The ubiquity of pleiotropy in human disease

Kevin Chesmore¹ · Jacqueline Bartlett² · Scott M. Williams^{2,3}

Received: 10 August 2017 / Accepted: 14 November 2017 / Published online: 21 November 2017
© The Author(s) 2017. This article is an open access publication

Abstract

Pleiotropy has long been thought to be a common phenomenon in the human genome; however, until recently appropriate data was unavailable to test this hypothesis. Prior studies focused on assessing the prevalence of pleiotropy in only small subsets of phenotypes (≤ 53 phenotypes), without a truly comprehensive assessment of pleiotropy in the human genome. In this study, we determined the prevalence of pleiotropy, using the entire GWAS catalog (1094 disease phenotypes, 14,459 genes), as well as investigate the relationship between the degree of pleiotropy and the average effect size for each associating gene. The number of associating phenotypes per gene ranged from 1 to 53, with 44% of genes reported in the GWAS catalog associating with more than one phenotype. The proportion of genes shown to be pleiotropic has continued to increase as more studies are added to the catalog. We also found the degree of pleiotropy scales positively with a gene's average effect size ($r = 0.04$, p value = 0.0003) and negatively with the variance of effect sizes in genes with a given number of associating phenotypes ($r = -0.590$, p value = 0.0019). Based on this and prior work, it is becoming evident that pleiotropy is a common, if not ubiquitous, phenomenon. These results have implications in understanding disease etiologies, potentially common biology underlying even disparate diseases, and in elucidating the genotype–phenotype map.

Introduction

Pleiotropy, defined here as one gene affecting more than one phenotype, has become increasingly important in interpreting both genotype–phenotype maps and underlying factors that affect comorbidity, and has been integrated into several methods for genome wide association studies (GWAS) (Chung et al. 2014; Hill and Zhang 2012; Lee et al. 2012; Li et al. 2014a, b; Pendergrass et al. 2013; Rzhetsky et al. 2007; Sivakumaran et al. 2011; Solovieff et al. 2013; Stearns 2010; Wagner et al. 2008; Wagner and Zhang 2011; Yang

et al. 2015). Studies of pleiotropy can impact our understanding of underlying disease processes, and implicitly, the prevention and treatment of disease as well as predicting adverse reactions to drugs. Such studies can also greatly aid in our understanding of the relationships between diseases based on the shared genetic architecture and allow us to distinguish between horizontal (multiple independent phenotypic effects) and vertical (cross-phenotype effects) forms of pleiotropy (Tyler et al. 2009). However, using pleiotropy to understand disease etiologies will only be productive if pleiotropy is common. While pleiotropy has been suspected to be common in the human genome, it is only recently that large-scale data to test this claim have been available. In a recent study, researchers were able to identify 341 pleiotropic loci across 42 traits (Pickrell et al. 2016; Visscher and Yang 2016). Considering how many pleiotropic loci were found from such a limited number of traits, it was concluded that pleiotropy is likely a common phenomenon in the human genome. However, while this evidence does support the hypothesis that pleiotropy is common, it is far from a comprehensive analysis of genetic association studies.

Despite the recognized importance of pleiotropy in developing a comprehensive understanding of gene–disease relationships (Chung et al. 2014; Hill and Zhang 2012; Lee et al. 2012; Li et al. 2014a, b; Pendergrass et al. 2013; Rzhetsky

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s00439-017-1854-z>) contains supplementary material, which is available to authorized users.

✉ Scott M. Williams
smw154@case.edu

¹ Department of Genetics, Geisel School of Medicine, Dartmouth College, Hanover, NH, USA

² Departments of Population and Quantitative Health Sciences, Institute of Computational Biology, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106, USA

³ Department of Genetics and Genome Sciences, Case Western Reserve University, Cleveland, OH, USA

et al. 2007; Sivakumaran et al. 2011; Solovieff et al. 2013; Stearns 2010; Wagner et al. 2008; Wagner and Zhang 2011; Yang et al. 2015), little has been done to investigate the relationship between pleiotropy and risk variation as measured by effect sizes. Studies investigating this question have focused on exploring the correlation of effect sizes of different traits associated with a single locus; however, no study has investigated the relationship between effect size and the degree of pleiotropy in humans. One study, using QTL associations addressed the relationship in sticklebacks between the degree of pleiotropy and allelic effect sizes, found a positive correlation (Albert et al. 2008; Wagner and Zhang 2011). Similar results have also been found in a variety of model organisms (Wang et al. 2010). While such studies are useful, no comparable study has ever been done to address this relationship with human disease.

