towards increasing the signal by using markers that do not achieve genome-wide significance. In this regard, polygenic risk scores (PRS) constitute a more powered alternative to GRS. PRS are calculated using SNPs selected through less stringent significance thresholds (e.g. $P < 10^{-3}$). However, even for PRS using thousands of SNPs the AUC for most traits is still below 0.7 [49]. These low figures indicate that GWAS have not yet mustered the large sample sizes – in the millions – that are needed to refine effect size estimates and achieve accurate individual risk profiles [50].

Although larger GWAS will lead to better characterization of the genetic architecture of disease, the field of genetic prediction is moving towards using evidence from genome-wide markers. Among others, a flurry of new methods is already using publicly available summary statistics from large studies [51], incorporating evidence from LD patterns to refine the effect size estimates (instead of simply focusing on the most significant variant for each associate loci) [52], and more recently, combining association evidence from several traits to refine trait prediction [53, 54]. Strategies integrating GWAS with gene expression data, such as **transcriptome-wide association studies** (TWAS) [55], or directly based on the latter such as the transcriptional risk score or TRS [56] also hold great promise for increasing prediction of phenotypes from genotypes.

Still, it is important to ensure the community maintains realistic hopes about the potential of genetic predictors. Ultimately, even if all causal variants and their true effects were known with exactitude, the heritability of each disease constitutes an upper-bound for the maximum phenotypic variance that could be accounted by a genetic predictor [50]. For instance, PRS for T2D will top off with AUC ~0.8 as GWAS progressively grow into sample size in the low millions [57].

More importantly, the foreseeable improvements in terms of AUC do not imply immediate translational potential in the clinical setting. Most complex diseases have low prevalence, and therefore even a specific cut off value in risk score that maximizes both the detection of future cases and the exclusion of healthy individuals can be offset by the fact that a large proportion of the detected individuals end up being false positives. The latter point is summarized by the positive predictive value (PPV), namely the proportion of ascertained individuals that will truly develop the disease, which can remain low even if the AUC goes up. A low PPV hurts the diagnostic power of risk scores because prediction of disease is not accurate at the individual level, and thus neutralizes some of the social and economic benefits of genetic testing for complex disease.

At present, the most immediate clinical application of risk scores lies in its potential for risk stratification and complement the information provided by other risk factors. To illustrate this point, we calculated a PRS for type 2 diabetes in the Kaiser RPGEH GERA cohort. As shown in Figure 4, the classification of individuals according to PRS deciles effectively captures an increasing fraction of T2D cases, which hints at its potential for stratification according to individual risk of disease. In fact, it is often underappreciated that many epidemiological risk factors have effect sizes that are like those of risk alleles discovered by GWAS, such as dietary factors that are known to increase risk of cancer [58]. This suggests that, rather than genomics information substituting classical epidemiology as the initial fad seemed to indicate, we are entering an era of integrative predictive medicine. In conclusion, while it is still too soon, we can anticipate that genetic profiling will help pinpointing individuals at high risk, which overall should lead to targeted lifestyle interventions and better decision making in the clinic.

Concluding Remarks and Future Perspectives

Genome-wide association studies have profoundly revolutionized both medical and complex trait genomics. To a large extent, this has been possible because findings in these fields, for the first time, are trustworthy. Although a variety of reasons account for this success, it is important to highlight the beneficial effect of having access to a robust methodology. GWAS led to the genetics community taking measures to avoid practices that lead to false discoveries, for example by enforcing strict p-value thresholds corrected for multiple testing to declare findings as statistically significant. Although we have witnessed amazing discoveries, we need new visions and methodologies to fully tackle the questions about the genetic architecture of complex traits. In particular, larger studies and improved methods will be needed to keep improving phenotype prediction to the level where it is justifiably commonplace in clinical settings. How quickly will this be fully achieved, rather than whether it will or not, is one of many unanswered questions (see Outstanding questions).

Outstanding Questions Box

- Will the field grow tired of GWAS in response to new paradigms, and, if so, will this decrease the interest in replicability that the field has seen in the last decade?
- If instead GWAS are scaled up to sample sizes in the order of millions, how will replicability be ensured in face of the increased levels of phenotypic and genetic heterogeneity? On the other hand, what will be the accuracy of genetic risk predictors based on these large-scale studies?
- What will be the replicability of findings arising from new methodologies such as machine learning or studies inspecting complex phenomena such as gene-by-environment and epistatic interactions? The extent to which good practices that fueled GWAS success can be implemented upon these studies is unknown.

