Published in final edited form as:

Genet Epidemiol. 2013 February ; 37(2): . doi:10.1002/gepi.21705.

Importance of different types of prior knowledge in selecting genome-wide findings for follow-up

Cosetta Minelli^{1,*}, Alessandro De Grandi^{1,*}, Christian X. Weichenberger¹, Martin Gögele¹, Mirko Modenese¹, John Attia^{2,3}, Jennifer H. Barrett⁴, Michael Boehnke⁵, Giuseppe Borsani⁶, Giorgio Casari⁷, Caroline S. Fox^{8,9}, Thomas Freina¹, Andrew A. Hicks¹, Fabio Marroni¹⁰, Giovanni Parmigiani^{11,12}, Andrea Pastore¹³, Cristian Pattaro¹, Arne Pfeufer¹⁴, Fabrizio Ruggeri¹⁵, Christine Schwenbacher¹, Daniel Taliun¹, Peter P. Pramstaller^{1,16,17}, Francisco S. Domingues¹, and John R. Thompson¹⁸

¹Center for Biomedicine, European Academy Bozen/Bolzano (EURAC), Bolzano, Italy (Affiliated Institute of the University of Lübeck, Lübeck, Germany)

²Centre for Clinical Epidemiology and Biostatistics, The University of Newcastle, Hunter Medical Research Institute, Newcastle, NSW, Australia

³Department of General Medicine, John Hunter Hospital, Newcastle, NSW, Australia

⁴Section of Epidemiology and Biostatistics, Leeds Institute of Molecular Medicine, University of Leeds, Leeds, UK

⁵Department of Biostatistics and Center for Statistical Genetics, University of Michigan, Ann Arbor, MI, USA

⁶Department of Biomedical Sciences and Biotechnology, University of Brescia, Brescia, Italy

⁷Vita-Salute San Raffaele University and Center for Translational Genomics and Bioinformatics, Milan, Italy

⁸Caroline S. Fox: The National Heart, Lung, and Blood Institute's Framingham Heart Study, Framingham, Massachusetts, USA

⁹Center for Population Studies, Framingham, Massachusetts, USA

¹⁰Institute of Applied Genomics, Udine, Italy

¹¹Department of Biostatistics and Computational Biology, Dana-Farber Cancer Institute, Boston, Massachusetts

¹²Department of Biostatistics, Harvard School of Public Health, Boston, Massachusetts

¹³Department of Economics, Ca' Foscari Venezia University, Venezia, Italy

¹⁴Department of Bioinformatics and Systems Biology IBIS, Helmholtz Zentrum Munich, German Research Center for Environmental Health (GmbH), Neuherberg, Germany

¹⁵National Research Council, Institute of Applied Mathematics and Information Technology, Milan, Italy

Correspondence to: Cosetta Minelli, Center for Biomedicine, EURAC Research, Viale Druso, 1 - 39100 Bolzano (Italy).

*Joint first authors

Supplemental Data

Supplemental Data include six tables and three figures.

Conflict of Interest

None of the authors declares any conflict of interest.

¹⁶Department of Neurology, Central Hospital, Bolzano, Italy

¹⁷Department of Neurology, University of Lübeck, Lübeck, Germany

¹⁸Department of Health Sciences, University of Leicester, Leicester, UK

Abstract

Biological plausibility and other prior information could help select genome-wide association (GWA) findings for further follow-up, but there is no consensus on which types of knowledge should be considered or how to weight them. We used experts' opinions and empirical evidence to estimate the relative importance of 15 types of information at the single nucleotide polymorphism (SNP) and gene levels. Opinions were elicited from ten experts using a two-round Delphi survey. Empirical evidence was obtained by comparing the frequency of each type of characteristic in SNPs established as being associated with seven disease traits through GWA meta-analysis and independent replication, with the corresponding frequency in a randomly selected set of SNPs. SNP and gene characteristics were retrieved using a specially developed bioinformatics tool. Both the expert and the empirical evidence rated previous association in a meta-analysis or more than one study as conferring the highest relative probability of true association, while previous association in a single study ranked much lower. High relative probabilities were also observed for location in a functional protein domain, while location in a region evolutionarily conserved in vertebrates was ranked high by the data but not by the experts. Our empirical evidence did not support the importance attributed by the experts to whether the gene encodes a protein in a pathway or shows interactions relevant to the trait. Our findings provide insight into the selection and weighting of different types of knowledge in SNP or gene prioritization, and point to areas requiring further research.

Keywords

Gene prioritization; Genome-wide association studies; Bioinformatics databases

INTRODUCTION

In genome-wide association (GWA) studies, the choice of which single nucleotide polymorphisms (SNPs) should be followed up for replication in independent samples or for functional investigation can either be based purely on discovery p-values or can incorporate prior knowledge about the SNP and its possible association with the phenotype of interest. Selection of SNPs for replication based purely on discovery p-values is currently the most common approach [Gögele et al., 2012], but this strategy tends to have low power to identify good candidates when the discovery sample is relatively small, particularly for SNPs with low minor allele frequency [Liu et al., 2008]. Despite current efforts to increase power by pooling GWA data from different studies, small discovery sample size can still be a critical issue for rarer disease outcomes or phenotypes that are difficult to measure, and the presence of heterogeneity across studies can further reduce statistical power [Ioannidis, 2007, Greene et al., 2009, Kraft, Zeggini and Ioannidis, 2009]. Incorporating prior information on biological gene function or findings from previous genetic epidemiological studies can help select the most promising SNPs in a more informed way, thus potentially increasing the yield of downstream studies [Moreau and Tranchevent, 2012]. Such information may derive from very different sources, including gene expression and proteomics studies, genetic studies in animal models, and previous association or linkage studies in humans.