Methods

To investigate the ubiquity of pleiotropy and its relationship to effect size in humans, we examined the the GWAS catalog (gwas_catalog_v1.0-associations_e89, collected in June, 2017, <https://www.ebi.ac.uk/gwas/>), a comprehensive database of genetic association studies (Visscher et al. 2012; Welter et al. 2014). Each significant SNP (p value $< 10e-6$) was assigned a gene based on the closest mapped gene(s) (annotations are based on GRCh38.p10 reference genome build and dbSNP Build 150) as listed in the GWAS catalog. We also performed our analyses using the more standard and conservative threshold for significance ($5 \times 10e-8$). In the case of an intergenic SNP, the SNP was assigned to both flanking genes. In total 14,459 genes were identified in the GWAS catalog. Phenotypes were manually curated to combine similar or identical phenotypes listed under different names [i.e., “Fasting plasma glucose” and “Fasting plasma glucose (childhood)”]. Studies of phenotypes in a diseased population, such as “Cardiovascular heart disease in diabetics” were listed separately from either the primary phenotype of interest (“Cardiovascular heart disease”) or the disease context of the population (“Diabetes”). This distinction was kept because these studies do not address the genetics of the population’s disease context and the genetics of the primary phenotype is likely to be distinct from that in an otherwise healthy population. This approach resulted in 1094 phenotypes and 46,211 gene–phenotype associations (of which 16,867 are dichotomous phenotypes) from the GWAS catalog.

Effect sizes were analyzed for dichotomous traits only, as we are particularly interested in investigating disease effects. Additionally, effect sizes for dichotomous traits are not directly comparable to those of continuous variables without standardizing the raw data from the association studies.

The average effect size for each gene was calculated as follows: i = each phenotype associated with gene “ x ”
 j = each SNP associated with phenotype “ i ” and gene “ x ”

$$\text{Effect size} = \frac{\sum_i^n \left(\frac{\sum_j^m |\log(\text{Odds Ratio}_{ij})|}{m_i} \right)}{n}$$

The purpose of using the absolute value of the log was to place effects that increase or decrease risk on the same scale. For example, odds ratios of 0.5 and 2 have the same magnitude of effect; however, they are acting in opposite directions. However, if we were to take the $\log(\text{Odds Ratio})$, both effect sizes would be equal to 0.3. This calculation also ensures that each phenotype contributes to the gene’s average effect size equally.

To address how the above data have changed as GWAS studies have become more frequent and methods have improved, we calculated the number of phenotypes and average effect size for each gene based on the data available for each year between 2005 and 2017. In this study, we explore the yearly changes in the distribution of the number of phenotypes and the average effect size for each gene in our dataset, as well as the correlation between these two variables.

Results

Currently, the number of diseases with which each gene in the GWAS catalog associated ranges from 1 to 53 (Fig. 1a). Of these genes, 44% associated with more than one disease (Fig. 1a). The ten most pleiotropic genes (between 27 and 53 disease associations) included *ABO* (53 diseases), *GCKR* (49), *MTCO3P1* (46), *HLA-DQB1* (43), *LOC107986589* (42), *HLA-DRB1* (41), *CSMD1* (35), *CDKN2B-AS1* (35), *LOC101929163* (33), and *PTPRD* (27) (Table S1). Eight of these ten genes remain when we use a more stringent significance threshold and the number of phenotypes associating with the most pleiotropic genes ranges from 17 to 39 (Table S2). In addition, using $5 \times 10e-8$ as the significance threshold did not substantially affect the percentage of pleiotropic genes as it only dropped to 39% (Figure S2A). As expected, from 2005 to 2017, the distribution of the number of genes for which evidence of pleiotropy existed increased monotonically. In 2005 and 2006, all genes in the catalog associated with only a single phenotype. In 2007, the first instances of pleiotropic loci were identified (Fig. 1b, Figure S2B, Table S3). The fraction of genes associated with more than one disease has continuously increased 2008 through 2017 (Fig. 1b, Figure S1, Figure S2B for the more stringent significance threshold).

