

Phenotype Similarity Regression for Identifying the Genetic Determinants of Rare Diseases

Daniel Greene,^{1,2} NIHR BioResource, Sylvia Richardson,^{2,3} and Ernest Turro^{1,2,3,*}

Rare genetic disorders, which can now be studied systematically with affordable genome sequencing, are often caused by high-penetrance rare variants. Such disorders are often heterogeneous and characterized by abnormalities spanning multiple organ systems ascertained with variable clinical precision. Existing methods for identifying genes with variants responsible for rare diseases summarize phenotypes with unstructured binary or quantitative variables. The Human Phenotype Ontology (HPO) allows composite phenotypes to be represented systematically but association methods accounting for the ontological relationship between HPO terms do not exist. We present a Bayesian method to model the association between an HPO-coded patient phenotype and genotype. Our method estimates the probability of an association together with an HPO-coded phenotype characteristic of the disease. We thus formalize a clinical approach to phenotyping that is lacking in standard regression techniques for rare disease research. We demonstrate the power of our method by uncovering a number of true associations in a large collection of genome-sequenced and HPO-coded cases with rare diseases.

Introduction

There is widespread interest in the study of rare diseases as a way of understanding the genetic architecture of biological processes. Consequently, tens of thousands of subjects are being phenotyped extensively and enrolled to genome-sequencing studies worldwide. To discover the cause of disease, these subjects would ideally be grouped a priori into clusters with a shared (though unknown) genetic etiology, but this is often hindered by extensive phenotypic and genetic heterogeneity (see [Web Resources](#) and examples^{1–9}). Rare variant association tests, even those accounting for some degree of genetic heterogeneity, typically summarize the clinical manifestations of a disease with a single variable,¹⁰ which can limit power when multiple phenotypic traits contain complementary information about the same causal genotype. Methods for modeling pleiotropy have proven successful in the context of genome-wide association studies^{11,12} but they are ill suited for rare disease studies in which the phenotype data are typically of mixed type and collected with variable detail and completeness.

The Human Phenotype Ontology (HPO)¹³ addresses the need for a standardized vocabulary for rare disease phenotypes and is being used to code patients in several large international projects^{14–16} (see also [Web Resources](#)). The HPO is a directed acyclic graph representing more than 10,000 phenotypic abnormalities in which the nodes (HPO terms) are connected to each other through “is-a” relations, represented as edges. The HPO was created with the support of experts in many areas of medicine to accommodate coding of phenotypic data derived from diverse sources, such as laboratory assays, images, graphs, and clinical interpretation. Methods exist that compare patient HPO data with HPO-coded profiles corresponding to

known diseases for the purpose of differential diagnosis.^{17,18} The HPO-coded profiles can be supplemented with functional gene-specific information to prioritize genes.^{19,20} If genotype data are available, these and other methods^{21,22} can be used to prioritize variants and potentially to suggest new causes of disease.^{19,20,23} However, the existing approaches do not share information between individually coded patients and as such are not statistical association methods.

Here, we present a regression-based method for discovering associations between arbitrarily diverse sets of HPO-coded phenotypes and genotypes at rare variant sites. To overcome the difficulty of modeling sparse and ontologically structured phenotype data, we treat the HPO-coded phenotypes of the subjects as the explanatory variables and their corresponding genotypes as the response. This is an example of “inverse regression” and is adequate in our setting because we are not interested in interpreting the regression coefficients per se but only in evaluating the probability of association. We define a subject’s “genotype” γ as a binary label that can take on the values “rare” (1) or “common” (0) according to a pre-specified function of the genetic data. For example, we could define the label “rare genotype” to mean that there is at least one rare variant in a particular gene (dominant inheritance) or at least two rare variants in a particular gene (recessive inheritance).

Our method then seeks to compare two models for the data, indexed by γ . Under the baseline model ($\gamma = 0$), the probability of observing the rare genotype is the same for each case. Under the alternate model ($\gamma = 1$), the probability of observing the rare genotype depends on the “phenotypic similarity” S (to be defined later) of the case to a latent *characteristic* HPO phenotype ϕ .

¹Department of Haematology, University of Cambridge, NHS Blood and Transplant, Cambridge Biomedical Campus, Cambridge CB2 0PT, UK; ²Medical Research Council Biostatistics Unit, Cambridge Biomedical Campus, Cambridge CB2 0SR, UK

³These authors contributed equally to this work

*Correspondence: et341@cam.ac.uk

<http://dx.doi.org/10.1016/j.ajhg.2016.01.008>. ©2016 The Authors

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Figure 1. Example HPO Coding of a Subject with Wiskott-Aldrich Syndrome

The nodes in blue imply the presence of the more general ancestral phenotypes depicted as gray nodes. No blue node has a directed path to any other, which means that the blue nodes comprise a *minimal* set of HPO terms. The graph has been simplified by removing nodes that link together only two other nodes.

We adopt a Bayesian inference framework, where the model selection indicator γ and characteristic phenotype ϕ are estimated through their posterior distributions. Of particular interest is the posterior mean of γ , which represents the probability that $\gamma = 1$, thus indicating the strength of evidence for an association.

A crucial element of our approach is the construction of an appropriate function for quantifying the semantic similarity of the characteristic phenotype ϕ to the phenotypes of the subjects. The choice of function is motivated by the need to optimally discriminate between subjects having clinical features that are pertinent to a disorder from those having overlapping or unrelated phenotypes due to a different disorder. To achieve this, we have chosen a function that accounts for the ontological structure of the HPO and induces a parsimonious characteristic phenotype: it selects the required terms to distinguish patient groups while avoiding overfitting and is robust to coding of patients with spurious or sporadic terms. Importantly, the function is flexible with respect to the phenotypic variability of disease and robust to the HPO coding practices of clinicians.

Our Bayesian approach provides a natural means of incorporating information from the scientific literature into our prior belief about the characteristic phenotype. In this work, we focus on gene-specific inference and up-weight the prior probability of characteristic phenotypes that are similar to clinical²³ and murine phenotypes²⁴ relevant to the gene.

We demonstrate the effectiveness of our method in identifying associations between genotype and phenotype through a simulation study, whereby phenotypes are simulated given genotypes in such a way as to emulate the effect of a hypothetical set of pathogenic variants. We go on to apply our inference procedure to a real dataset of more than 2,000 unrelated individuals enrolled to a variety of rare-disease sequencing studies under the auspices of the BRIDGE projects run by the NIH BioResource – Rare Diseases ([Web Resources](#)). We show that our method, implemented in the SimReg software package, can identify genes with rare variants responsible for a diverse set of pathologies in a single application and can estimate recognized disease phenotypes.

Material and Methods

Model Specification

We use a logistic regression framework to specify the two models under comparison:

$$\begin{aligned}
 y_i &\sim \text{Bernoulli}(p_i), \\
 \gamma = 0 : \quad &\log\left(\frac{p_i}{1-p_i}\right) = \alpha, \\
 \gamma = 1 : \quad &\log\left(\frac{p_i}{1-p_i}\right) = \alpha + \beta S(\phi, x_i).
 \end{aligned}
 \tag{Equation 1}$$

Here, y_1, \dots, y_N are the genotypes of the N subjects in the collection, where $y_i = 1$ if subject i possesses the rare genotype and $y_i = 0$ if subject i possesses the common genotype. x_1, \dots, x_N are the corresponding phenotypes of the subjects, where x_i comprises the minimal set of HPO terms required to describe the phenotypic abnormalities of subject i . Loosely speaking, a set of terms is minimal if it describes a patient's phenotype without redundancy (e.g., it does not include both "Abnormal bleeding" and "Joint hemorrhage"). More formally, a set of terms is said to be minimal if and only if it lacks elements implied by other terms in the set through directed edges in the HPO. The terms highlighted blue in [Figure 1](#) comprise such a set, because there is no directed path between any pair of blue nodes.