In practice, prior knowledge has been used to aid gene prioritization in many different ways. Sometimes investigators add to the list of GWA top hits sent to replication additional SNPs

within genes known by the authors to have been previously linked to the phenotype [Gögele et al., 2012], although this leaves the reader in doubt about whether other SNPs with even higher support from prior knowledge might have been omitted. Other authors have used more systematic ways of identifying relevant prior information for SNP selection, either focusing on a single type of evidence, such as pathway analysis, or combining different types of evidence [Thomas et al., 2009, Cantor, Lange and Sinsheimer, 2010, Saccone et al., 2008]. Lack of evidence on the relative informativeness of different types of prior knowledge means that decisions on what information is worth retrieving and how much weight should be attributed to different types of knowledge are inevitably subjective. This work provides suitable weights for estimating the likelihood of true association given certain types of prior knowledge, and contrasts the views of experts with empirical evidence taken from published GWA meta-analyses.

MATERIAL AND METHODS

Elicitation of experts' opinions

Ten experts in the field of GWA investigations from different backgrounds (molecular biology, genetic epidemiology, statistical genetics) were asked to participate in the study, without being told the identity of the other experts [Akins, Tolson and Cole, 2005]. Upon acceptance, a two-round Delphi survey was used to elicit their opinions through pre-prepared questionnaires circulated by e-mail. The Delphi method is a form of structured group communication process, consisting of an expert survey organized in two (or more) rounds [Adler and Ziglio, 1996]. In the second round, the anonymous results for all experts in the previous stage were given as feedback, and the same experts were asked to reassess their answers to the same set of questions in the light of their colleagues' opinions. These questions referred either to the specific SNP or to the gene(s) lying within 5kb of the SNP. The questions did not refer to any specific phenotype, and the experts were asked to think in general terms.

In the first round, experts were presented with a list of 20 items (Supplementary Table 1), and were asked to provide their "best guess" on how many times more likely a SNP was to be truly associated with the phenotype given a certain characteristic, when compared with a SNP with no such characteristic (hereafter referred to as the "relative probability" associated with that characteristic). Experts were asked to answer as if each type of evidence were the only external information available, so that all types of evidence were treated as independent. To help ensure consistent interpretation of the scale, the experts were provided with an example for which the probability that a random SNP was truly associated with the disease was around 1 in 10,000, so that an answer of 5 times more likely would translate into a probability of true association of 5 in 10,000.

In the second round, the number of questions was reduced based on findings from the first round (see "Statistical analyses") and the experts were asked to provide a revised answer to each question, together with a 95% interval representing their uncertainty, reflecting both their own experience and the results of the first round averaged across experts (Supplementary Table 1). Estimates based on the experts' opinions are hereafter referred to as "opinions".

Empirical evidence

Estimates of empirical relative probabilities—We obtained empirical estimates of the relative probability of association for a SNP given a certain type of prior knowledge using a "case-control" approach. We chose seven disease traits for which a set of SNPs, referred to as "true SNPs", had been identified through large GWA investigations and had

been replicated (Table I). With “true SNPs” representing our “cases” and a set of 1,000 random SNPs as our “controls”, we estimated relative probabilities of association by comparing the proportion of true SNPs for which a certain type of evidence was present versus the proportion in random SNPs (see Statistical analysis). A different set of 1,000 random SNPs was selected for each trait throughout the genome, from about 2,500,000 SNPs with minor allele frequency greater than 0.01 (the sampling frame was that of all SNPs used in the eGFRcrea meta-analysis [Köttgen et al., 2010] and can be found at: <https://intramural.nhlbi.nih.gov/labs/CF/Pages/CKDGenConsortium.aspx>). Since true and random SNPs were not matched by allele frequency, we also performed a sensitivity analysis adjusting for allele frequency.

The seven selected traits were estimated glomerular filtration rate (as a measure of renal dysfunction), Crohn’s disease, coronary artery disease, rheumatoid arthritis, primary biliary cirrhosis, type 2 diabetes and body mass index (as a measure of obesity). For each trait, the list of true SNPs was compiled based on the most recent (one or more) GWA meta-analysis (Table 1), after excluding SNPs that had been selected for replication based on prior knowledge rather than GWA evidence. Estimates based on empirical evidence are hereafter referred to as “data”.

To enable a comparison with the relative probabilities given by the experts, the data-based results for each SNP characteristic were also considered independently. However, the empirical evidence also enables us to investigate the dependence between the questions, and to determine the weights that should be used in a combined estimate that incorporates all of the SNP characteristics, as shown in the companion paper by Thompson et al. [Thompson et al., 2012] (see Discussion).

Bioinformatics retrieval of information—Information on each of the types of prior knowledge was retrieved for both the true SNPs and for 1,000 random SNPs in a standardized and automatic way, by use of a bioinformatics tool developed for this project. Ensembl [Stabenau et al., 2004, Flicek et al., 2011] was the main data source queried by the tool, but additional public databases were used to answer specific questions, in particular HUGE Navigator [Yu et al., 2008], Pfam [Finn et al., 2010], cisRED [Robertson et al., 2006], VISTA [Visel et al., 2007], miRanda [John et al., 2004], Mouse Genome Informatics (MGI) database [Blake et al., 2011], BioGPS [Wu et al., 2009, Su et al., 2004], KEGG [Kanehisa and Goto, 2000] and IntAct [Aranda et al., 2010]. Supplementary Table 2 summarizes the structure of each query, and the source code is available on our website at <https://gemex.eurac.edu/downloads/stats/GenEpi2012>.