The current distribution of effect sizes is exponentially distributed, ranging between 0.004 and 2.125 with a mean

Fig. 1 Distribution of the number of phenotypes each gene is associated with in the GWAS catalog. **a** Histogram showing the proportion of genes that associate with phenotypes in the GWAS catalog (y-axis) and the number of phenotypes with which they associate (x-axis). **b** The proportion of genes associated with 2 or more phenotypes (y-axis) for each year (x-axis)

of 0.245 (Fig. 2a, and Figure S3A between 0.006 and 2.65 with a mean of 0.166 for a p value threshold of 5×10^{-8}). Average log transformed effect sizes have remained stable since 2008 (Fig. 2b, Figure S3B); however, variances have increased steadily (Fig. 2c; $r = 0.750$, p value = 0.008, Figure S3C; $r = 0.602$, p value = 0.0498). Additionally, the variance of effect sizes appears to scale negatively with the number of phenotype associations (Fig. 3; $r = -0.590$, p value = 0.0019, Figure S4; $r = -0.487$, p value = 0.0472).

The correlation between the number of phenotypes and the average log transformed effect sizes of each gene is significant, and trends in a positive direction (Fig. 4a; $r = 0.04$, p value = 0.0003, Figure S5A; $r = 0.09$, p value = 3.23×10^{-7} using the 5×10^{-8} threshold). Removing genes with very few associations from this analysis results in an improved correlation where it was found that genes associating with more than 9 phenotypes had a correlation coefficient on 0.14 (Fig. 4b; p value = 0.0017). However, this correlation loses its significance under our more stringent cutoff of 5×10^{-8} (Figure S5B; p value = 0.2021).

Fig. 2 Distribution of effect sizes in the GWAS catalog. **a** Histogram showing the proportion of genes in the GWAS catalog (y-axis) with a given effect size (x-axis). **b** Boxplot of $\log_{10}$ (effect sizes) (y-axis) for each year from 2005 to 2017 (x-axis). **c** Variance of $\log_{10}$ (effect sizes) (y-axis) for each year from 2005 to 2017 (x-axis)

Fig. 3 Variance of $\log_{10}$ (effect size) (y-axis) for all genes with a given number of phenotypes (x-axis)

Fig. 4 Relationship between Pleiotropy and effect size. **a** the number of phenotypes (x-axis) vs. the log₁₀ (effect size) (y-axis). **b** Pearson's *r* correlation [number of phenotypes vs log₁₀ (effect size)] (y-axis) for all genes with more than the number of phenotypes (> X) (x-axis)

Discussion

Prior work attempting to catalog the prevalence of pleiotropy in the human genome has suggested that pleiotropy is common, but the scope of earlier studies was limited by the number of phenotypes investigated (≤ 53 phenotypes) (Bulik-Sullivan et al. 2015; Pickrell et al. 2016; Visscher and Yang 2016). We expand upon these prior findings by investigating all 1094 phenotypes listed in the GWAS catalog as of mid-2017. In this study, we found that nearly half of genes reported in the GWAS catalog are pleiotropic, and it is reasonable to conclude that this percentage will continue to increase as additional studies are added to the GWAS catalog. The prevalence of pleiotropy identified in this study is comparable to that found in a previous study of the GWAS catalog [for all GWAS completed up to 2011 (Sivakumaran et al. 2011)]. Minor differences between this study and ours can be explained by differences in gene and phenotype curation. Another study in 2017 estimated that 12% of protein-coding genes listed in UniProt are pleiotropic (Ittisoponpisan et al. 2017). While our findings differ, it is worth noting that the database and underlying hypotheses differ from ours and the findings may not be comparable.

From these data, we were also able to identify the most pleiotropic genes in the human genome; among these, we identified genes that are primarily involved in immune and metabolic functions as did prior studies (Ittisoponpisan et al. 2017; Sivakumaran et al. 2011), although this may reflect a bias in the classes of diseases frequently examined in GWAS (Table S1). Some notable highly pleiotropic genes include: ABO, members of the HLA gene family, and the mitochondrial protein MTCO3P1. Additionally, it is worth noting that two of the most pleiotropic genes are non-coding RNAs: *LOC107986589*, *LOC101929163*. These findings indicate that non-coding RNAs can be just as pleiotropic as protein-coding genes. Interestingly, despite being two of the most pleiotropic (and potentially clinically relevant) genes in the human genome, nothing is known regarding the functions of these non-coding RNAs; neither has an entry in OMIM.