The term $S(\phi, x_i)$ denotes a chosen measure of phenotypic similarity between the characteristic phenotype and subject i 's phenotype. Note that our response and predictor are inverted compared to classical regression methods to avoid having to treat sparse and structured HPO data as the response. Under the baseline model, the intercept α is the global rate of rare genotypes. Under the alternate model, there is an additional parameter β , which is strictly positive and captures the effect of a unit increase in phenotypic similarity to the characteristic phenotype ϕ on the log odds of having the rare genotype. Thus, the probability that $\gamma = 1$ is greater in expectation when $S(\phi, x_i)$ is larger if $y_i = 1$ than if $y_i = 0$.

Similarity Measure

Our chosen similarity measure S is built with consideration for (1) quantification of the similarity of terms, (2) quantification of the similarity of a patient phenotype x_i to the characteristic phenotype ϕ , and (3) flexible transformation of the similarity between phenotypes.

Consistent with the ontological literature, we base our measure for the similarity of terms on the information content (IC) of each individual term,

$$\text{IC}(t) = -\log(\text{frequency}(t)),$$

where the frequency of term t can be derived from its appearance in the case collection, including instances in which this is implied by the presence of more specific terms in the ontology.

We use Lin's²⁵ similarity function to compare two different terms:

$$s(t_1, t_2) = \frac{2 \times \max_{t \in \text{anc}(t_1) \cap \text{anc}(t_2)} \text{IC}(t)}{\text{IC}(t_1) + \text{IC}(t_2)},$$

where $\text{anc}(t)$ denotes the union of term t and its ancestral nodes in the HPO graph. For example, for the hypothetical subject shown in Figure 1, the expression $\max_{t \in \text{anc}(t_1) \cap \text{anc}(t_2)} \text{IC}(t)$ if t_1 were "Thrombocytopenia" and t_2 were "Joint hemorrhage" would correspond to the IC of "Abnormality of blood and blood-forming tissues." Because terms cannot have a higher IC than their descendants, the similarity s between two terms can range between zero and one. Next, we consider asymmetric measures of similarity between a case phenotype and ϕ inspired by the best-match-average (BMA) function,¹⁷ which computes the best match for each term and takes the mean:

$$S_\phi(\phi \rightarrow x_i) = \frac{1}{|\phi|} \sum_{t_\phi \in \phi} \max_{t_x \in x_i} s(t_\phi, t_x) \mathbf{1}_{t_\phi \in \text{anc}(t_x)},$$

$$S_x(x_i \rightarrow \phi) = \frac{1}{|x_i|} \sum_{t_x \in x_i} \max_{t_\phi \in \phi} s(t_x, t_\phi) \mathbf{1}_{t_\phi \in \text{anc}(t_x)}.$$

The standard BMA function does not include the indicator variable above, which evaluates to 1 only if the node in ϕ is among the ancestors of the node in x_i . We prefer to include this restriction, which penalizes similarity to ϕ when it includes over-specific terms, in order to concentrate the posterior weight of ϕ preferentially on nodes that are present among the subjects.

The presence of a term in ϕ that is absent from x_i has the effect of lowering S_ϕ , whereas the presence of a term in x_i that is absent from ϕ has the effect of lowering S_x . Summation of two asymmetric similarities, as used in BMA, would allow reasonably high overall similarities to be obtained even when one of the two asymmetric similarities is close to zero. We prefer to multiply rather than add up the two similarity measures to obtain an expression for the overall similarity function used in Equation 1 because it ensures that the overall similarity can be high only when there is a high asymmetric similarity in both directions. However, because the values of S_x and S_ϕ are influenced by factors such as how frequent terms are in the reference database (which affects nodal IC) and the structure of the HPO graph, there is no guarantee that a linear function of their product optimally distinguishes subjects with objectively distinct clinical features. To ensure the model is robust to the choice of S , we allow modulation of the shapes of the similarity parameters, S_ϕ and S_x , through transformations f and g , respectively. A reasonable choice for f and g is the beta cumulative distribution function (CDF), because it maps $[0,1]$ to $[0,1]$ monotonically and allows a wide variety of shapes:

$$f(z, a_f, b_f) = I_z(a_f, b_f),$$

$$g(z, a_g, b_g) = I_z(a_g, b_g),$$

where I_z is the regularized incomplete beta function (see [Supplemental Note](#)) and a_f, a_g, b_f , and b_g are unknown parameters to be estimated.

Finally, the overall similarity function is given by

$$S(\phi, x_i) = f(S_\phi(\phi \rightarrow x_i), a_f, b_f) \cdot g(S_x(x_i \rightarrow \phi), a_g, b_g). \quad (\text{Equation 2})$$

Priors

We propose the following prior distributions for the model indicator and the regression parameters:

$$\gamma \sim \text{Bernoulli}(\pi),$$

$$\alpha \sim \text{Normal}(\text{mean} = 0, \text{sd} = 5),$$

$$\log \beta \sim \text{Normal}(\text{mean} = 2, \text{sd} = 1).$$

The value of π indicates how likely we believe a priori that there is a true association. All of the analyses in this paper assume $\pi = 0.05$. We place a vague prior on α around 0. Additionally, we include an offset on α by a constant \hat{h}_i for each individual that can take into account batch effects and factors affecting the background rate of rare genotypes (not shown in Equation 1 for clarity of exposition, see [Supplemental Note](#)). The prior distribution on β is positive because the probability of $y_i = 1$ increases with $S(\phi, x_i)$, given $\gamma = 1$. The prior variance of β allows for a wide range of effect sizes given the range of S . The priors on the beta CDF transformations are discussed in the [Supplemental Note](#). In brief, the choice of prior for f favors parsimonious characteristic phenotypes and the prior for g allows for an indeterminate number of nodes appearing sporadically among patients.