To allow retrieval of evidence using these databases, the formulation of the query had, in a few cases, to be modified slightly from the initial question presented to the experts (Table II). All types of evidence regarding relationships between genes and phenotypes were retrieved using MESH terms linked to UMLS CUIs (Unified Medical Language System Concept Unique Identifier) directly referring to the particular phenotype (Table I; Supplementary Table 2), while questions to the experts had been phrased more generally as “the same/closely related phenotype”. The potential impact of this difference was assessed in sensitivity analyses performed on three of the seven traits by repeating the retrieval of evidence after extending the list of UMLS CUI terms to cover “closely related” phenotypes. One important difference was related to evidence from previous genetic association studies (Q8, Q9 and Q11). Since HuGE Navigator only provides information on whether a gene-phenotype association has been investigated and not on whether it has been established as true, the search of HuGE Navigator includes publications with negative findings [Yu et al., 2008]. This problem is typical of all search engines based on text mining of published literature, and we are not aware of any alternative public resource, covering both candidate-

gene and GWA studies, which also provides the result of each investigated association. Formulation of the question on “functional models” (Q12) also had to be modified, from “evidence from *in vitro* and animal studies” in the question to experts to “evidence from mouse models (MGI database)” in the bioinformatics query, because we could not find public databases from which we could retrieve the required genome-wide information from other functional models. Finally, the question on the importance of whether the SNP is in a gene which shows gene/protein (gene-gene, gene-protein, or protein-protein) interactions relevant to the phenotype (Q15) was restricted to protein-protein interactions in the bioinformatics query.

We could not obtain by automated methods empirical evidence on the relative probability of association given supporting knowledge from linkage studies (Q10: “The SNP is in a gene ($\pm 5\text{kb}$) which is under a linkage peak that has been associated with the same/closely related phenotype”), due to the lack of electronically processable public databases summarizing published genome-wide linkage findings.

Statistical analyses

For experts’ opinions, correlations between items of the original questionnaire in the first Delphi round were analyzed to help reduce the list of types of evidence to be evaluated, by dropping questions that appeared not to convey much additional information.

We estimated empirical relative probabilities as odds ratios of true association with logistic regression analysis, modeling the probability of being a “true SNP” given the presence of each type of evidence.

For both opinions and data, relative probabilities were analyzed after log transformation. Inverse variance meta-analysis based on either a fixed- or random-effect model, depending on absence or presence of heterogeneity respectively, was used to pool opinions across experts and data across traits. Between-expert and between-trait heterogeneity was tested using chi-square tests with statistical significance defined at $p\text{-value} < 0.10$ [Fleiss, 1993], and the magnitude of the heterogeneity was estimated using the I^2 , representing the percentage of variability in estimates explained by heterogeneity rather than sampling error [Higgins et al., 2003].

In this paper, the term “relative probability”, used to indicate the probability of true association for a SNP given a certain type of prior knowledge compared with a SNP with no such evidence, refers to a relative risk for opinions and to an odds ratio for data. However, the impact of such difference on the comparison between opinions and data should be minimal, given the very low frequency of the outcome, represented by the *a priori* probability of true association for any given SNP in the genome [Davies, Crombie and Tavakoli, 1998].

RESULTS

Experts’ opinions

Based on the findings of the first Delphi round and the correlation coefficients between items of the original questionnaire (Supplementary Figures 1 and 2), the number of questions was reduced from 20 to 15. An item was dropped when it did not seem to convey additional information compared with another more general or relevant one, but only if the two were highly correlated and their relative probabilities very similar across experts (Supplementary Table 1). Although we had planned to drop types of evidence showing relative probabilities close to one (i.e. no relevance at all in support of a true association) consistently across experts, no such types of evidence were identified. Rewording of some

questions which could have been interpreted either as mutually exclusive or overlapping was also performed to improve clarity in the second round (Supplementary Table 1).

Nine of the ten experts completed the second round. Some experts were more prone than others to change their opinion towards the average, and the extent of the changes from the first to the second round also varied across questions (Supplementary Table 1). Pooled estimates of relative probabilities across experts are presented in Table II, and were obtained using a random-effect meta-analysis model due to the presence of moderate to large between-expert heterogeneity for most questions. Subgroup analyses by experts' background, biological vs. non-biological, could not explain the heterogeneity observed (data not shown). Heterogeneity disappeared in two questions after excluding one outlying expert, who provided much higher relative probabilities for most questions (Supplementary Table 3).

It turned out that the types of evidence considered as the most important (relative probabilities >10) were related to information at the gene level rather than to the SNP itself (Table II). In decreasing order of perceived importance, they were: gene previously associated with the phenotype in a meta-analysis or more than one study (Q8); gene previously associated with the phenotype in functional models (Q12); gene encoding for a protein in a pathway relevant to the phenotype (Q14); gene which shows gene/protein interactions relevant to the phenotype (Q15).

Empirical evidence

The number of true SNPs varied across traits between 18 and 71 (Table I). Table II reports the frequency of each type of evidence in truly associated and random SNPs, calculated as a weighted average across traits, while histograms for each trait are presented in Supplementary Figure 3 (the full set of results is available on our website at <https://gemex.eurac.edu/downloads/stats/GenEpi2012>). Some types of evidence were commonly observed in the sample of SNPs and their relative probabilities could be accurately estimated, while others, often with high relative probabilities, were rare (genuine low frequency or limited coverage of the bioinformatics query/data source) and had imprecise estimates. Pooled estimates of relative probabilities across traits were obtained using a fixed-effect model given the absence of substantial heterogeneity for most questions. Pooled estimates are presented and compared with those from experts' opinions in Table II and in the forest plots in Figure 1, while findings for the individual traits are reported in Supplementary Table 4.