- In what form will GWAS findings from genomic medicine be incorporated in the day-to-day clinical setting? The extent of accuracy they will bring to practice still remains open.
- Regarding other scientific fields, whether a generalized push for larger sample sizes will be enough to ensure replicability remains to be seen. More complex measures such as registering studies prior to completion (as it is done in clinical trials) should also help

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

AN and JAR were supported by Ministerio de Ciencia e Innovación, Spain (BFU2015-68649-P, MINECO/FEDER, UE), by Direcció General de Recerca, Generalitat de Catalunya (2014SGR1311 and 2014SGR866), and by the Spanish National Institute of Bioinformatics (PT13/0001/0026), the REEM (RD16/0015/0017) of the Instituto de Salud Carlos III. These authors have also received funding from the European Union's Horizon 2020 research and innovation programme 2014–2020 under Grant Agreement No. 634143 (*MedBioinformatics*). UMM and GG were funded by US NIH grants 1-P01-GM099568 (Project 3) and 2-R01-DK087694.

Glossary

Candidate-gene studies

The main strategy used in the prior to the GWAS era. These studies tested only a small set of mutations in a genic region with plausible involvement with disease.

Common disease-common variant

The model that disease is attributable to a moderate number of variants of common frequency. This was the main paradigm under the studies in the early 2000s and it played a key role in fuelling the genome-wide strategy, though it is no longer considered relevant.

Effect size

A quantitative measure of the strength of a phenomenon. In GWAS it has been generally reported as odds ratio for binary traits (diseases) and as regression coefficients (β) for quantitative traits.

Genetic architecture of disease

The underlying genetic basis of a phenotypic trait, which includes the total number of causal variants, their frequency and magnitude of effects, and degree of interactions with one another or environmental factors.

Genetic variance

The proportion of total variance in a certain phenotype that is explained by the combination factors.

Genome-wide significance

A p-value that is generally used as a threshold to consider a positive finding. In GWAS it is usually declared when a variant is associated with a $P < 5 \times 10^{-8}$. This value is considered

robust to avoid multiple testing issues and is roughly the Bonferroni value for 0.05 divided by ~1 million independent genomic intervals.

GWAS

Acronym for genome-wide association studies. GWAS analyse hundreds of thousands of markers, generally SNPs, comparing their frequencies across very large samples of individuals that share phenotypic characteristics.

Linkage disequilibrium (LD)

The non-random association of alleles at different loci. The levels of linkage disequilibrium vary among loci depending on local recombination rates. In humans, nearby genetic markers are usually linked in blocks of high LD, a feature that facilitates disease variant discovery through LD-based association mapping.

PAGE multi-ethnic cohort

The Population Architecture using Genomics and Epidemiology (PAGE) study is an ongoing consortium formed by several multi-ethnic population-based studies. The PAGE consortium has published several pioneer studies exploring the consistency of risk effects across genetic ancestries.

Publication bias

Phenomenon in which only a fraction of studies performed is eventually published and available to other researchers, usually involving that those studies failing to find an effect in the expected direction remain unpublished.

Re-discovery

finding of SNP-trait associations with suggestive statistical evidence or even genome-wide significance that is consistent with the discovery from the original study

Replicability

Property of a study or field whereby it produces the same or similar results when reproduced in a different sample, location or time.

TWAS

Acronym for transcriptome-wide association studies, a new methodology that uses genotypes and summary-statistic level data to impute gene expression in individuals participating in GWAS. This permits to identify significant gene-trait associations without the harsh multiple testing correction typical of GWAS.

Type I error

Statistical term that refers to the situation in hypothesis testing in which the null hypothesis is rejected despite being true (often called also as false positive). Given the large number of independent tests performed, the probability of type I errors increase exponentially in GWAS if multiple testing corrections are not considered.