By default, our prior distribution on the characteristic phenotype ϕ places a uniform prior probability on all minimal sets of terms of size less than or equal to k . We choose $k = 3$ on the grounds that three nodes should adequately distinguish between the primary features of most rare diseases. If γ is set based on variants in a particular feature, such as a gene, then our prior can up-weight HPO phenotypes comprising terms annotated to that feature on the basis of reports in the scientific literature. Thus, the prior on ϕ is given by

$$\mathbb{P}(\phi) = \begin{cases} \frac{1}{|\Phi^{(k)}|} & \text{No literature phenotype} \\ \frac{S'(M \rightarrow \phi)}{\sum_{\psi \in \Phi^{(k)}} S'(M \rightarrow \psi)} & \text{Literature phenotype } M \end{cases} \quad (\text{Equation 3})$$

where $\Phi^{(k)}$ denotes the set of all minimal sets of up to k HPO terms and S' is an unstandardized similarity function (see [Supplemental Note](#)). In practice, the literature phenotype could be obtained from OMIM or from the Mouse Genome Informatics (MGI) database²⁴ after mapping murine phenotypes coded using the Mammalian Phenotype Ontology²⁶ to HPO terms through a cross-species phenotype ontology.²⁷

Inference

We perform model comparison using the Markov chain Monte Carlo (MCMC)-based method of Carlin and Chib.²⁸ We sample the model selection parameter γ from its full conditional distribution while the remaining parameters are sampled using the




Figure 2. Results of Inference on Simulated Data

Phenotype data were simulated using three levels of expressivity r of the disease terms. The plots within each panel correspond to different frequencies $\sum_i y_i$ of the rare genotype. In each plot, the red dots mark the estimated posterior mean of γ for 64 datasets simulated under $\tilde{\gamma} = 1$ and the gray dots show an equivalent set of estimates for datasets simulated under $\tilde{\gamma} = 0$ (i.e., whereby phenotypes for subjects having $y_i = 1$ are sampled from the same distribution as for subjects having $y_i = 0$). The position of points on the x axis within a plot is arbitrary.

Metropolis-Hastings algorithm or from a pseudoprior distribution, depending on the value of γ at each iteration.

It is not straightforward to sample from the space of minimal sets $\Phi^{(k)}$ when $\gamma = 1$ because not all possible HPO term combinations comprise such a minimal set. To overcome this difficulty, we propose an unrestricted vector of k HPO terms $\tilde{\phi}$ and then derive the associated underlying phenotype ϕ by applying a mapping function v . We therefore need to impose a prior distribution on the unrestricted space which is compatible with the desired prior for ϕ (Equation 3) on the restricted space. To be precise, the prior on $\tilde{\phi}$ is given by:

$$\mathbb{P}(\tilde{\phi}) = \frac{\mathbb{P}(v(\tilde{\phi}))}{|\{\tilde{\phi}' \in H^k : v(\tilde{\phi}') = v(\tilde{\phi})\}|},$$

where H^k is the space of all vectors of k HPO terms and v maps an arbitrary such vector of terms to its corresponding minimal set. The denominator accounts for the number of unrestricted vectors that map to the same minimal set. For further details on the method used to calculate $\mathbb{P}(\phi)$ and $\mathbb{P}(\tilde{\phi})$, the MCMC algorithm, and the tuning of the pseudopriors, refer to the [Supplemental Note](#).

Results

Simulation Study

We assessed the performance of SimReg by analyzing datasets generated under two scenarios, labeled by $\tilde{\gamma}$. Under $\tilde{\gamma} = 1$, the HPO phenotypes $x_{1,\dots,N}$ were simulated conditional on the genotypes $y_{1,\dots,N}$ of N individuals whereas under $\tilde{\gamma} = 0$ they were simulated independently of the genotypes. When $\tilde{\gamma} = 1$, phenotypes for all subjects having $y_i = 1$ were formed by selecting terms from an arbitrarily chosen disease template (“Decreased mean platelet

volume,” “Thrombocytopenia,” and “Autism”). Each term was selected with a pre-specified probability r , termed “expressivity,” and m further noise terms drawn at random from a set of approximately 1,000 HPO terms were appended, where $m \sim \text{Poisson}(\lambda = 5)$. The set of terms from which the noise terms were drawn was created by selecting 200 HPO terms at random, taking the union with the ancestral terms. Phenotypes for subjects having $y_i = 0$ were drawn at random using terms from the above set with $\lambda = 8$ and then mapped to minimal sets. When $\tilde{\gamma} = 0$, all phenotypes were sampled from the noise term set with $\lambda = 8$. This ensures that on average individuals have approximately 8 terms, irrespective of y_i and $\tilde{\gamma}$. The simulation was performed with the set of disease template terms and set of noise terms fixed but with different numbers of individuals carrying the rare genotype ($\sum_i y_i \in \{2, 4, 6, 8, 10, 20\}$ out of $N = 1,000$) and varied levels of expressivity $r \in \{1/3, 2/3, 1\}$. The low expressivity set-ups capture situations in which a fraction of the individuals having a rare genotype can be considered to carry a neutral variant with respect to the disease in question because they have none of the template terms. For the same reason, they capture scenarios of incomplete penetrance of a subset of the underlying rare variants. Furthermore, a degree of genetic heterogeneity is built into our simulation setup, because there is a non-zero probability of a template phenotype term being randomly allocated to an individual with the common genotype.

The results of repeating the simulation 64 times for each value of $\tilde{\gamma}$ and combination of r and $\sum_i y_i$, depicted in [Figure 2](#), show that power to detect a true association, as assessed by the posterior mean of γ , increases with the expressivity of the disease terms r and also with the frequency of the rare genotype in the study sample $\sum_i y_i$ (red dots). Under $\tilde{\gamma} = 0$, the posterior mean of γ remains very close to zero in all circumstances (gray dots). Specifically, we find that 2, 6, and 20 cases out of 1,000 subjects are

Table 1. Studies from which Genetic and Phenotypic Data Were Obtained

Study	Phenotype	Unrelated Subjects	Known Genes
Bleeding and Platelet Disorders (BPD)	detailed patient-specific HPO terms	709	74
Primary ImmunoDeficiency (PID)	Abnormality of the immune system (HP:0002715)	201	131
Pulmonary Arterial Hypertension (PAH)	Pulmonary hypertension (HP:0002092)	422	9
Specialist Pathology Evaluating Exomes in Diagnostics (SPEED)	Retinal dystrophy (HP:0000556)	384	241
	Abnormality of the nervous system (HP:0000707)	215	689
	Abnormality of the nervous system and Retinal dystrophy (HP:0000707, HP:0000556)	7	
	Phenotypic abnormality (HP:0000118)	107	

Note that the SPEED project has a branch dealing with retinal dystrophy and another branch dealing with abnormalities of the nervous system and that 7 individuals are included in both branches. In addition, 107 subjects could not be assigned to a specific sub-project at the time of writing due to lack of information and we assigned them a single abstract HPO term “Phenotypic abnormality” (HP:0000118).

sufficient to obtain perfect or near-perfect discrimination between the two models when the expressivity is 1, 2/3, and 1/3, respectively. When the number of subjects with the rare genotype is equal to 6 and the expressivity is 2/3, which implies that any two individuals with the rare genotype have only a 0.17 chance of having exactly the same template terms, our method can achieve a positive predictive value of 1, even when the negative predictive value is as high as 0.95, by thresholding at $\mathbb{P}(\gamma = 1 | y) \geq 0.25$. Under this set-up, we expect 1.78 of the 6 individuals with the rare genotype to have none of the template terms at all, which indicates that the method has some resilience to the presence of $\gamma_i = 1$ induced by neutral rather than pathogenic variants. In order to assess the specificity of the method more accurately, we have simulated 20,000 datasets under the scenario in which there is no association and $\sum_i \gamma_i = 6$ and found that only 7 datasets yield $\mathbb{P}(\gamma = 1 | y) > 0.25$, which equates to a specificity of 99.97% for this chosen cut-off (Supplemental Note). We have also extended our simulation study to include a variable controlling genetic heterogeneity, whereby many individuals are drawn from the same template but only a subset have the rare genotype. Power is maintained even in challenging scenarios in which there is substantial genetic heterogeneity and moderate phenotypic expressivity (Supplemental Note). Overall, the results of our simulation study show that our method produces accurate results even in the presence of significant phenotypic or genetic heterogeneity and low expressivity of the rare genotype’s characteristic terms. Because these are typical hallmarks of many rare disease studies, our evaluation substantiates the utility of our approach.