The type of evidence with the highest relative probability was previous association of the gene with the phenotype in a meta-analysis or in more than one study (Q8), while previous association in a single study (Q9) showed a much lower relative probability. Although this was in line with experts' opinions, estimates for the two types of evidence were substantially lower in the data, 21 (95% confidence interval: 17 to 27) vs. 49 (95% interval: 20 to 123) for association in meta-analysis/more than one study, and 2 (1 to 4) vs. 6 (5 to 9) for association in single study. Similarly, the relative probability for whether the SNP is in a locus within which other SNPs have been previously associated with the phenotype (Q11) was significantly lower in data compared with opinions, 4 (3 to 5) vs. 10 (5 to 21). The difference between the data and the opinions for these three types of evidence might be partly explained by the fact that the bioinformatics tool retrieved evidence on "previous investigation" rather than "previous association". This represents a measurement error in the assessment of the exposure (presence or absence of previous association), and as such is more likely to introduce bias towards the null, leading to underestimation of the relative probability in the data.

The type of evidence with the second highest relative probability in the data was whether the SNP is in a functional protein domain (Q4), with relative probability higher than in experts' opinions, 10 (6 to 15) vs. 6 (4 to 10). Similarly to the opinions, relative probabilities for the other three questions on the SNP's possible functional role (transcribed, Q1; translated, Q2; changes the amino acid, Q3) and the two questions on location in regulatory regions (not transcribed, Q5; transcribed, Q6) reflected their hierarchical structure. Relative probabilities for questions dealing with regulatory regions were significantly lower than those based on opinions, and only the question addressing whether the SNP is located in a transcribed region was statistically different from one.

Previous association of the gene with the phenotype in functional models (Q12) had the third highest relative probability. The scarcity of the data available for this type of evidence, which was limited to mouse models, made the estimate imprecise, 10 (2 to 39), and the observed difference with the experts' opinions (22; 12 to 40) could be due to chance.

Location of the SNP in a genomic region evolutionarily conserved in vertebrates (Q7) had a relative probability of 6 (4 to 8), much higher than in experts' opinions. Data and opinions gave very similar estimates for high gene expression in a tissue relevant to the phenotype (Q13), with relative probabilities around 3. Finally, whether the gene encodes for a protein which is in a pathway (Q14) or shows protein/protein interactions (Q15) relevant to the phenotype had relative probabilities around 2, significantly lower than in experts' opinions (16 and 10, respectively).

Results of the sensitivity analyses performed for three of the seven traits by extending the list of UMLS CUI terms to cover "closely related" phenotypes, as formulated in the questions to experts, were similar to the main results (Supplementary Table 5). Similarly, adjusting the analyses for allele frequency did not change the results (data not shown).

Dependence between questions—The degree of dependence between questions, as expressed by the correlation matrix of 7,000 random SNPs, is shown in Supplementary Table 6. The correlation of Q7 (SNP in a region evolutionary conserved in vertebrates) with Q2, Q3 and Q4 (indicating a possible functional role, from translated without amino acid change to translated in functional protein domain) may be due to functional regions being more likely to be conserved [Levenstien and Klein, 2011]. Similarly, the correlation between Q14 (SNP in a gene encoding for a protein in a pathway relevant to the phenotype) and Q15 (SNP in a gene showing protein-protein interactions relevant to the phenotype) might be explained by the fact that proteins in a pathway may also interact with each other [Kirouac et al., 2012]. The interdependence of the different types of prior knowledge needs to be accounted for when they are used together, through conditional analyses that jointly model them as discussed in the companion paper by Thompson et al. [Thompson et al., 2012].

DISCUSSION

The use of prior knowledge may improve the selection of GWA signals for follow-up, thus increasing the probability of a successful replication or functional investigation. Studies which have systematically incorporated prior knowledge from multiple data sources using bioinformatics tools have attributed equal importance to the different types of evidence [Sookoian et al., 2009, Chen et al., 2011, Aerts et al., 2006], and yet our findings suggest that this may be suboptimal. Our study convincingly shows that, for commonly investigated traits, evidence from previous association studies on the phenotype of interest represents the most informative type of knowledge for gene prioritization, although it does not help discover novel genes. The empirical findings suggest that SNPs in genes previously investigated in relation with the phenotype in a meta-analysis or in more than one study are

21 times more likely to represent true associations, with this being reduced to 2 times if previous investigation is limited to a single study.

Our findings suggest that location of the SNP in a functional protein domain may increase the probability of true association up to 10 times, with progressively decreasing effect for whether the SNP changes the amino acid but is not in a functional protein domain, and whether the SNP is in a translated region but does not change the amino acid. Despite the very low proportion of SNPs with these characteristics, between 1% and 5% in the SNPs associated with our seven traits, information on SNP characteristics can be retrieved easily and accurately so that these types of evidence are worth considering in the prioritization of GWA signals. Similarly, location of the SNP in a gene previously associated with the phenotype in functional models could substantially increase the probability of true association, by 9 times according to our empirical findings limited to mouse data, but up to 23 times in experts' opinions regarding animal and *in vitro* models in general. The mouse model is widely used [Hardouin and Nagy, 2000, Rosenthal and Brown, 2007], and the observed empirical estimates were highly consistent across the seven traits considered, suggesting that retrieving such information can still be useful when other functional data cannot be accessed. However, its frequency was very low (1% in our "true SNPs"), and the practical importance of incorporating functional evidence in gene prioritization would increase by considering additional models.

