Learning causal biological networks with the principle of Mendelian randomization

Md. Bahadur Badsha[1*] and Audrey Qiuyan Fu[1*]

[1] Department of Statistical Science, Institute for Bioinformatics and Evolutionary Studies, Center for Modeling Complex Interactions, University of Idaho, Moscow, ID, USA.

*Corresponding author (mdbadsha@uidaho.edu; audreyf@uidaho.edu)

Although large amounts of genomic data are available, it remains a challenge to reliably infer causal relationships among molecular phenotypes (such as gene expression), especially when many phenotypes are involved. We present MRPC, which learns a causal biological network efficiently and robustly from integrating genotype and molecular phenotype data, in which directed edges indicate causal directions. MRPC is the first machine learning algorithm that incorporates the Principle of Mendelian randomization (PMR) in classical algorithms for learning causal graphs in computer science. We demonstrate through simulation that MRPC outperforms existing general-purpose network inference methods and methods using the PMR. We apply MRPC to distinguish direct and indirect targets among multiple genes associated with expression quantitative trait loci (eQTLs). We also construct a causal network for frequently altered cancer genes.

Whereas experiments (e.g., temporal transcription or protein expression assays, gene knockouts or knockdowns) have been conducted to understand the causal relationships among genes[1,2], or between an expression quantitative trait loci (eQTL) and its direct and indirect target genes[3], it remains a challenge to learn causality directly from genomic data. It is even harder to learn (i.e., infer) a causal network, which may represent which genes regulate which other genes. We address this problem in this paper. Correlation (or association) is often used as a proxy of a potentially causal relationship, but similar levels of correlation can arise from different causal mechanisms (Models 1-4 in **Fig. 1a**). For example, between two genes with correlated expression levels, it is plausible that one gene regulates the other gene (Models 1 and 2 in **Fig. 1a**); it is also plausible that they do not regulate each other directly, but both are regulated by a common genetic variant (Model 3 in **Fig. 1a**).

Correlation between the expression, or any molecular phenotype, of two genes is symmetrical – we cannot infer which of the two genes is the regulator and which the target. However, if a genetic variant (e.g., a SNP) is significantly associated with the expression of one of the two genes, then we may assign a directed edge from the variant to the gene, as it is reasonable to assume that the genotype causes changes in the phenotype (expression), not the other way around. This additional, directed edge breaks the symmetry between the two genes, and makes it possible to infer the causal direction (e.g., compare Models 1 and 2 in **Fig. 1a**). This is the rationale behind the Principle of Mendelian Randomization (PMR). The randomization principle in experimental design (e.g., clinical trials) is critical in establishing causality: only when subjects are randomly assigned to different exposures is it possible to draw causal connections between exposure and outcome. As a randomization principle, the PMR assumes that the alleles

of a genetic variant are randomly assigned to individuals in a population, analogous to a natural perturbation experiment and therefore achieving the goal of randomization[4]. The PMR has been widely used in epidemiology studies, where genetic variants are used as instrumental variables to facilitate the estimate of causal effect between a mediator (or exposure, such as gene expression) and an outcome (e.g., a disease phenotype[4]). It received increasing attention in genetics in recent years[5-17]

Large consortia, such as the GEUVADIS consortium[18] and subsequently the GTEx consortium[19], have established the widespread genetic variation (i.e. eQTLs) in human genome that may regulate gene expression, making PMR-based methods increasingly relevant and important for understanding interactions among genes. Furthermore, genome-wide association studies (GWASs) have identified a large number of genetic variants that are potentially causal to diseases[20]. Understanding the roles of these GWAS-significant variants is key to understanding the mechanisms underlying diseases. Interestingly, likely half of the GWAS-significant variants genetic variants are eQTLs[21]. As it becomes more common nowadays to collect gene expression data in disease studies[6,11], studying eQTLs (which may also be GWAS-significant SNPs) and their associated genes provides a powerful approach for a deeper understanding of diseases.

The research of complex diseases often focuses on moderately-sized networks of dozens or hundreds of disease-relevant genes, aiming to identify key regulators and understand the processes involved[6,11]. Being able to accurately reconstruct the causal network of a moderate set of genes will help biologists generate testable hypotheses (e.g., which genes are the key

regulators and may potentially serve as drug targets) and relate these networks to other phenotypes, such as drug response.