References

1. Visscher PM, et al. 10 Years of GWAS Discovery: Biology, Function, and Translation. *Am J Hum Genet.* 2017; 101(1):5–22. [PubMed: 28686856]
2. Maher B. Personal genomes: The case of the missing heritability. *Nature.* 2008; 456(7218):18–21. [PubMed: 18987709]
3. Wellcome Trust Case Control C. Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature.* 2007; 447(7145):661–78. [PubMed: 17554300]
4. Boyle EA, et al. An Expanded View of Complex Traits: From Polygenic to Omnigenic. *Cell.* 2017; 169(7):1177–1186. [PubMed: 28622505]
5. Ioannidis JP. Why most published research findings are false. *PLoS Med.* 2005; 2(8):e124. [PubMed: 16060722]
6. Hilgard, J., Jamieson, KH. Science as “Broken” Versus Science as “Self-Correcting”: How Retractions and Peer-Review Problems Are Exploited to Attack Science. In: Jamieson, KH., et al., editors. *The Oxford handbook on the science of science communication.* Oxford University Press; 2017.
7. Plomin R, et al. Top 10 Replicated Findings From Behavioral Genetics. *Perspect Psychol Sci.* 2016; 11(1):3–23. [PubMed: 26817721]
8. Kraft P, et al. Replication in genome-wide association studies. *Stat Sci.* 2009; 24(4):561–573. [PubMed: 20454541]
9. Visscher PM, et al. Five years of GWAS discovery. *Am J Hum Genet.* 2012; 90(1):7–24. [PubMed: 22243964]
10. Marigorta UM, Navarro A. High trans-ethnic replicability of GWAS results implies common causal variants. *PLoS Genet.* 2013; 9(6):e1003566. [PubMed: 23785302]
11. Palmer C, Pe’er I. Statistical correction of the Winner’s Curse explains replication variability in quantitative trait genome-wide association studies. *PLoS Genet.* 2017; 13(7):e1006916. [PubMed: 28715421]
12. Waters KM, et al. Generalizability of associations from prostate cancer genome-wide association studies in multiple populations. *Cancer Epidemiol Biomarkers Prev.* 2009; 18(4):1285–9. [PubMed: 19318432]
13. Waters KM, et al. Consistent association of type 2 diabetes risk variants found in europeans in diverse racial and ethnic groups. *PLoS Genet.* 2010; 6(8)
14. Carlson CS, et al. Generalization and dilution of association results from European GWAS in populations of non-European ancestry: the PAGE study. *PLoS Biol.* 2013; 11(9):e1001661. [PubMed: 24068893]
15. Fu J, et al. Multi-ethnic studies in complex traits. *Hum Mol Genet.* 2011; 20(R2):R206–13. [PubMed: 21890495]
16. de Candia TR, et al. Additive genetic variation in schizophrenia risk is shared by populations of African and European descent. *Am J Hum Genet.* 2013; 93(3):463–70. [PubMed: 23954163]
17. Manolio TA. In Retrospect: A decade of shared genomic associations. *Nature.* 2017; 546(7658):360–361. [PubMed: 28617469]
18. Martin AR, et al. Human Demographic History Impacts Genetic Risk Prediction across Diverse Populations. *Am J Hum Genet.* 2017; 100(4):635–649. [PubMed: 28366442]
19. Nakagome S, et al. Crohn’s disease risk alleles on the NOD2 locus have been maintained by natural selection on standing variation. *Mol Biol Evol.* 2012; 29(6):1569–85. [PubMed: 22319155]
20. Wang YF, et al. Clinical manifestations of inflammatory bowel disease: East and West differences. *J Dig Dis.* 2007; 8(3):121–7. [PubMed: 17650222]
21. consortium C. Sparse whole-genome sequencing identifies two loci for major depressive disorder. *Nature.* 2015; 523(7562):588–91. [PubMed: 26176920]
22. Klein RJ, et al. Complement factor H polymorphism in age-related macular degeneration. *Science.* 2005; 308(5720):385–9. [PubMed: 15761122]