It is interesting to note how studies which have tried to incorporate prior knowledge have often disregarded knowledge of association of the gene with the phenotype of interest from human and animal studies, but rather focused on information on pathways or protein-protein interactions, SNP characteristics and gene expression data [Saccone et al., 2008, Parikh, Lyssenko and Groop, 2009, Chen, Aronow and Jegga, 2009, Zhong et al., 2010]. Use of gene pathway analysis for gene prioritization has received much attention and bioinformatics tools have been developed to allow retrieval of such information at genome-wide level [Cantor, Lange and Sinsheimer, 2010, Zhong et al., 2010, Elbers et al., 2009]. This is reflected by experts' opinion, which ranked this type of information as the third most important. However, our empirical findings based on pathway information retrieved using KEGG [Kanehisa and Goto, 2000] do not support this view. This may be partly explained by the difficulty of defining the boundaries of a pathway, but nonetheless suggests that more investigation is needed to evaluate the potential value of pathway information and how it should be modeled. Similarly, our empirical findings did not support the importance attributed by the experts to information on whether the gene product shows evidence of interactions relevant to the phenotype. On the other hand, our empirical findings suggest that presence of the SNP in a genomic region evolutionarily conserved in vertebrates could increase the probability of true association by 6 times, contrary to experts' opinion that ranked this as the least important type of knowledge. Although it occurred in only 10% of our "true SNPs", this type of evidence can be easily and accurately retrieved and may well be incorporated in gene prioritization. An interesting follow-up of our study will be to investigate the impact of the choice of a 5kb window for mapping SNPs to genes, and to provide evidence on what window might be the most informative. Such choice is likely to influence the estimated relative probability of association of many types of knowledge, including whether the gene encodes a protein in a pathway.

Limitations of the study

Our empirical findings are based on only seven examples of gene-disease associations, but their generalization is supported by the high consistency observed across traits for most types of evidence. The precision of our empirical weights could be improved by considering more traits and increasing the number of "true SNPs" analyzed. This could be done systematically using publicly available databases such as the NHGRI GWAS Catalog, a

continuously updated catalog of findings from published GWA investigations [Hindorff et al., 2009]. As for the selection of “true SNPs” for each of the seven traits, while the completeness of our lists is not an issue, its representativeness is, and it may well be that SNPs identified by GWA investigations and replicated are not representative of all true genetic associations, particularly those with weaker effects.

Many of the “true SNPs” in our seven traits may be SNPs in linkage disequilibrium with the real causal variants, so that types of evidence referring to the characteristics of the SNP (e.g. “the SNP is in a functional protein domain”) may be negative for the “true SNP” only because this is in fact only a proxy of the causal one. This, which can be interpreted as measurement error in the definition of our “cases”, is likely to introduce bias towards the null and therefore lead to underestimation of relative probabilities, particularly for types of evidence referring to SNP characteristics. Other forms of mis-classification could in theory lead to a bias away from the null, for example if, because of the way in which the GWA studies have been conducted, a “true SNP” is more likely to have been identified if it has certain of the characteristics, or due to inaccuracies in the SNP annotations. Empirical estimates of relative probabilities of true associations will be different in the future, when GWA findings will be based on newer sequencing and genotyping technology resulting in higher genome coverage and improved reliability. In general, empirical estimates of the relative importance of different types of evidence will depend on current knowledge and data availability, so requiring continual synchronization of the query databases.

Regarding the retrieval of the evidence, our study shows the limitations of using bioinformatics tools that search for prior knowledge at genome-wide level from publicly available databases, and the practical limits on certain types of information, such as evidence from linkage studies, functional studies other than mouse models, and eQTL databases for gene expression from multiple tissue sources. Moreover, data quality strongly depends on the coverage provided by the interrogated databases, which suggests that integrating information on a certain type of evidence from multiple databases may be preferable to relying on a single one.

Practical use of our findings

Investigators willing to incorporate prior information on biological function or evidence from previous studies in the selection of GWA hits for follow up encounter a few practical issues: What types of prior knowledge are worthwhile considering? How can prior knowledge be retrieved in a systematic way? How can prior knowledge be combined with the discovery p-values? How should different types of knowledge be differentially weighted to provide an overall *a priori* probability of association for each SNP? Our findings answer the question about the relative importance of different types of prior knowledge and show the feasibility of automatic retrieval of such information using a bioinformatics tool that queries multiple data sources. A companion paper by Thompson et al. [Thompson et al., 2012] demonstrates the use of prior knowledge in combination with discovery p-values within a Bayesian framework to provide a posterior probability of replication, which can be used to rank the most promising SNPs for follow up. That work combines our estimates of relative probabilities for the 14 types of knowledge and calculates the overall prior probability of association for a given SNP.

Thompson et al. demonstrate that the success of replication is increased when the selection of SNPs incorporates prior knowledge using a simple approximate Bayesian analysis, compared with the classical approach purely based on discovery p-values.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

All researchers from the Center for Biomedicine at EURAC were supported by the Department for Promotion of Educational Policies, Universities and Research of the Autonomous Province of Bolzano, South Tyrol, Italy; Michael Boehnke was supported by NIH grant HG000376.