Results from Real Data

Our dataset comprises HPO phenotypes and corresponding variant call data for 2,045 unrelated individuals enrolled to a variety of rare-disease sequencing studies (Table 1). Detailed HPO data were available only for subjects enrolled to the Bleeding and Platelet Disorders (BPD) project.¹⁴ BPDs are a heterogeneous group of

diseases, including polysymptomatic examples, making them an interesting use-case for the modeling we present. For the other projects, only high-level HPO terms were used (Table 1). A set of genes within which variants are known to be implicated in each class of disorders was provided by BRIDGE collaborators to assess the performance of the model (Supplemental Note).

We used variant call data from 686 sequenced exomes and 1,359 sequenced whole genomes. To account for biases that might alter the baseline rate of rare genotypes (e.g., sequencing platform), we use a plug-in offset in the regression Equation 1, estimated a priori (see Supplemental Note). Variants were retained only if they were predicted to alter protein sequence and were either absent from ExAC (Web Resources) or had an allele frequency therein below 1/1,000 or 1/10,000 when a recessive or dominant mode of inheritance, respectively, was assumed in the analysis. Rare variants were aggregated within genes to account for genetic heterogeneity and increase power. We defined the binary genotypes y based on three different aggregation approaches corresponding to the following hypothetical modes of inheritance: (1) dominant, i.e., presence of at least one rare allele; (2) recessive, i.e., presence of at least two rare alleles; or (3) high-impact dominant, i.e., presence of at least one rare allele predicted²⁹ to introduce a splice site aberration, frameshift, start loss, or stop gain.

ACTN1 as Exemplar Gene

We now describe the properties of SimReg’s output by focusing on a gene, *ACTN1* (MIM: 102575), that has recently been reported to harbor rare variants responsible for reduced platelet number and increased platelet size (macrothrombocytopenia).³⁰ We note that data for *ACTN1* were used to inform and motivate our choice for the similarity measure given in Equation 2 (Supplemental Note). Once learnt on the *ACTN1* data, this choice has then been used universally for all genes. We observe strong evidence that the rare genotype for *ACTN1* is associated with similarity to a characteristic phenotype ($\mathbb{P}(\gamma = 1 | y) = 1$), as expected. The estimated characteristic

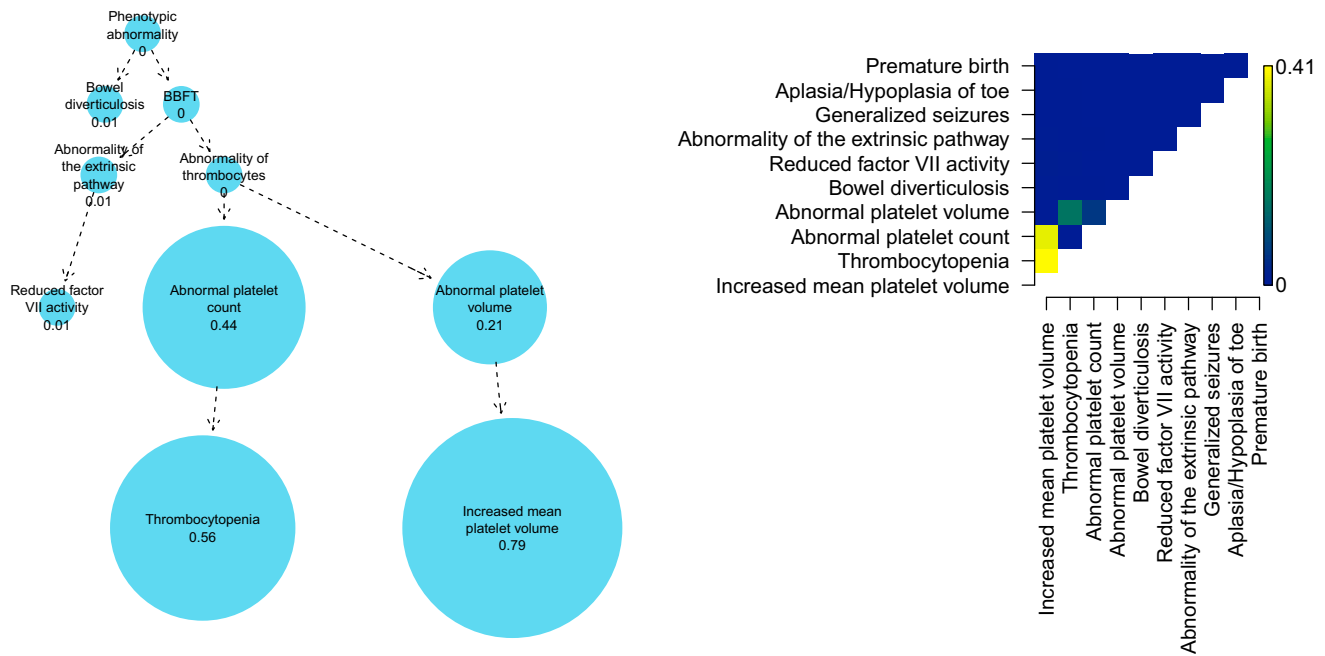


Figure 3. Results for ACTN1

The panels show results obtained by applying the SimReg method to phenotype data for all subjects and genotype data for *ACTN1*. There were 43 individuals in our dataset coded with the rare genotype for this gene, of which 22 were coded with “Thrombocytopenia” and “Increased mean platelet volume.” The graph shows the estimated probabilities of inclusion of individual terms in ϕ (only the seven terms with the highest probabilities of inclusion and their ancestors are shown). The acronym “BBFT” refers to “Abnormality of blood and blood-forming tissues.” The heatmap shows the estimated probabilities of pairs of terms co-occurring in ϕ , for pairs composed from the ten most frequently included individual terms.

phenotype focuses primarily on phenotypes that include “Thrombocytopenia” and “Increased mean platelet volume” (Figure 3), which together correspond to macrothrombocytopenia. The slightly more general terms “Abnormal platelet count” and “Abnormal platelet volume” also have substantial marginal posterior weight whereas the rest of the nodes in the HPO have a marginal posterior probability of inclusion less than 0.02. As can be seen in a two-dimensional matrix of the marginal posterior on pairs of terms (Figure 3), there is a high degree of co-occurrence of the two primary terms representing the *ACTN1*-related phenotype, which implies that they are not alternatives but rather complements that together produce a good model fit.