However, existing methods adopting the PMR (e.g. the mediation-based methods[12,13], and the MR methods[22]) are not directly applicable to inference of a causal network of gene expression. This is because these methods typically examine the graph of $V_1 \rightarrow T_1 \rightarrow T_2$ (i.e., Model 1 in **Fig. 1a**), where V_1 is the genetic variant, T_1 may represent gene expression, and T_2 a clinical trait. This graph, called the "causal model" by existing PMR-based methods, is sensible when T_2 is a potential outcome of T_1. However, when we examine relationships among gene expression or other molecular phenotypes, it is usually not known beforehand which of T_1 and T_2 is more likely to be the outcome of the other, and Model 1 alone does not have the flexibility of examining other possibilities. As a result, these methods are limited in the causal relationships they can recover. In this paper, we generalize the interpretation of the PMR to account for a variety of causal relationships.

Additionally, applications of the PMR in genomics has not been efficient: existing methods generally work with a small number of nodes, may require spatial (e.g. locations of genetic variants on the genome) or temporal information, or tend to add many false positive edges. There has been some recent effort to address the efficiency issue[13]. Meanwhile, in machine learning, a class of algorithms, such as those based on the classic PC algorithm[23-27], have been developed in over a decade to efficiently learn causal graphs for a large number of nodes. These algorithms typically consist of two main steps (**Fig. 1b**): i) *inferring the graph skeleton* through a series of statistical independence tests. The graph skeleton is the same as the final graph except that the

edges are undirected; and ii) *determining the direction of the edges* in the skeleton. Variants of the original PC algorithm have been developed to reduce the impact of the ordering of the nodes on inference (e.g., the R package pcalg[26,27]), or to reduce the number of statistical tests needed for inferring the skeleton (e.g., the R package bnlearn[24,25]).

Here we develop a new method, namely MRPC, which incorporates the PMR into PC algorithms and learns a causal graph where the nodes are genetic variants and molecular phenotypes (such as gene expression), and where the edges between nodes are undirected or directed, with the direction indicating causality. Crucially, by combining the PMR with machine learning, our method is efficient and accurate. Our extended interpretation of the PMR can be thought of as a way of introducing useful constraints in graph learning and effectively reducing the search space of possible topologies. We demonstrated the performance of MRPC on simulated and real data.

Results

Multiple causal relationships under the Principle of Mendelian Randomization (PMR). We extended the interpretation of the PMR to consider five causal relationships in a triplet of a genetic variant and two phenotypes, including the "causal model" (**Fig. 1a**). Under the assumption that genotype influences phenotype and not the other way around, these five models are mutually exclusive and encompass all possibilities, with Model 0 being the null model where the two phenotype nodes are not related, and the other four models being non-null models. As mentioned in the Introduction, Model 1 ($V_1 \rightarrow T_1 \rightarrow T_2$) is typically referred to as the causal model under standard use of the PMR with T_1 being the exposure (e.g., gene expression) and T_2 being

the outcome (e.g., clinical phenotype). cit[12,28] and findr[13], two existing PMR-based methods for example, both focus on testing Model 1.

However, Model 1 is limited. Among other possible causal relationships, Model 2 ($V_1 \rightarrow T_1 \leftarrow T_2$) defines a v-structure where both edges point to the same node. This model is suitable when no genetic variant is available for T_2 in the data. Model 3 ($V_1 \rightarrow T_1$ and $V_1 \rightarrow T_2$) captures the scenario where T_1 and T_2 are not directly related, but both regulated by V_1. The current interpretation of the PMR in other methods typically rejects these two models in search of the "causal" model (Model 1). However, under our interpretation of the PMR, Models 2 and 3 describe alternative regulatory mechanisms between two genes, and therefore should also be allowed when constructing the network of molecular phenotypes. Model 4 ($V_1 \rightarrow T_1$; $V_1 \rightarrow T_2$; $T_1 - T_2$) refers to the case where the two phenotypes T_1 and T_2 have additional dependence (represented by the undirected edge) on top of that induced by the sharing genetic variant. We consider undirected and bidirected edges to be equivalent for simplicity, in that an undirected can be thought of as an average of two equally likely directions, namely M5 ($V_1 \rightarrow T_1$; $V_1 \rightarrow T_2$; $T_1 \rightarrow T_2$) and M6 ($V_1 \rightarrow T_1$; $V_1 \rightarrow T_2$; $T_1 \leftarrow T_2$); in other words, $M4 = 1/2 \times M5 + 1/2 \times M6$. M5 and M6 are indistinguishable in terms of their dependence relationships (i.e., they are Markov equivalent[29]): all pairs of nodes can be marginally dependent and conditionally dependent given the remaining node. It is plausible that a hidden variable regulates both T_1 and T_2, although we currently do not consider hidden variables in our inference.