References

- Adler, M.; Ziglio, E. Gazing into the oracle: the Delphi method and its application to social policy and public health. London: Jessica Kingsley; 1996.
- Aerts S, Lambrechts D, Maity S, et al. Gene prioritization through genomic data fusion. *Nat Biotechnol.* 2006; 24:537–44. [PubMed: 16680138]
- Akins RB, Tolson H, Cole BR. Stability of response characteristics of a Delphi panel: application of bootstrap data expansion. *BMC Med Res Methodol.* 2005; 5:37. [PubMed: 16321161]
- Aranda B, Achuthan P, Alam-Faruque Y, et al. The IntAct molecular interaction database in 2010. *Nucleic Acids Res.* 2010; 38:D525–31. [PubMed: 19850723]
- Blake JA, Bult CJ, Kadin JA, Richardson JE, Eppig JT. Mouse Genome Database Group. The Mouse Genome Database (MGD): premier model organism resource for mammalian genomics and genetics. *Nucleic Acids Res.* 2011; 39:D842–8. [PubMed: 21051359]
- Cantor RM, Lange K, Sinsheimer JS. Prioritizing GWAS results: A review of statistical methods and recommendations for their application. *Am J Hum Genet.* 2010; 86:6–22. [PubMed: 20074509]
- Chen J, Aronow BJ, Jegga AG. Disease candidate gene identification and prioritization using protein interaction networks. *BMC Bioinformatics.* 2009; 10:73. [PubMed: 19245720]
- Chen Y, Wang W, Zhou Y, et al. In silico gene prioritization by integrating multiple data sources. *PLoS One.* 2011; 6:e21137. [PubMed: 21731658]
- Coronary Artery Disease (C4D) Genetics Consortium. A genome-wide association study in Europeans and South Asians identifies five new loci for coronary artery disease. *Nat Genet.* 2011; 43:339–44. [PubMed: 21378988]
- Davies HT, Crombie IK, Tavakoli M. When can odds ratios mislead? *BMJ.* 1998; 316:989–91. [PubMed: 9550961]
- Elbers CC, van Eijk KR, Franke L, et al. Using genome-wide pathway analysis to unravel the etiology of complex diseases. *Genet Epidemiol.* 2009; 33:419–31. [PubMed: 19235186]
- Finn RD, Mistry J, Tate J, et al. The Pfam protein families database. *Nucleic Acids Res.* 2010; 38:D211–22. [PubMed: 19920124]
- Fleiss JL. The statistical basis of meta-analysis. *Stat Methods Med Res.* 1993; 2:121–45. [PubMed: 8261254]
- Flicek P, Amode MR, Barrell D, et al. Ensembl 2011. *Nucleic Acids Res.* 2011; 39:D800–6. [PubMed: 21045057]
- Franke A, McGovern DP, Barrett JC, et al. Genome-wide meta-analysis increases to 71 the number of confirmed Crohn's disease susceptibility loci. *Nat Genet.* 2010; 42:1118–25. [PubMed: 21102463]
- Gögele M, Minelli C, Thakkeinstian A, et al. Methods for meta-analyses of genome-wide association studies: critical assessment of empirical evidence. *Am J Epidemiol.* 2012; 175:739–49. [PubMed: 22427610]
- Greene CS, Penrod NM, Williams SM, Moore JH. Failure to replicate a genetic association may provide important clues about genetic architecture. *PLoS One.* 2009; 4:e5639. [PubMed: 19503614]
- Hardouin SN, Nagy A. Mouse models for human disease. *Clin Genet.* 2000; 57:237–44. [PubMed: 10845564]
- Higgins JP, Thompson SG, Deeks JJ, Altman DG. Measuring inconsistency in meta-analyses. *BMJ.* 2003; 327:557–60. [PubMed: 12958120]