DIAPH1 and RASGRP2

Under the high-impact dominant mode of inheritance described above, one of the genes with the highest estimated value of γ that also has a BPD-like inferred phenotype is *DIAPH1* (MIM: 602121) ($\gamma = 0.87$). We recently showed, through an application of our similarity regression approach, that the introduction of a premature stop codon present in two unrelated individuals in the BPD project truncates *DIAPH1*’s 3’ auto-inhibitory domain and causes macrothrombocytopenia, hearing loss, and mild bleeding.³¹ As shown in Figure 4 (left), the salient terms in ϕ relate to hearing impairment and abnormality of blood and blood-forming tissues, with the latter driven mainly by thrombocytopenia and bleeding. The high

posterior estimate of γ was obtained in part because a sensorineural hearing loss phenotype had previously been reported in the literature,³² which up-weighted hearing abnormality terms in the prior for ϕ (Table 2). However, even without using an informative prior on ϕ , a high posterior probability of an association ($\gamma = 0.59$) could be found for *DIAPH1*.

RASGRP2 (MIM: 605577) was recently implicated in a new form of Glanzmann’s-like thrombasthenia (MIM: 273800) based on data from a single pedigree.³³ Glanzmann’s is characterized by impaired platelet aggregation, leading to excessive bleeding. Under a recessive mode of inheritance, our similarity regression successfully detects an association ($\gamma = 0.75$) for *RASGRP2* and estimates a characteristic phenotype concentrated around “Abnormal platelet aggregation” (Figure 4). It is characteristic of Glanzmann’s that platelet aggregation is impaired in response to multiple agonists because their common downstream effect—the binding of platelets to fibrinogen—is impeded by the presence of reduced numbers of fibrinogen receptors. Here we also observe this phenomenon but only collagen-induced platelet aggregation carries significant weight in the characteristic phenotype because it is the only specific aggregation term that is shared by all the cases of this recently discovered disorder. There is also a very low probability of inclusion of two rare terms that are not related to the disease—“Atypical scarring of skin” and “Intracranial meningioma”—because of a chance comorbidity in one of the affected cases.

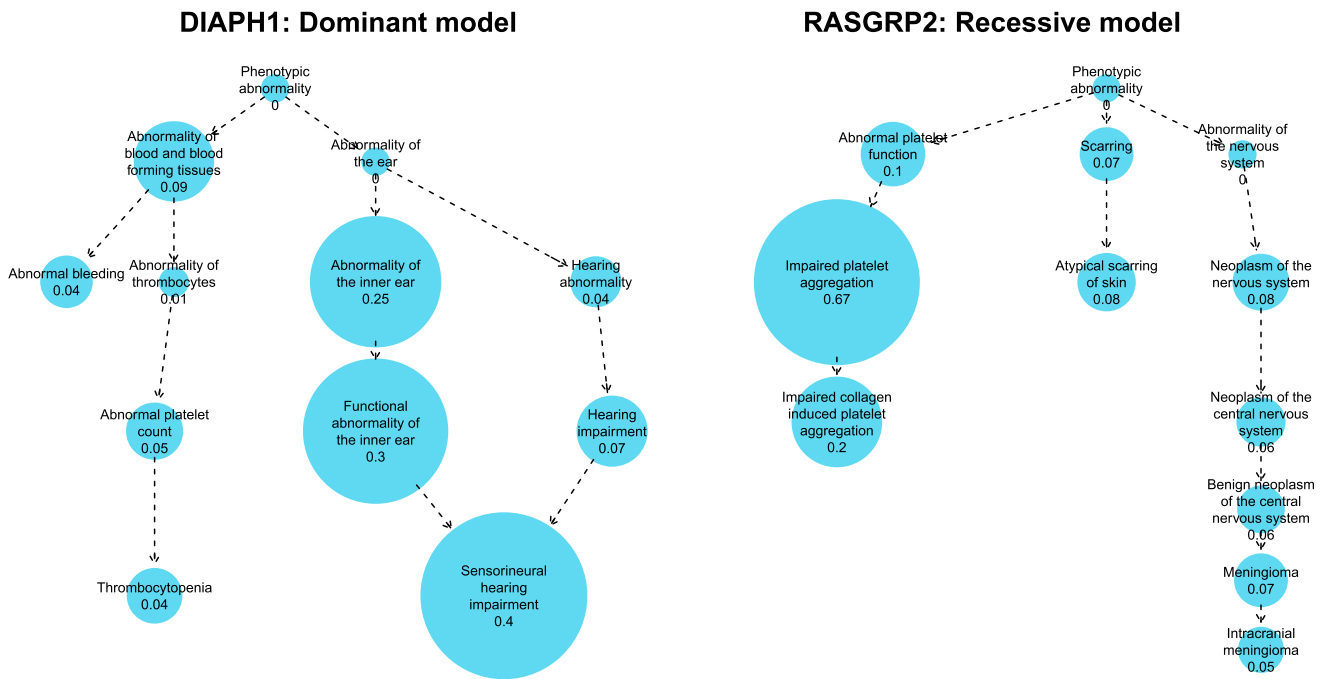


Figure 4. Results for *DIAPH1* and *RASGRP2*

Estimated posterior probabilities of individual terms being included in the characteristic phenotype ϕ using phenotype data for all subjects and variant data for *DIAPH1* ($\sum_i \gamma_i = 2$) encoded under a high-impact dominant model and *RASGRP2* ($\sum_i \gamma_i = 7$) encoded under a recessive model. The ten terms with the highest marginal posterior probability are shown. The estimated posterior probability that $\gamma = 1$ is equal to 0.872 and 0.750 for *DIAPH1* and *RASGRP2*, respectively.

Overall Results

Finally, we turn our attention to the overall results of applying the inference procedure to data for all genes under the three modes of inheritance considered, subject to $\sum_i \gamma_i \geq 2$. In total, we applied the inference to 19,573, 3,134 and 9,733 genes for the dominant, recessive, and high-impact dominant modes of inheritance, respectively. The estimates of $\mathbb{P}(\gamma = 1 | \gamma)$ are shown as vertical density plots in Figure 5. For the majority of genes (65%), $\mathbb{P}(\gamma = 1 | \gamma) < \mathbb{P}(\gamma = 1) = 0.05$, which implies that no characteristic phenotype can be found that helps distinguish carriers of the rare genotype from other subjects. This result is consistent with the expectation that variants in only a small proportion of genes are implicated in these rare diseases and indicates that specificity is largely controlled.

Strikingly, under all three assumed modes of inheritance, most of the highly confident results (i.e., the genes for which the estimates of $\mathbb{P}(\gamma = 1 | \gamma)$ are close to one) are for genes known to be relevant to the pathologies of the patients (indicated by red labels in Figure 5). In all but one case (*KIF1A* [MIM: 601255]), where a gene had $\mathbb{P}(\gamma = 1 | \gamma) > 0.25$ and was in one of the projects' set of known genes, a characteristic phenotype similar to the known phenotype was inferred (Table 2). Above a threshold of $\mathbb{P}(\gamma = 1 | \gamma) = 0.25$, there was a significant enrichment for known genes (Fisher exact test $p = 2.39 \times 10^{-4}$, 1.98×10^{-4} , and 2.23×10^{-7} for the dominant, recessive, and high-impact dominant modes of in-

heritance, respectively). Some of the inferred known genes are highlighted more than once across the three modes of inheritance in Figure 5 because there is power to detect the association even when the mode of inheritance is misspecified. For example, *RASGRP2*-related Glanzmann's is recessive, yet $\mathbb{P}(\gamma = 1 | \gamma) > 0.25$ even if a high-impact dominant mode of inheritance is assumed.