23. Fritsche LG, et al. A large genome-wide association study of age-related macular degeneration highlights contributions of rare and common variants. *Nat Genet.* 2016; 48(2):134–43. [PubMed: 26691988]
24. Cotsapas C, et al. Pervasive sharing of genetic effects in autoimmune disease. *PLoS Genet.* 2011; 7(8):e1002254. [PubMed: 21852963]
25. Moser G, et al. Simultaneous discovery, estimation and prediction analysis of complex traits using a bayesian mixture model. *PLoS Genet.* 2015; 11(4):e1004969. [PubMed: 25849665]
26. Luciano M, et al. Association analysis in over 329,000 individuals identifies 116 independent variants influencing neuroticism. *Nat Genet.* 2018; 50(1):6–11. [PubMed: 29255261]
27. Diabetes Genetics Initiative of Broad Institute of H et al. Genome-wide association analysis identifies loci for type 2 diabetes and triglyceride levels. *Science.* 2007; 316(5829):1331–6. [PubMed: 17463246]
28. Scott LJ, et al. A genome-wide association study of type 2 diabetes in Finns detects multiple susceptibility variants. *Science.* 2007; 316(5829):1341–5. [PubMed: 17463248]
29. Timpson NJ, et al. Adiposity-related heterogeneity in patterns of type 2 diabetes susceptibility observed in genome-wide association data. *Diabetes.* 2009; 58(2):505–10. [PubMed: 19056611]
30. Kulminski AM, et al. Explicating heterogeneity of complex traits has strong potential for improving GWAS efficiency. *Sci Rep.* 2016; 6:35390. [PubMed: 27739495]
31. Magosi LE, et al. Identifying systematic heterogeneity patterns in genetic association meta-analysis studies. *PLoS Genet.* 2017; 13(5):e1006755. [PubMed: 28459806]
32. Hirschhorn JN, et al. A comprehensive review of genetic association studies. *Genet Med.* 2002; 4(2):45–61. [PubMed: 11882781]
33. Lohmueller KE, et al. Meta-analysis of genetic association studies supports a contribution of common variants to susceptibility to common disease. *Nat Genet.* 2003; 33(2):177–82. [PubMed: 12524541]
34. Ioannidis JP, et al. The false-positive to false-negative ratio in epidemiologic studies. *Epidemiology.* 2011; 22(4):450–6. [PubMed: 21490505]
35. Siontis KC, et al. Replication of past candidate loci for common diseases and phenotypes in 100 genome-wide association studies. *Eur J Hum Genet.* 2010; 18(7):832–7. [PubMed: 20234392]
36. Ioannidis JP, et al. Implications of small effect sizes of individual genetic variants on the design and interpretation of genetic association studies of complex diseases. *Am J Epidemiol.* 2006; 164(7):609–14. [PubMed: 16893921]
37. Weiss KM, Terwilliger JD. How many diseases does it take to map a gene with SNPs? *Nat Genet.* 2000; 26(2):151–7. [PubMed: 11017069]
38. Contreras JL. NIH's genomic data sharing policy: timing and tradeoffs. *Trends Genet.* 2015; 31(2):55–7. [PubMed: 25620797]
39. Asking for more. *Nat Genet.* 2012; 44(7):733. [PubMed: 22735581]
40. Chanock SJ, et al. Replicating genotype-phenotype associations. *Nature.* 2007; 447(7145):655–60. [PubMed: 17554299]
41. Yang J, et al. Genetic variance estimation with imputed variants finds negligible missing heritability for human height and body mass index. *Nat Genet.* 2015; 47(10):1114–20. [PubMed: 26323059]
42. Fadista J, et al. The (in)famous GWAS P-value threshold revisited and updated for low-frequency variants. *Eur J Hum Genet.* 2016; 24(8):1202–5. [PubMed: 26733288]
43. Keller MC. Gene x environment interaction studies have not properly controlled for potential confounders: the problem and the (simple) solution. *Biol Psychiatry.* 2014; 75(1):18–24. [PubMed: 24135711]
44. Marigorta UM, Gibson G. A simulation study of gene-by-environment interactions in GWAS implies ample hidden effects. *Front Genet.* 2014; 5:225. [PubMed: 25101110]
45. Mieth B, et al. Combining Multiple Hypothesis Testing with Machine Learning Increases the Statistical Power of Genome-wide Association Studies. *Sci Rep.* 2016; 6:36671. [PubMed: 27892471]