- Hindorff LA, Sethupathy P, Junkins HA, et al. Potential etiologic and functional implications of genome-wide association loci for human diseases and traits. *Proc Natl Acad Sci U S A*. 2009; 106:9362–7. [PubMed: 19474294]
- Ioannidis JP. Non-replication and inconsistency in the genome-wide association setting. *Hum Hered*. 2007; 64:203–13. [PubMed: 17551261]
- John B, Enright AJ, Aravin A, Tuschl T, Sander C, Marks DS. Human MicroRNA targets. *PLoS Biol*. 2004; 2:e363. [PubMed: 15502875]
- Kanehisa M, Goto S. KEGG: kyoto encyclopedia of genes and genomes. *Nucleic Acids Res*. 2000; 28:27–30. [PubMed: 10592173]
- Kirouac DC, Saez-Rodriguez J, Swantek J, Burke JM, Lauffenburger DA, Sorger PK. Creating and analyzing pathway and protein interaction compendia for modelling signal transduction networks. *BMC Syst Biol*. 2012; 6:29. [PubMed: 22548703]
- Köttgen A, Pattaro C, Böger CA, et al. New loci associated with kidney function and chronic kidney disease. *Nat Genet*. 2010; 42:376–84. [PubMed: 20383146]
- Kraft P, Zeggini E, Ioannidis JP. Replication in genome-wide association studies. *Stat Sci*. 2009; 24:561–73. [PubMed: 20454541]
- Levenstien MA, Klein RJ. Predicting functionally important SNP classes based on negative selection. *BMC Bioinformatics*. 2011; 12:26. [PubMed: 21247465]
- Liu YJ, Papasian CJ, Liu JF, Hamilton J, Deng HW. Is replication the gold standard for validating genome-wide association findings? *PLoS One*. 2008; 3:e4037. [PubMed: 19112512]
- Mells GF, Floyd JA, Morley KI, et al. Genome-wide association study identifies 12 new susceptibility loci for primary biliary cirrhosis. *Nat Genet*. 2011; 43:329–32. [PubMed: 21399635]
- Moreau Y, Tranchevent LC. Computational tools for prioritizing candidate genes: boosting disease gene discovery. *Nat Rev Genet*. 2012; 13:523–36. [PubMed: 22751426]
- Parikh H, Lyssenko V, Groop LC. Prioritizing genes for follow-up from genome wide association studies using information on gene expression in tissues relevant for type 2 diabetes mellitus. *BMC Med Genomics*. 2009; 2:72. [PubMed: 20043853]
- Robertson G, Bilenky M, Lin K, et al. cisRED: a database system for genome-scale computational discovery of regulatory elements. *Nucleic Acids Res*. 2006; 34:D68–73. [PubMed: 16381958]
- Rosenthal N, Brown S. The mouse ascending: perspectives for human-disease models. *Nat Cell Biol*. 2007; 9:993–9. [PubMed: 17762889]
- Saccone SF, Saccone NL, Swan GE, et al. Systematic biological prioritization after a genome-wide association study: an application to nicotine dependence. *Bioinformatics*. 2008; 24:1805–11. [PubMed: 18565990]
- Schunkert H, König IR, Kathiresan S, et al. Large-scale association analysis identifies 13 new susceptibility loci for coronary artery disease. *Nat Genet*. 2011; 43:333–8. [PubMed: 21378990]
- Sookoian S, Gianotti TF, Schuman M, Pirola CJ. Gene prioritization based on biological plausibility over genome wide association studies renders new loci associated with type 2 diabetes. *Genet Med*. 2009; 11:338–43. [PubMed: 19346957]
- Speliotes EK, Willer CJ, Berndt SI, et al. Association analyses of 249,796 individuals reveal 18 new loci associated with body mass index. *Nat Genet*. 2010; 42:937–48. [PubMed: 20935630]
- Stabenau A, McVicker G, Melsopp C, Proctor G, Clamp M, Birney E. The Ensembl core software libraries. *Genome Res*. 2004; 14:929–33. [PubMed: 15123588]
- Stahl EA, Raychaudhuri S, Remmers EF, et al. Genome-wide association study meta-analysis identifies seven new rheumatoid arthritis risk loci. *Nat Genet*. 2010; 42:508–14. [PubMed: 20453842]
- Su AI, Wiltshire T, Batalov S, et al. A gene atlas of the mouse and human protein-encoding transcriptomes. *Proc Natl Acad Sci U S A*. 2004; 101:6062–7. [PubMed: 15075390]
- Thomas DC, Conti DV, Baurley J, Nijhout F, Reed M, Ulrich CM. Use of pathway information in molecular epidemiology. *Hum Genomics*. 2009; 4:21–42. [PubMed: 21072972]
- Thompson JR, Gögele M, Weichenberger CX, et al. A Bayesian method for calculating the probability of replication by combining SNP information and genome-wide data. *Genet Epidemiol* (Companion paper). 2012

- Visel A, Minovitsky S, Dubchak I, Pennacchio LA. VISTA Enhancer Browser--a database of tissue-specific human enhancers. *Nucleic Acids Res.* 2007; 35:D88–92. [PubMed: 17130149]
- Voight BF, Scott LJ, Steinthorsdottir V, et al. Twelve type 2 diabetes susceptibility loci identified through large-scale association analysis. *Nat Genet.* 2010; 42:579–89. [PubMed: 20581827]
- Wu C, Orozco C, Boyer J, et al. BioGPS: an extensible and customizable portal for querying and organizing gene annotation resources. *Genome Biol.* 2009; 10:R130. [PubMed: 19919682]
- Yu W, Gwinn M, Clyne M, Yesupriya A, Khoury MJ. A navigator for human genome epidemiology. *Nat Genet.* 2008; 40:124–5. [PubMed: 18227866]
- Zhong H, Yang X, Kaplan LM, Molony C, Schadt EE. Integrating pathway analysis and genetics of gene expression for genome-wide association studies. *Am J Hum Genet.* 2010; 86:581–91. [PubMed: 20346437]

Figure 1.

Plots of the relative probabilities of association for each type of prior knowledge from empirical evidence (with 95% confidence intervals) compared with those based on experts' opinions (with 95% intervals). Q10 is omitted as empirical evidence is not available (data could not be processed electronically)

*Questions for which the difference in estimates from data vs. opinions are statistically significant (p-value<0.05)

Table I

Data source and number of “true SNPs” for the seven disease traits used to estimate empirical relative probabilities. “True SNPs” were SNPs established as being associated with the seven traits through GWA meta-analysis and independent replication. MESH terms are those used for the bioinformatics retrieval of evidence on the relationships between the SNPs and the phenotype.

Trait	Identification of “true SNPs”		MESH terms used in bioinformatics tool
	Data source	Number	
eGFR	Köttgen (2010) [Köttgen et al., 2010]	28	Kidney Diseases; Kidney Failure; Kidney Failure, Acute; Kidney Failure, Chronic; Renal Insufficiency, Acute; Renal Insufficiency, Chronic
Crohn’s Disease	Franke (2010) [Franke et al., 2010]	71	Crohn Disease
Coronary Artery Disease	Schunkert (2011) [Schunkert et al., 2011]; C4D Genetics Consortium (2011) [Coronary Artery Disease (C4D) Genetics Consortium, 2011]	30	Coronary Artery Disease; Angina Pectoris; Coronary Disease; Coronary Restenosis; Coronary Stenosis; Coronary Thrombosis; Myocardial Infarction
Rheumatoid Arthritis	Stahl (2010) [Stahl et al., 2010]	18	Arthritis, Rheumatoid; Arthritis, Juvenile Rheumatoid
Primary Biliary Cirrhosis	Mells (2011) [Mells et al., 2011]	18	Liver Cirrhosis, Biliary
Type 2 Diabetes	Voight (2010) [Voight et al., 2010]	26	Diabetes Mellitus, Type 2
BMI	Speliotes (2010) [Speliotes et al., 2010]	32	Obesity

eGFR: estimated Glomerular Filtration Rate; BMI: Body Mass Index

Table II

Pooled relative probabilities of association across the nine experts for experts' opinion and across the seven traits for empirical evidence. Example of how questions were formulated (Q13): "How many times more likely to be associated with the phenotype is a SNP in a gene ($\pm 5kb$) which is highly expressed in a tissue relevant to the phenotype compared with a SNP which is not?". Frequencies for each type of evidence were calculated across all 7 traits for "True SNPs", and across all 7,000 SNPs for "Random SNPs".