The black dashes in Figure 5 correspond to unknown genes for which the inferred $\mathbb{P}(\gamma = 1)$ is greater than 0.25, of which there were 8, 1, and 5 found for the dominant, recessive, and high-impact dominant model of inheritance, respectively. These candidates are genes with potentially novel roles in disease and are being actively explored.

Discussion

We have described a method for identifying the genetic determinants of rare diseases that does not require the disease phenotype to be specified a priori. The method uncovers associations between rare genotypes and the similarities between subject phenotypes and a latent characteristic phenotype. Throughout this paper, rare variants have been aggregated within genes according to a hypothesized mode of inheritance in order to define presence or absence of a rare genotype. However, the unit of analysis could be a set of interacting domains or any other arbitrary genomic grouping. During final review of this work, a prioritization

Table 2. Known Genes for which $\mathbb{P}(\gamma=1 y) > 0.25$ and the Inferred Phenotype Was Compatible with the Known Disorder					
Gene	MIM No.	Mode of Inheritance	Known Disorder	$\mathbb{P}(\gamma=1 y)$	Highest Marginal Posterior Probability Terms in ϕ
<i>ACTN1</i>	102575	dominant	bleeding and platelet disorder	1.00	increased mean platelet volume (0.79), thrombocytopenia (0.56), platelet count (0.44)
<i>BMP2</i>	600799	dominant	pulmonary arterial hypertension	1.00	pulmonary hypertension (0.34), elevated pulmonary artery pressure (0.31), pulmonary artery (0.11)
<i>ABCA4</i>	601691	recessive	retinal dystrophy	0.99	retinal dystrophy (0.22), retina (0.22), fundus (0.16)
<i>USH2A</i>	608400	recessive	retinal dystrophy	0.99	retina (0.23), retinal dystrophy (0.2), fundus (0.17)
<i>CRB1</i>	604210	recessive	retinal dystrophy	0.97	retinal dystrophy (0.21), retina (0.18), fundus (0.18)
<i>F11</i>	264900	high-impact dominant	bleeding and platelet disorder	0.95	reduced factor XI activity (0.89), intrinsic pathway (0.11), platelet aggregation (0.07)
<i>RASGRP2</i>	605577	recessive	bleeding and platelet disorder	0.75	platelet aggregation (0.67), collagen-induced platelet aggregation (0.2), platelet function (0.1)
<i>EYS</i>	612424	high-impact dominant	retinal dystrophy	0.70	retinal dystrophy (0.2), retina (0.17), fundus (0.14)
<i>F7</i>	613878	high-impact dominant	bleeding and platelet disorder	0.68	extrinsic pathway (0.5), reduced factor vii activity (0.46), white hair (0.1)
<i>RPGR</i>	312610	high-impact dominant	retinal dystrophy	0.42	retina (0.2), retinal dystrophy (0.17), posterior segment of the eye (0.16)

We display the mode of inheritance under which the association was found, the known disorder, the probability of association, and the top three HPO terms (shown in abbreviated form) in the inferred phenotypes. The marginal posterior probability of inclusion in the characteristic phenotype is shown in brackets next to each term. When an association was found under multiple modes of inheritance, only the true mode is shown. Note that the inferred phenotypes are influenced by prior phenotypic information in the form of OMIM and MGI annotations.

procedure was proposed that combines a standard measure of strength of phenotypic clustering among individuals having two loss-of-function variants in a gene and the probability of the variants appearing in opposite haplotypes in an outbred population.³⁴ In contrast, our inference procedure is based on statistical principles and the

formulation of a model that is flexible with regards to phenotypic expressivity and genetic architecture and robust to noisy clinical coding and moderate genetic heterogeneity. Our Bayesian model naturally accounts for prior evidence of disease phenotypes associated with variants in particular genes by differentially weighting the

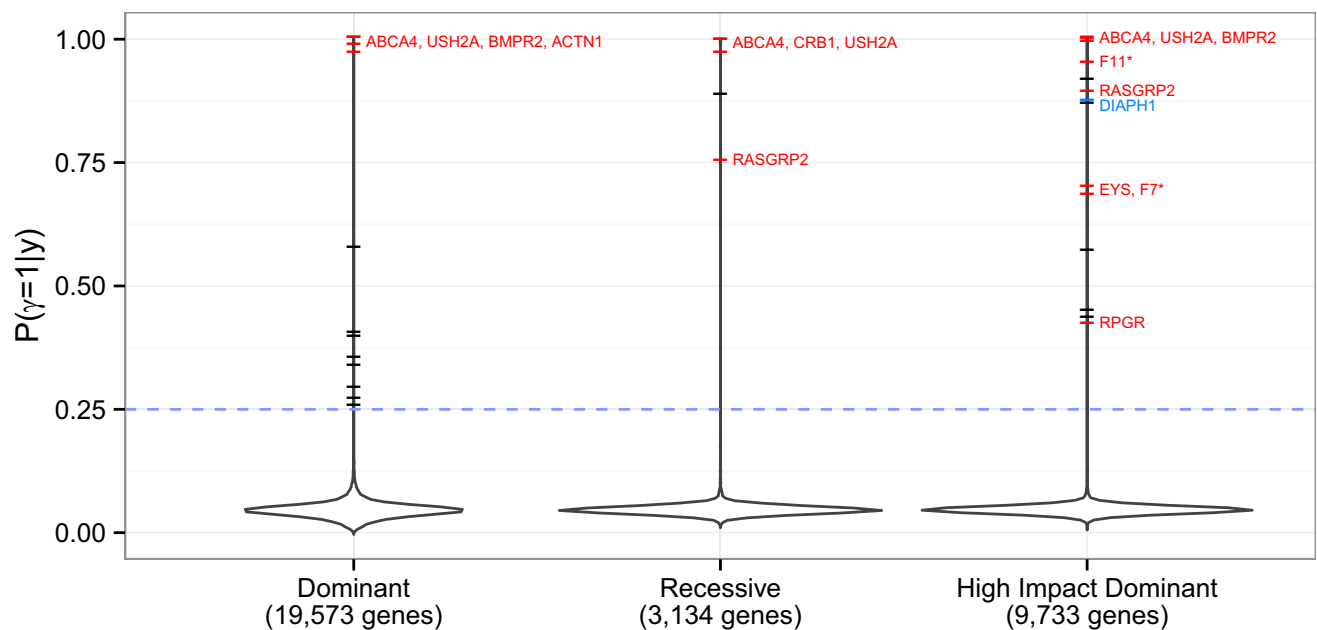


Figure 5. Overall Results

Distributions of the estimated posterior means of γ obtained by applying the SimReg method to each gene under three different modes of inheritance. The tails are truncated at the most extreme values. The dashes indicate values greater than 0.25. The known genes for the BRIDGE project disorders having $\mathbb{P}(\gamma = 1 | y) > 0.25$ and a compatible inferred phenotype are labeled and colored in red. An asterisk indicates that a posterior mean of γ greater than 0.25 was estimated only with the use of a prior on ϕ that was informed by the literature of human and murine heritable disorders.

prior probability of inclusion of HPO terms in the characteristic phenotype. Our finding that variants in *DIAPH1* can cause macrothrombocytopenia is an example of how this up-weighting can improve the inference.