MRPC, a novel causal network learning algorithm. Our method, namely MRPC, is a novel causal network inference method for genomic data (**Fig. 1b**; **Supplementary Figs. 1, 2)**. This method analyzes a data matrix with each row being an individual, and each column a genetic

variant or a molecular phenotype. Our method also consists of the two main steps as described above. The first step of learning the graph skeleton is similar to that of other PC algorithms, but with an online control of the false discovery rate (FDR), which is explained in detail below. We incorporated the PMR in the second step of edge orientation (**Fig. 1b; Supplementary Fig. 2**), which involves three scenarios: i) MRPC first identify edges involving the genetic variants and orient these edges to point to the molecular phenotype; ii) MRPC then looks for three nodes with a potential v-structure (e.g., Model 2 in **Fig. 1a**, or among three molecular phenotypes, $T_1 \rightarrow T_2 \leftarrow T_3$). MRPC conducts additional conditional independence tests if no such test has been performed in the first step; and iii) among the remaining edges, MRPC iteratively finds node triplets with only one undirected edge. It examines the results from the independence tests from the first step to identify which of the five basic topologies is consistent with the test results for this triplet. In MRPC, we use Fisher's z transformation for Pearson correlation in all the marginal tests and for the partial correlation in all the conditional tests, consistent with the default method in pcalg (see "Conditional independence tests based on partial correlations" in Methods). However, other parametric or nonparametric tests for marginal and conditional independence tests may be performed in place of Fisher's z transformation test.

Existing network inference algorithms (such as those implemented in R packages pcalg and bnlearn) control the type I error rate for each individual statistical test, but not the family-wise error rate (FWER) or the FDR, as most methods controlling both the FWER and FDR require the knowledge of the total number of tests, which is not known in advance in graph learning. Lack of correction for multiple comparison often leads to too many false edges in the inferred graph, especially when the graph is large (see our simulation results below). We implemented in MRPC

159 the LOND (Levels based on Number of Discoveries) method for controlling the FDR in an
160 online manner[30] (see "Sequential FDR control" in Methods). The LOND method estimates the
161 expected FDR conditioned on the number of tests performed so far and the number of rejections
162 from these tests.
163
164 Furthermore, genomic data may contain outliers[31], which can greatly distort the inferred graph
165 (see our simulation results below). Like pcalg, MRPC uses the correlation matrix, rather than the
166 individual-feature matrix, as input. We implemented in MRPC a method for calculating the
167 robust correlation matrix[31] (see "Calculation of robust correlation" in Methods) in place of
168 Pearson correlation to alleviate the impact of outliers if they are present.
169
170 **MRPC outperforms existing network inference algorithms and PMR-based methods on**
171 **synthetic data in overall accuracy.** We compared MRPC with two popular network inference
172 algorithms: the pc method (implemented in pcalg) and the mmhc method (implemented in
173 bnlearn), and three PMR-based methods, namely cit, findr and QPSO[32]. Except for QPSO, which
174 is implemented in MATLAB, all the methods are implemented in R. We simulated data using
175 linear models for the five basic topologies, three common topologies in biology[33,34] (such as
176 multi-parent, star, and layered graphs), as well as a complex topology with over 20 nodes (**Fig.**
177 **2**). We varied the sample size, as well as the signal strength through the coefficients in the linear
178 models (see "Generating simulated data" in Methods).
179
180 For each topology, we generated 1000 data sets with different combinations of signal strength
181 and sample size, and ran each method with their default parameters. Specifically, we ran MRPC

with FDR=0.05, Pearson correlation ($\beta = 0$; see "Sequential FDR control" in Methods) and the LOND method ($a = 2$; see "Calculation of robust correlation" in Methods). We ran mmhc and pc with the type I error rate being the default value of 0.05. We explained the procedures for running other PMR-based methods in the next section.