From 2005 to 2007, GWASs were only able to detect the strong effects (Fig. 2b). However, as experimental designs and methodologies improved GWASs were able to detect subtler effects. From 2008 to 2017, the mean log₁₀ effect sizes have remained fairly stable. However, the variance of effect sizes has increased each year. The increasing variance may be explained by two trends in GWAS data: (1) higher density SNP chips allow for higher likelihood of tagging the causal variant or another SNP in close proximity to it, resulting in higher effect sizes, and (2) larger sample sizes, allow for the detection of SNPs with lower effect sizes. Together, these two phenomena will spread the distribution of the effect sizes.

We found that the variance of effect sizes among genes with fewer phenotype associations was higher than those with more associations. Two hypotheses may explain this observation: (1) many, if not all, genes used in this study likely have more phenotype associations than are currently known. As additional studies are added to the GWAS catalog the distribution of the number of phenotypes that genes are associated with should converge on the actual level of pleiotropy. Along with this change in detected pleiotropy, we further hypothesize that the variance of effect sizes for genes with some given number phenotypes will decrease; (2) in accordance with the law of large numbers the more phenotypes a gene associates, the more accurate the estimated “average effect size”. Inversely, for genes with very few associations the estimate of the average effect size will not accurately represent the true average effect sizes, resulting in a larger variance among genes with fewer phenotypic associations.

The correlation between the log₁₀ (effect size) and the number of phenotypes showed highly significant, albeit weak correlation. This weak correlation is not unexpected because the GWAS catalog is an incomplete dataset, as many gene–disease associations are still unknown. As more studies are added to the catalog, we will gain a more accurate

view of the degree of pleiotropy and average effect sizes for each gene; this should refine this relationship. Many of the issues encountered in this dataset occur in those genes with very few phenotype associations; by ignoring these highly variable data points, the correlation between pleiotropy and effect size significantly improves. We acknowledge that there may be confounding variables that are unaccounted for in this analysis, such as the effect of gene length. Nonetheless, based on the current results, we conclude that effect size correlates positively with pleiotropy. This conclusion is also supported by previous work from non-human studies (Pickrell et al. 2016; Visscher and Yang 2016; Wang et al. 2010). Three (non-mutually exclusive) hypotheses may explain why effect size scales positively with level of pleiotropy:

1. Observation bias—since we know that it is easier to detect associations with larger effect sizes, we would be more likely to detect many associations with genes that have (on average) larger effect sizes, and fewer associations for those genes which have smaller effect sizes. For example, in a hypothetical scenario, if 2 genes both contributed to exactly 10 phenotypes and “gene X” had an average Odds Ratio of 2.00 and “gene Y” had an average Odds Ratio of 1.05, we might expect to identify more of the phenotypes associating with “gene X” than with “gene Y”.
2. Pleiotropy drives effect size—the more phenotypes with which a gene is associated, the more profoundly that gene affects those phenotypes. A gene that associates with many diseases will likely be involved in many biological processes (a physiological hub). Therefore, mutations in these genes will likely have a stronger phenotypic effect.
3. Effect size drives pleiotropy—the larger the effect size, the more likely the gene is to be involved in multiple phenotypes, perhaps due to secondary effects of one phenotypic outcome. The observation of scaling of effects as a function of level of pleiotropy has been previously argued, albeit differently, in previous studies (Li et al. 2014a).

The findings in this study build upon prior research, giving support to the hypothesis that pleiotropy is probably ubiquitous in the human genome. In this study, it was shown that nearly half of disease associated genes reported in the GWAS catalog are pleiotropic. There is also a positive correlation between degree of pleiotropy and average effect size, a relationship which until now has only been explored in model systems. As additional studies are added to the GWAS catalog, we will accurately determine the true levels of pleiotropy and its relationship to effect size. Understanding this relationship is a crucial step towards understanding

genotype–phenotype maps and the complex relationship that exists between human diseases.

Acknowledgements This was funded by NIH Grants T32GM008704, P20GM103534 and 2R01LM010098.