The approach we have described is a natural and powerful way of modeling many rare disease phenotypes because it accounts for phenotypic abnormalities across all organ systems encoded with variable precision. Studies of syndromic diseases in particular can benefit from this way of uncovering associations. Our model can also be used for predicting the log odds of the rare genotype using solely phenotype data by means of a function implemented in our SimReg software. This could be used to aid diagnosis by indicating which of a patient's genes should be prioritized for sequencing based on his or her HPO terms. Finally, our regression approach might prove useful for performing inference using notions of similarity between terms in other ontologies where a binary response can be encoded.

Although our method improves significantly on modeling of phenotypic heterogeneity, our treatment of genetic heterogeneity can still be refined, because we currently rely on aggregation of genetic information into single binary variables. In the future we will explore improved modeling of genetic heterogeneity, in which the possibility of a mixture of pathogenic and neutral variants is accounted for explicitly. This would be applicable to genes in which different variants can cause drastically different clinical pathologies (e.g., *LMNA* [MIM: 150330]). Allele frequency, conservation, and functional information could also be used to modulate prior distributions.

In summary, our work represents an advancement in the statistical modeling of ontological heterogeneity that might prove useful at a time in which large collections of deeply phenotyped and sequenced cases are being assembled to uncover hitherto elusive causes of rare heterogeneous diseases.

Supplemental Data

Supplemental Data include Supplemental Note (Diagram Representing the $\gamma = 1$ model; Detailed Model Specification; Estimation of the Offset \hat{h}_i ; Prior on f and g ; Genetic Heterogeneity; Specificity; Inference using Markov Chain Monte Carlo; Calculation of Prior Probability for ϕ and $\hat{\phi}$; Ethics; SimReg Performance; and Lists of Known Genes for the BRIDGE Projects) and one table (listing additional members and collaborators of the NIH BioResource) and can be found with this article online at <http://dx.doi.org/10.1016/j.ajhg.2016.01.008>.

Acknowledgments

This work was supported by NIH award RG65966 (D.G. and E.T.) and the Medical Research Council programme grant MC_UP_0801/1 (D.G. and S.R.). The NIH BioResource – Rare Diseases projects were approved by Research Ethics Committees in the UK and appropriate national ethics authorities in non-UK enrollment centers (see [Supplemental Note](#)). We are grateful to Dr. William J. Astle for advice on the statistical model and for providing comments on

the manuscript. We are particularly thankful to the BPD project members for granting access to detailed HPO terms of subjects.

Received: August 25, 2015

Accepted: January 8, 2016

Published: February 25, 2016

Web Resources

The URLs for data presented herein are as follows:

Care for Rare, <http://care4rare.ca>

Comprehensive R Archive Network, <http://cran.r-project.org>

ExAC Browser, <http://exac.broadinstitute.org/>

Genomics England, <http://www.genomicsengland.co.uk>

HPO, <http://compbio.charite.de/jenkins/job/hpo/>

Mouse Genome Informatics, <http://www.informatics.jax.org/>

NIHR BioResource – Rare Diseases, <https://bioresource.nihr.ac.uk/rare-diseases/>

OMIM, <http://www.omim.org/>

RetNet – Retinal Information Network, <https://sph.uth.edu/retnet/home.htm>

Undiagnosed Diseases Network, <http://www.genome.gov/27550959>

References

1. Seri, M., Cusano, R., Gangarossa, S., Caridi, G., Bordo, D., Lo Nigro, C., Ghiggeri, G.M., Ravazzolo, R., Savino, M., Del Vecchio, M., et al.; The May-Hegglin/Fechtner Syndrome Consortium (2000). Mutations in *MYH9* result in the May-Hegglin anomaly, and Fechtner and Sebastian syndromes. *Nat. Genet.* 26, 103–105.
2. Murayama, S., Akiyama, M., Namba, H., Wada, Y., Ida, H., and Kunishima, S. (2013). Familial cases with *MYH9* disorders caused by *MYH9* S96L mutation. *Pediatr. Int.* 55, 102–104.
3. Feng, L., Seymour, A.B., Jiang, S., To, A., Peden, A.A., Novak, E.K., Zhen, L., Rusiniak, M.E., Eicher, E.M., Robinson, M.S., et al. (1999). The β 3A subunit gene (*Ap3b1*) of the AP-3 adaptor complex is altered in the mouse hypopigmentation mutant pearl, a model for Hermansky-Pudlak syndrome and night blindness. *Hum. Mol. Genet.* 8, 323–330.
4. Anikster, Y., Huizing, M., White, J., Shevchenko, Y.O., Fitzpatrick, D.L., Touchman, J.W., Compton, J.G., Bale, S.J., Swank, R.T., Gahl, W.A., and Toro, J.R. (2001). Mutation of a new gene causes a unique form of Hermansky-Pudlak syndrome in a genetic isolate of central Puerto Rico. *Nat. Genet.* 28, 376–380.
5. Suzuki, T., Li, W., Zhang, Q., Karim, A., Novak, E.K., Sviderskaya, E.V., Hill, S.P., Bennett, D.C., Levin, A.V., Nieuwenhuis, H.K., et al. (2002). Hermansky-Pudlak syndrome is caused by mutations in *HPS4*, the human homolog of the mouse light-ear gene. *Nat. Genet.* 30, 321–324.
6. Zhang, Q., Zhao, B., Li, W., Oiso, N., Novak, E.K., Rusiniak, M.E., Gautam, R., Chintala, S., O'Brien, E.P., Zhang, Y., et al. (2003). *Ru2* and *Ru* encode mouse orthologs of the genes mutated in human Hermansky-Pudlak syndrome types 5 and 6. *Nat. Genet.* 33, 145–153.
7. Morgan, N.V., Pasha, S., Johnson, C.A., Ainsworth, J.R., Eady, R.A., Dawood, B., McKeown, C., Trembath, R.C., Wilde, J., Watson, S.P., and Maher, E.R. (2006). A germline mutation in *BLOC1S3/reduced pigmentation* causes a novel variant of