We compared the recall and precision (see "Recall and precision" in Methods) across methods (**Fig. 3a**; **Supplementary Figs. 3, 4; Supplementary Tables 1, 2)**. Recall (i.e., power, or sensitivity) measures how many edges from the true graph a method can recover, whereas precisions (i.e., 1-FDR) measures how many correct edges are recovered in the inferred graph. Across different topologies and parameter settings, MRPC has the highest median recall and precision, with both median recall and median precision above 85%. MRPC is followed by mmhc, findr, QPSO, pc, with cit trailing far behind (**Fig. 3a**). MRPC recovers the true graph particularly well at moderate or stronger signal with a medium or larger sample size. For the complex topology, MRPC performs consistently better than pc and mmhc. This is still the case when the signal strength is heterogeneous across the complex topology (see "Simulation under the complex topology with heterogeneous signal strengths" in Methods; **Supplementary Fig. 5**). Examination of inferred graphs from different methods shows that pc is unable to determine edge directions or wrongly identifies v-structures when the true model contains none (**Fig. 3b**; **Supplementary Figs. 6, 7**). PMR-based methods, such as findr and cit, can infer too many or too few edges, whereas QPSO cannot identify the direction correctly. In the presence of outliers, MRPC with robust correlation as input substantially outperforms pc and mmhc (**Supplementary Fig. 8)**.

Existing PMR-based methods cannot deal with complex causal relationships. We examine the performance of PMR-based methods more closely in this section. Since cit and findr focus on Model 1, the topologies they can identify are limited to those that involve primarily Model 1, such as the star graph and the layered graph: the star graph consists of four M1s, and the layered graph five (**Fig. 2**). For method comparison, we limited the true graphs to those that can be analyzed by findr or cit, specifically, M0, M1, M3, star and layered graphs for findr, and M1, star and layered graphs for cit (see "Application of findr and cit" in Methods).

Unlike MRPC, which is agnostic about which genes may be potential regulators and which potential targets, findr and cit are applied to ordered gene pairs iteratively, requiring specification of which of the two genes is the potential regulator and which the target. For example, to test whether the data are simulated under M1, then findr and cit will be performed twice, on (V_1, T_1, T_2) and then on (V_1, T_2, T_1). The number of ordered gene pairs is $2 \times \binom{5}{2} = 20$ for the star graph and $2 \times \binom{7}{2} = 42$ for the layered graph. We applied Bonferroni correction with a familywise type I error rate of 0.05. Take again the star model with a sample size of 1000 for example, where we varied the signal strengths in simulation. Although Bonferroni correction is already a conservative method for multiple testing, findr still sometimes infers more edges than there are (summarized by the lower precision in **Fig 3a**, also see **Supplementary Fig. 7**), whereas cit may infer a very dense graph or no edges at all (summarized by low recall and low precision in **Fig. 3a**; also see **Fig. 3b** and **Supplementary Figs. 6, 7**).

Next, we investigated the performance of findr and cit when the graph skeleton is known, such that the number of tests is reduced to one on simple models (M0, M1 and M3), and to four in the

star graph and to five in the layered graph (**Supplementary Fig. 9**). In other words, potential regulators and targets are known to findr and cit. For MRPC we continued to assume that the skeleton was unknown. With known skeletons, both findr and cit performed similarly to, and in almost all the cases not better than MRPC. The performance of cit can still be much worse than the other two when the signal strength is low or the sample size is small.