46. Pickrell JK. Joint analysis of functional genomic data and genome-wide association studies of 18 human traits. *Am J Hum Genet.* 2014; 94(4):559–73. [PubMed: 24702953]
47. Wood AR, et al. Defining the role of common variation in the genomic and biological architecture of adult human height. *Nat Genet.* 2014; 46(11):1173–86. [PubMed: 25282103]
48. Rodriguez JA, et al. Integrating genomics into evolutionary medicine. *Curr Opin Genet Dev.* 2014; 29:97–102. [PubMed: 25218863]
49. So HC, Sham PC. Exploring the predictive power of polygenic scores derived from genome-wide association studies: a study of 10 complex traits. *Bioinformatics.* 2017; 33(6):886–892. [PubMed: 28065900]
50. Wray NR, et al. Pitfalls of predicting complex traits from SNPs. *Nat Rev Genet.* 2013; 14(7):507–15. [PubMed: 23774735]
51. Robinson MR, et al. Genetic evidence of assortative mating in humans. *Nature Human Behaviour.* 2017; 1:0016.
52. Vilhjalmsón BJ, et al. Modeling Linkage Disequilibrium Increases Accuracy of Polygenic Risk Scores. *Am J Hum Genet.* 2015; 97(4):576–92. [PubMed: 26430803]
53. Turley P, et al. Multi-trait analysis of genome-wide association summary statistics using MTAG. *Nat Genet.* 2018; 50(2):229–237. [PubMed: 29292387]
54. Maier RM, et al. Improving genetic prediction by leveraging genetic correlations among human diseases and traits. *Nat Commun.* 2018; 9(1):989. [PubMed: 29515099]
55. Gusev A, et al. Integrative approaches for large-scale transcriptome-wide association studies. *Nat Genet.* 2016; 48(3):245–52. [PubMed: 26854917]
56. Marigorta UM, et al. Transcriptional risk scores link GWAS to eQTLs and predict complications in Crohn’s disease. *Nat Genet.* 2017; 49(10):1517–1521. [PubMed: 28805827]
57. Chatterjee N, et al. Projecting the performance of risk prediction based on polygenic analyses of genome-wide association studies. *Nat Genet.* 2013; 45(4):400–5. 405e1–3. [PubMed: 23455638]
58. Figueiredo JC, et al. Genome-wide diet-gene interaction analyses for risk of colorectal cancer. *PLoS Genet.* 2014; 10(4):e1004228. [PubMed: 24743840]
59. Risch N, Merikangas K. The future of genetic studies of complex human diseases. *Science.* 1996; 273(5281):1516–7. [PubMed: 8801636]
60. Carney DR, et al. Power posing: brief nonverbal displays affect neuroendocrine levels and risk tolerance. *Psychol Sci.* 2010; 21(10):1363–8. [PubMed: 20855902]
61. Ranehill E, et al. Assessing the robustness of power posing: no effect on hormones and risk tolerance in a large sample of men and women. *Psychol Sci.* 2015; 26(5):653–6. [PubMed: 25810452]
62. Rousseau DL. “Polywater” and sweat:similarities between the infrared spectra. *Science.* 1971; 171(3967):170–2. [PubMed: 5538826]
63. Caspi A, et al. Influence of life stress on depression: moderation by a polymorphism in the 5-HTT gene. *Science.* 2003; 301(5631):386–9. [PubMed: 12869766]
64. Szucs D, Ioannidis JP. Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biol.* 2017; 15(3):e2000797. [PubMed: 28253258]
65. Benjamin DJ, et al. Redefine statistical significance. *Nature Human Behaviour.* 2017
66. Patil P, et al. What Should Researchers Expect When They Replicate Studies? A Statistical View of Replicability in Psychological Science. *Perspect Psychol Sci.* 2016; 11(4):539–44. [PubMed: 27474140]
67. Jager LR, Leek JT. An estimate of the science-wise false discovery rate and application to the top medical literature. *Biostatistics.* 2014; 15(1):1–12. [PubMed: 24068246]

Box 1**Quick glance at the history and rationale of GWAS**

The two-page perspective by Risch and Merikangas in 1996 [59] can be considered as the start of the GWAS era. This landmark paper demonstrated that association studies based on polymorphisms spanning all over the genome would be powered to detect susceptibility variants of small effects. The key idea for association mapping lies in exploiting the correlation structure among markers that are nearby in our genomes (linkage disequilibrium, LD). This property allows carrying out indirect association tests, whereby it is only necessary to type and test a single variant (tagSNP) selected to act as a surrogate for all variants in the block of LD.

GWAS look for alleles transmitted with disease status and therefore with significantly different frequencies between cases and controls. After imputation of variants, GWAS test millions of SNPs and hence it is important to account for multiple testing to avoid false positives. The most used threshold to declare statistical evidence of association is $P < 5 \times 10^{-8}$, which arises from adjusting a 5% type I error rate (false positives) to the roughly one to ten million independent tests performed.