Type of prior knowledge		Frequency		Experts' opinions ¹		Empirical evidence ²	
Questions to Experts (Questionnaire Round 2)	Bioinformatics Query (Changes from Questionnaire)	True SNPs (%)	Random SNPs (%)	Pooled RP (95%CI)	Heterogeneity % I ² (p-value) ³	Pooled RP (95%CI)	Heterogeneity % I ² (p-value) ³
Q1: SNP in a transcribed but not translated region?	No change	58.7	50.6	2.1 (1.6–2.7)	33 (0.150)	1.4 (1.1–1.8)	0 (0.669)
Q2: SNP in a translated region but does not change the amino acid?	No change	1.4	0.6	2.8 (1.9–4.2)	69 (0.001)	4.0 (1.8–8.6)	0 (0.700)
Q3: SNP changes the amino acid but not in a functional protein domain?	No change	3.4	0.4	5.3 (3.2–8.8)	74 (< 0.001)	7.7 (4.2–14.1)	0 (0.854)
Q4: SNP in a functional protein domain?	No change	4.5	0.4	6.4 (4.1–9.8)	61 (0.009)	9.6 (6.1–15.1)	0 (0.822)
Q5: SNP in a regulatory region which is not transcribed?	No change	3.6	3.5	3.5 (2.5–4.9)	41 (0.094)	1.2 (0.7–2.3)	0 (0.935)
Q6: SNP in a transcribed regulatory region?	No change	11.2	5.1	4.9 (3.8–6.4)	7 (0.377)	2.4 (1.6–3.5)	20 (0.281)
Q7: SNP in a genomic region evolutionary conserved in vertebrates?	No change	10.3	1.5	1.8 (1.4–2.3)	11 (0.344)	5.7 (3.8–8.4)	0 (0.888)
Q8: SNP in a gene ($\pm 5kb$) that has been associated with same/closely related phenotype in a meta-analysis or in > 1 study?	SNP in a gene investigated with the phenotype in meta-analysis or > 1 study	33.6	1.3	49.3 (19.7–123.2)	86 (< 0.001)	21.1 ⁴ (16.6–26.8)	64 (0.010)
Q9: SNP in a gene ($\pm 5kb$) that has been associated with same/closely related phenotype in a single study?	SNP in a gene investigated with the phenotype in a single study	8.1	2.4	6.4 (4.7–8.8)	40 (0.104)	2.4 ⁵ (1.4–4.2)	57 (0.030)
Q10: SNP in a gene ($\pm 5kb$) under a linkage peak that has been associated with same/closely related phenotype?	NA	-	-	5.2 (3.7–7.4)	52 (0.032)		
Q11: SNP in a locus within which other SNPs have been associated with the same/closely related phenotype?	SNP in a locus where other SNPs investigated with the phenotype	75.8	46.9	9.8 (4.6–21.0)	84 (< 0.001)	3.5 (2.6–4.7)	43 (0.106)

Type of prior knowledge		Frequency		Experts' opinions ¹		Empirical evidence ²	
Questions to Experts (Questionnaire Round 2)	Bioinformatics Query (Changes from Questionnaire)	True SNPs (%)	Random SNPs (%)	Pooled RP (95%CI)	Heterogeneity % I ² (p-value) ³	Pooled RP (95%CI)	Heterogeneity % I ² (p-value) ³
Q12: SNP in a gene (± 5 kb) that has been associated with same/closely related phenotype in functional models (animal or in-vitro studies)?	SNP in a gene associated with the phenotype in <i>mouse</i> models	1.4	0.1	21.9 (12.1–39.5)	68 (0.002)	9.5 (2.3–38.5)	0 (0.984)
Q13: SNP in a gene (± 5 kb) which is highly expressed in a tissue relevant to the phenotype?	No change	10.8	2.7	3.6 (1.9–6.7)	83 (<0.001)	3.4 ⁶ (2.3–5.0)	52 (0.080)
Q14: SNP in a gene (± 5 kb) which encodes for a protein in a pathway relevant to the phenotype?	No change	26.0	13.3	16.0 (8.3–30.7)	84 (<0.001)	2.1 ⁷ (1.5–2.8)	47 (0.079)
Q15: SNP in a gene (± 5 kb) which shows gene/protein interactions relevant to the phenotype?	SNP in a gene which shows <i>protein-protein</i> interactions relevant to the phenotype	39.9	18.7	10.3 (5.1–20.8)	83 (<0.001)	2.5 (1.9–3.2)	0 (0.527)

¹ Pooled estimate obtained from random-effect meta-analysis;

² Pooled estimate obtained from fixed-effect meta-analysis;

³ P-value from the heterogeneity test (in bold: p-value <0.10);

⁴ Random effects (RE) meta-analysis: 24.1 (15.5–37.4);

⁵ RE meta-analysis: 7.3 (3.5–15.2);

⁶ RE meta-analysis: 3.9 (2.1–7.4);

⁷ RE meta-analysis: 2.4 (1.6–3.6)

RP=relative probability; 95%CI=95% confidence intervals