- Hermansky-Pudlak syndrome (HPS8). *Am. J. Hum. Genet.* 78, 160–166.
8. Li, W., Zhang, Q., Oiso, N., Novak, E.K., Gautam, R., O'Brien, E.P., Tinsley, C.L., Blake, D.J., Spritz, R.A., Copeland, N.G., et al. (2003). Hermansky-Pudlak syndrome type 7 (HPS-7) results from mutant dysbindin, a member of the biogenesis of lysosome-related organelles complex 1 (BLOC-1). *Nat. Genet.* 35, 84–89.
 9. Cullinane, A.R., Curry, J.A., Carmona-Rivera, C., Summers, C.G., Ciccone, C., Cardillo, N.D., Dorward, H., Hess, R.A., White, J.G., Adams, D., et al. (2011). A BLOC-1 mutation screen reveals that *PLDN* is mutated in Hermansky-Pudlak Syndrome type 9. *Am. J. Hum. Genet.* 88, 778–787.
 10. Lee, S., Abecasis, G.R., Boehnke, M., and Lin, X. (2014). Rare-variant association analysis: study designs and statistical tests. *Am. J. Hum. Genet.* 95, 5–23.
 11. O'Reilly, P.F., Hoggart, C.J., Pomyen, Y., Calboli, F.C., Elliott, P., Jarvelin, M.R., and Coin, L.J. (2012). MultiPhen: joint model of multiple phenotypes can increase discovery in GWAS. *PLoS ONE* 7, e34861.
 12. Stephens, M. (2013). A unified framework for association analysis with multiple related phenotypes. *PLoS ONE* 8, e65245.
 13. Köhler, S., Doelken, S.C., Mungall, C.J., Bauer, S., Firth, H.V., Bailleul-Forestier, I., Black, G.C., Brown, D.L., Brudno, M., Campbell, J., et al. (2014). The Human Phenotype Ontology project: linking molecular biology and disease through phenotype data. *Nucleic Acids Res.* 42, D966–D974.
 14. Westbury, S.K., Turro, E., Greene, D., Lentaigne, C., Kelly, A.M., Bariana, T.K., Simeoni, I., Pillois, X., Attwood, A., Austin, S., et al.; BRIDGE-BPD Consortium (2015). Human phenotype ontology annotation and cluster analysis to unravel genetic defects in 707 cases with unexplained bleeding and platelet disorders. *Genome Med.* 7, 36.
 15. Fitzgerald, T.W., Gerety, S.S., Jones, W.D., van Kogelenberg, M., King, D.A., McRae, J., Morley, K.I., Parthiban, V., Al-Turki, S., Ambridge, K., et al.; Deciphering Developmental Disorders Study (2015). Large-scale discovery of novel genetic causes of developmental disorders. *Nature* 519, 223–228.
 16. Philippakis, A.A., Azzariti, D.R., Beltran, S., Brookes, A.J., Brownstein, C.A., Brudno, M., Brunner, H.G., Buske, O.J., Carey, K., Doll, C., et al. (2015). The Matchmaker Exchange: a platform for rare disease gene discovery. *Hum. Mutat.* 36, 915–921.
 17. Köhler, S., Schulz, M.H., Krawitz, P., Bauer, S., Dölken, S., Ott, C.E., Mundlos, C., Horn, D., Mundlos, S., and Robinson, P.N. (2009). Clinical diagnostics in human genetics with semantic similarity searches in ontologies. *Am. J. Hum. Genet.* 85, 457–464.
 18. Bauer, S., Köhler, S., Schulz, M.H., and Robinson, P.N. (2012). Bayesian ontology querying for accurate and noise-tolerant semantic searches. *Bioinformatics* 28, 2502–2508.
 19. Singleton, M.V., Guthery, S.L., Voelkerding, K.V., Chen, K., Kennedy, B., Margraf, R.L., Durtschi, J., Eilbeck, K., Reese, M.G., Jorde, L.B., et al. (2014). Phevor combines multiple biomedical ontologies for accurate identification of disease-causing alleles in single individuals and small nuclear families. *Am. J. Hum. Genet.* 94, 599–610.
 20. Yang, H., Robinson, P.N., and Wang, K. (2015). Phenolyzer: phenotype-based prioritization of candidate genes for human diseases. *Nat. Methods* 12, 841–843.
 21. Robinson, P.N., Köhler, S., Oellrich, A., Wang, K., Mungall, C.J., Lewis, S.E., Washington, N., Bauer, S., Seelow, D., Krawitz, P., et al.; Sanger Mouse Genetics Project (2014). Improved exome prioritization of disease genes through cross-species phenotype comparison. *Genome Res.* 24, 340–348.
 22. Zemojtel, T., Köhler, S., Mackenroth, L., Jäger, M., Hecht, J., Krawitz, P., Graul-Neumann, L., Doelken, S., Ehmke, N., Spielmann, M., et al. (2014). Effective diagnosis of genetic disease by computational phenotype analysis of the disease-associated genome. *Sci. Transl. Med.* 6, 252ra123.
 23. Javed, A., Agrawal, S., and Ng, P.C. (2014). Phen-Gen: combining phenotype and genotype to analyze rare disorders. *Nat. Methods* 11, 935–937.
 24. Blake, J.A., Bult, C.J., Eppig, J.T., Kadin, J.A., and Richardson, J.E.; Mouse Genome Database Group (2014). The Mouse Genome Database: integration of and access to knowledge about the laboratory mouse. *Nucleic Acids Res.* 42, D810–D817.
 25. Lin, D. (1998). An information-theoretic definition of similarity. In Shavlik, J.W., ed., *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998)*, Madison, WI, USA, July 24–27, 1998. (Morgan Kaufmann) pp. 296–304.
 26. Smith, C.L., and Eppig, J.T. (2009). The mammalian phenotype ontology: enabling robust annotation and comparative analysis. *Wiley Interdiscip. Rev. Syst. Biol. Med.* 1, 390–399.
 27. Köhler, S., Doelken, S.C., Ruef, B.J., Bauer, S., Washington, N., Westerfield, M., Gkoutos, G., Schofield, P., Smedley, D., Lewis, S.E., et al. (2013). Construction and accessibility of a cross-species phenotype ontology along with gene annotations for biomedical research. *F1000Res.* 2, 30.
 28. Carlin, B.P., and Chib, S. (1995). Bayesian model choice via markov chain monte carlo methods. *J. R. Stat. Soc., B* 57, 473–484.
 29. Cingolani, P., Platts, A., Wang, L., Coon, M., Nguyen, T., Wang, L., Land, S.J., Lu, X., and Ruden, D.M. (2012). A program for annotating and predicting the effects of single nucleotide polymorphisms, SnpEff: SNPs in the genome of *Drosophila melanogaster* strain w1118; iso-2; iso-3. *Fly (Austin)* 6, 80–92.
 30. Kunishima, S., Okuno, Y., Yoshida, K., Shiraiishi, Y., Sanada, M., Muramatsu, H., Chiba, K., Tanaka, H., Miyazaki, K., Sakai, M., et al. (2013). *ACTN1* mutations cause congenital macrothrombocytopenia. *Am. J. Hum. Genet.* 92, 431–438.
 31. Stritt, S., Nurden, P., Turro, E., Greene, D., Jansen, S.B.G., Westbury, S.K., Petersen, R., Astle, W.J., Marlin, S., Bariana, T.K., et al. (2016). A gain-of-function variant in *DIAPH1* causes dominant macrothrombocytopenia and hearing loss. *Blood* <http://dx.doi.org/10.1182/blood-2015-10-675629>.
 32. Lynch, E.D., Lee, M.K., Morrow, J.E., Welsh, P.L., León, P.E., and King, M.C. (1997). Nonsyndromic deafness DFNA1 associated with mutation of a human homolog of the *Drosophila* gene *diaphanous*. *Science* 278, 1315–1318.
 33. Canault, M., Ghalloussi, D., Grosdidier, C., Guinier, M., Perret, C., Chelghoum, N., Germain, M., Raslova, H., Peiretti, F., Morange, P.E., et al. (2014). Human CalDAG-GEF1 gene (*RASGRP2*) mutation affects platelet function and causes severe bleeding. *J. Exp. Med.* 211, 1349–1362.
 34. Akawi, N., McRae, J., Ansari, M., Balasubramanian, M., Blyth, M., Brady, A.F., Clayton, S., Cole, T., Deshpande, C., Fitzgerald, T.W., et al.; DDD study (2015). Discovery of four recessive developmental disorders using probabilistic genotype and phenotype matching among 4,125 families. *Nat. Genet.* 47, 1363–1369.