Attention to other statistical matters is also important. Among others, extensive quality control steps include ensuring the reliability of the genotyping data, careful evaluation of genetic ancestry of the samples to avoid population stratification that may lead to p-value inflation if unaccounted for, and adjustment for clinical covariates such as age, gender and lifestyle to avoid spurious associations and biases in the estimation of effect size. Finally, GWAS are increasingly performed through large meta-analyses that combine statistical evidence from multiple cohorts. Even if all sub-studies test the same hypothesis, it is important to quantify the presence of heterogeneity to identify outliers that should be excluded.

Overall, the GWAS era has constituted an impressive tour de force. As it is beautifully shown in the famous karyotype diagram quarterly released by the GWAS Catalog (<https://www.ebi.ac.uk/gwas/diagram>), we now know of over 12,000 SNPs associated at genome-wide significance to a myriad of traits. Overall, GWAS have demonstrated a sweeping performance in human genetics and will likely remain as the main tool to discover genes involved in disease in the next future.

Box 2**Replicability crises in other scientific fields**

In 2010 Carney *et al.* reported that adopting an expansive body posture before facing a stressful situation could boost self-confidence and authority by increasing testosterone and decreasing cortisol levels [60]. In spite of attracting plenty of attention, this finding failed to replicate in a larger study on power posing and hormonal levels [61]. Similar to social psychology, neuroimaging has also been on the spot in recent times. Of note, similar false positives have been ubiquitous across scientific fields [5], including physics research in the 1960s through the discovery of the ‘polywater’, a hypothetical polymerized form of water that was eventually debunked [62].

Genomics itself has not escaped from the lack of replication problem, including the highly cited and yet still debated gene-by-environment interaction between the serotonin transporter gene, adverse life events and risk of depression [63]. In this regard, lack of correction for multiple testing is often the primary factor leading to false positives. In many fields, $P < 0.05$ is still deemed as enough evidence for calling an association as true. Given that studies are generally underpowered, using such a liberal threshold is a recipe for false positives [64]. Genomics and physics have shielded themselves against this problem by imposing highly stringent p-value cut-offs, setting an example that should be embraced by scientists in other communities [65].

Some scholars have stated that this replicability crisis is overblown [66]. A major issue is that controlling the false positive rate (FPR) depends on the fraction of hypotheses tested that is true [5]. Thus, if only 1% of hypotheses turn out to be correct, then with 80% power at $p < 0.05$ then only 8 of 1,000 tests will be true positives, relative to 50 false positives. Happily, statistical analysis of reported p-value distributions strongly implies that the true positive rate is much higher and the FPR is well controlled across much of science [67]. Also on the bright side, the consensus about the importance of ensuring enough sample sizes is steadily growing. In this regard, several initiatives to detect false positives have emerged, namely the Reproducibility Project, Pubpeer, clinicaltrials.gov, and Retraction Watch, among many others. With all this into account, if we want to assure a future of sustainable, reproducible and replicable science we need to take direct action and extend the basic principles that have been shown to work to as many fields of research as possible.

Highlights or Trends

- The decade of GWAS constitutes a clear improvement in the recent history of reproducibility in medical genetic research. The extremely high rates of replication imply that, for the first time, findings can be trusted.
- Large number of false positives and the tiny effect size of genetic risk variants induced a change in incentives during the GWAS era, priming the requirement for large sample sizes and a culture of data sharing.
- There is increasing interest and need of new methodologies to better understand the genetic architecture of complex traits.
- It is necessary to keep fostering a culture of compulsory replication to maintain the current high reliability in findings.
- Although the success of GWAS has not translated into an ability to predict phenotypes based on genetic markers, polygenic and transcriptional risk scores (PRS and TRS) hold potential for stratification according to risk.