Compliance with ethical standards

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Albert AY et al (2008) The genetics of adaptive shape shift in stickleback: pleiotropy and effect size. *Evolution* 62:76–85. <https://doi.org/10.1111/j.1558-5646.2007.00259.x>
- Bulik-Sullivan B et al (2015) An atlas of genetic correlations across human diseases and traits. *Nat Genet* 47:1236–1241. <https://doi.org/10.1038/ng.3406>
- Chung D, Yang C, Li C, Gelernter J, Zhao H (2014) GPA: a statistical approach to prioritizing GWAS results by integrating pleiotropy and annotation. *PLoS Genet* 10:e1004787. <https://doi.org/10.1371/journal.pgen.1004787>
- Hill WG, Zhang XS (2012) On the pleiotropic structure of the genotype–phenotype map and the evolvability of complex organisms. *Genetics* 190:1131–1137. <https://doi.org/10.1534/genetics.111.135681>
- Ittisoponpisan S, Alhuzimi E, Sternberg MJ, David A (2017) Landscape of pleiotropic proteins causing human disease: structural and system biology insights. *Hum Mutat* 38:289–296. <https://doi.org/10.1002/humu.23155>
- Lee SH, Yang J, Goddard ME, Visscher PM, Wray NR (2012) Estimation of pleiotropy between complex diseases using single-nucleotide polymorphism-derived genomic relationships and restricted maximum likelihood. *Bioinformatics* 28:2540–2542. <https://doi.org/10.1093/bioinformatics/bts474>
- Li C, Yang C, Gelernter J, Zhao HY (2014a) Improving genetic risk prediction by leveraging pleiotropy. *Hum Genet* 133:639–650. <https://doi.org/10.1007/s00439-013-1401-5>
- Li L et al (2014b) Disease risk factors identified through shared genetic architecture and electronic medical records. *Sci Transl Med* 6:234ra57. <https://doi.org/10.1126/scitranslmed.3007191>
- Pendergrass SA et al (2013) Phenome-wide association study (PheWAS) for detection of pleiotropy within the population architecture using genomics and epidemiology (PAGE) network. *PLoS Genet* 9:e1003087. <https://doi.org/10.1371/journal.pgen.1003087>
- Pickrell JK, Berisa T, Liu JZ, Segurel L, Tung JY, Hinds DA (2016) Detection and interpretation of shared genetic influences on 42 human traits. *Nat Genet* 48:709–717. <https://doi.org/10.1038/ng.3570>
- Rzhetsky A, Wajngurt D, Park N, Zheng T (2007) Probing genetic overlap among complex human phenotypes. *Proc Natl Acad Sci USA* 104:11694–11699. <https://doi.org/10.1073/pnas.0704820104>

- Sivakumaran S et al (2011) Abundant pleiotropy in human complex diseases and traits. *Am J Hum Genet* 89:607–618. <https://doi.org/10.1016/j.ajhg.2011.10.004>
- Solovieff N, Cotsapas C, Lee PH, Purcell SM, Smoller JW (2013) Pleiotropy in complex traits: challenges and strategies. *Nat Rev Genet* 14:483–495. <https://doi.org/10.1038/nrg3461>
- Stearns FW (2010) One hundred years of pleiotropy: a retrospective. *Genetics* 186:767–773. <https://doi.org/10.1534/genetics.110.122549>
- Tyler AL, Asselbergs FW, Williams SM, Moore JH (2009) Shadows of complexity: what biological networks reveal about epistasis and pleiotropy. *BioEssays* 31:220–227. <https://doi.org/10.1002/bies.200800022>
- Visscher PM, Yang J (2016) A plethora of pleiotropy across complex traits. *Nat Genet* 48:707–708. <https://doi.org/10.1038/ng.3604>
- Visscher PM, Brown MA, McCarthy MI, Yang J (2012) Five years of GWAS discovery. *Am J Hum Genet* 90:7–24. <https://doi.org/10.1016/j.ajhg.2011.11.029>
- Wagner GP, Zhang J (2011) The pleiotropic structure of the genotype-phenotype map: the evolvability of complex organisms. *Nat Rev Genet* 12:204–213. <https://doi.org/10.1038/nrg2949>
- Wagner GP, Kenney-Hunt JP, Pavlicev M, Peck JR, Waxman D, Cheverud JM (2008) Pleiotropic scaling of gene effects and the ‘cost of complexity’. *Nature* 452:470–472. <https://doi.org/10.1038/nature06756>
- Wang Z, Liao BY, Zhang JZ (2010) Genomic patterns of pleiotropy and the evolution of complexity. *Proc Natl Acad Sci USA* 107:18034–18039. <https://doi.org/10.1073/pnas.1004666107>
- Welter D et al (2014) The NHGRI GWAS catalog, a curated resource of SNP-trait associations. *Nucleic Acids Res* 42:D1001–D1006. <https://doi.org/10.1093/nar/gkt1229>
- Yang C, Li C, Wang Q, Chung DJ, Zhao HY (2015) Implications of pleiotropy: challenges and opportunities for mining Big Data in biomedicine. *Front Genet* 6:229. <https://doi.org/10.3389/fgene.2015.00229>