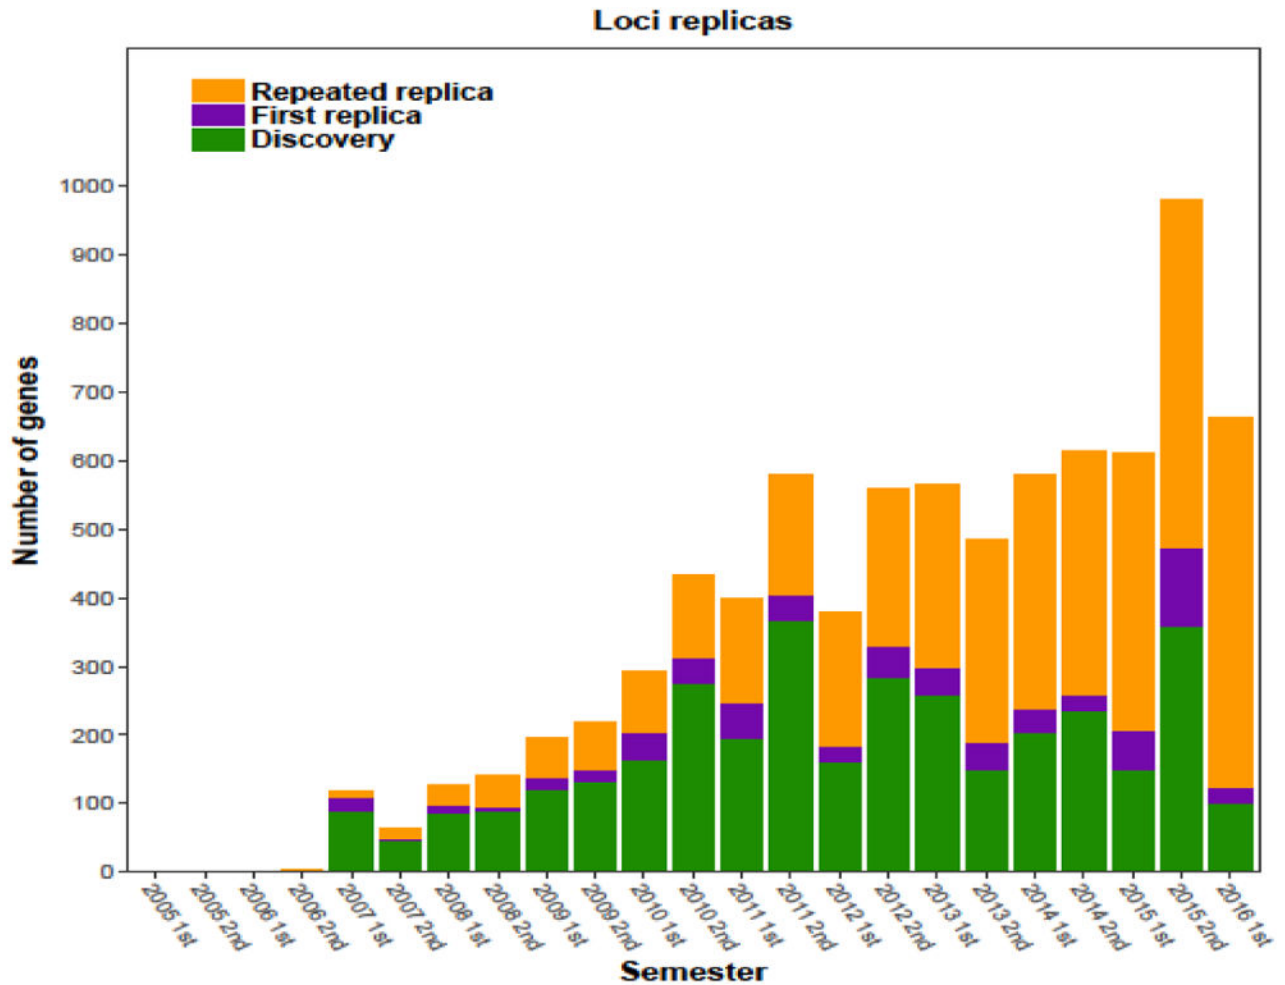


Figure 1, Key Figure. An increasing proportion of GWAS findings correspond to replications of SNP-trait associations that had been previously reported

The graph reports the number of discoveries of new loci and re-discovery of previously discovered loci for 60 diseases included in the GWAS Catalog (quantitative traits were not considered; last accessed 17th March 2017). Data are classified according to the semester of publication (from 2005 to 2016). Given that all records included in the Catalog achieve at least $P < 10^{-5}$, all newly included SNP-disease pairs that had already been included because of being discovered by previous GWAS can be considered as replications of the first finding. We labelled these instances as re-discoveries, or replica. The re-discovery figure for a given semester corresponds to the total cumulative number since 2005, separated according to whether the re-discovery event constitutes the first evidence for replication of a given locus (*'first replica'*) or if it had been previously replicated (*'repeated replica'*). Given the diversity of arrays before the recent generalization of imputation in GWAS, SNPs are considered replicated either when the same SNP is re-discovered or another SNP in $R^2 \geq 0.8$ (using 1KG Europeans) within a ± 500 kb window is reported. Only GWAS performed in Eurasian populations were evaluated.




Figure 2. Risk variants discovered by GWAS replicate similarly regardless of the Odds Ratio (OR)

X-axis: Classified by disease, percentage of SNPs with a discovery OR between 1.2 and 1.5 that have been re-discovered at least once and included in the GWAS Catalog (filled circles) or that achieve nominal evidence of association ($P < 0.05$) in a large meta-analysis published after the discovery GWAS (inverted triangles). Y-axis: for the same diseases, proportion of re-discovery for risk SNPs with discovery OR < 1.2. Although the percentages vary by disease, the average re-discovery of low-effect variants is not significantly lower than that of variants with larger effect ($P = 0.93$, Kolmogorov-Smirnov test) and overall re-discovery estimates are highly correlated (Spearman's rho: 0.54, $P = 10^{-16}$).

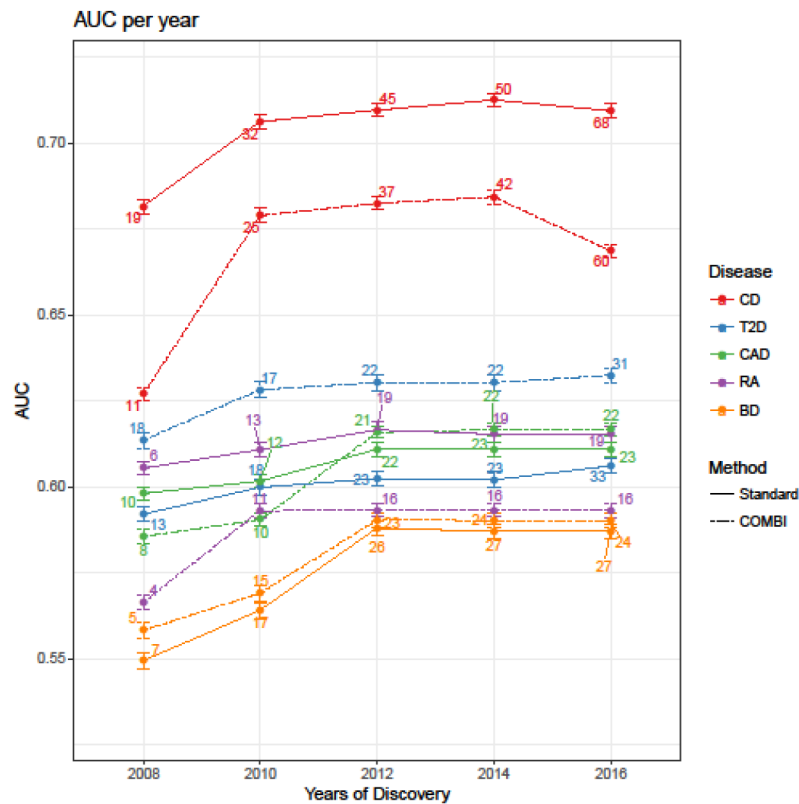


Figure 3. Increased availability of known risk variants and new methodologies are improving the ability to classify individuals according to disease

We calculated the area under the ROC curve (AUC) for five diseases using the samples from the WTCCC paper and SNPs discovered at different year ranges. X-axis: Discovery year of disease-associated variants included in the genetic risk score for each disease (based on year of inclusion in the GWAS Catalog). Y-axis: AUC using samples from the WTCCC case-control study. Performance of two predictors is shown, namely, a standard GRS (dashed lines) and a modified version of the GRS using the number of risk alleles times the weight for each SNP, as derived from the COMBI algorithm (solid lines). The number of SNPs used in each AUC calculation are indicated in the small numbers next to each point. Even if the number of risk SNPs has increased steadily, for both methodologies we observe a plateau in the predictive power of the genetic risk scores.




Figure 4. Polygenic Risk Scores (PRS) stratify individuals according to risk of disease
X-axis: 55,210 samples of European ancestry from the Kaiser GERA Cohort are classified according to deciles of PRS for type 2 diabetes (T2D). Y-axis: Percentage of patients in each category. The graph shows the impact of increasing deciles of a weighted genetic risk score based on 414 LD-pruned SNPs associated with T2D ($P < 10^{-3}$ in the 2014 trans-ethnic DIAGRAM GWAS). The PRS captures risk of disease according to genetic makeup of individuals, with a 2-fold enrichment of cases in the top vs. lowest decile.