We included QPSO in our comparison upon a reviewer's request. Unlike the other five methods discussed here, QPSO takes a graph skeleton as the input and seeks the optimal orientation of the edges, its performance therefore depending heavily on how well the skeleton is inferred. Whereas the authors of QPSO used pc to generate the skeleton, we used MRPC to generate the input, having observed the unsatisfactory performance of pc. As a result, the accuracy of MRPC in identifying the skeleton gives QPSO an advantage in the performance evaluation over other methods. A fair comparison is not to compare QPSO directly with all the other methods, but with MRPC alone. This comparison again shows that QPSO is lacking both in recall and in precision **(Fig. 3a)**. Additionally, QPSO takes much longer time than all the other methods. For example, the runtime is 21 minutes per data set with the complex topology, compared to 2.5 seconds for MRPC, 0.2 seconds for mmhc, and 0.3 seconds for pc (**Supplementary Table 3**). We therefore calculated recall and precision only for 20 (instead of 1000) data sets in simulation for QPSO.

Distinguishing direct and indirect targets of eQTLs. We next applied MRPC to two causal inference problems that are common in biology. First, we are interested in identifying true targets when a single SNP is statistically associated with the expression of multiple genes. Multiple

genes are potential targets often because these genes are physically close to one another on the genome, and the eQTL analysis usually examines the association between one SNP-gene pair at a time, ignoring dependence among genes. Indeed, among eQTLs identified from the GEUVADIS data[18] (i.e., gene expression measured in lymphoblastoid cell lines, or LCLs, of a subset of individuals genotyped in the 1000 Genomes Project), 62 eQTLs discovered under the most stringent criteria have more than one associated gene (see "Analysis of the GEUVADIS data" in Methods). We applied MRPC to each of these eQTLs and their associated genes in the 373 Europeans, and identified 11 types of topologies (**Fig. 4; Supplementary Table 4;** also see comparison with mmhc and pc for some of the eQTL-gene sets in **Supplementary Fig. 11)**. Three of these 11 types are Models 1, 3 and 4 shown in **Fig. 1a**. Seven other topologies are identified for eight eQTLs each with three associated genes (**Supplementary Table 4)**.

Although the multiple associated genes of the same eQTL are physically near one another, our method showed promise in teasing apart the different dependence (or regulatory relationships) among these genes. For example, the SNP rs479844 (chr11:65,784,486; GRCh38), one of the 62 eQTLs, turns out to be significant in at least three GWASs for atopic march and more specifically, atopic dermatitis (p values ranging from 10^{-10} to 10^{-18})[20, 35-37]. This SNP has been linked with two genes, AP5B1 (chr11:65,775,893-65,780,802) and OVOL1 (chr11:65,787,022-65,797,219), in these GWASs, but it is unclear which is the real target. Our MRPC infers Model 1 for the triplet: rs479844→OVOL1→AP5B1 (**Fig. 4a**), which suggests that OVOL1 is more likely to be the direct target, and AP5B1 the indirect one. Meanwhile, for HLA-DQA1 (chr6:32,637,403-32,654,846) and HLA-DQB1 (chr6:32,659,464-32,666,689), both genes are associated with the SNP rs9274660 and located in the major histocompatibility (MHC) region of

high linkage disequilibrium. As expected, MRPC infers an undirected edge between the two genes, as the information on the two genes is highly symmetric in the genotype and gene expression data. By contrast, mmhc and pc often misspecify edges or their directions (**Supplementary Fig. 11**). We focused on the European sample in this analysis, as the sample size of the Africans is small (89). However, we managed to replicate part of the topologies for the few eQTLs discovered in both populations (see "Analysis of the GEUVADIS data" in Methods).

Since the GTEx consortium[19] contains data also from LCLs, we next examined whether the causal relationships inferred from the GEUVADIS data may be replicated in the LCL samples from GTEx (see "Analysis of the GTEx data" in Methods). The sample size of 117 is much smaller in GTEx, though, which reduces the expected number of causal relationships to be replicated. We therefore focus on eQTL-gene sets that were inferred to have an M1 model in GEUVADIS by MRPC. We ran MRPC, findr and cit on the 16 eQTL-gene sets with an M1 model that have the genotype and gene expression data in both GEUVADIS and GTEx LCL samples. findr replicated 9 sets, MRPC 8 and cit only 1 (**Supplementary Table 5**). This result is consistent with simulation results (**Fig. 3a**): whether the graph skeleton is known or not, MRPC and findr have similar performance on M1 across different sample sizes and signal strengths, both much better than cit. In particular, we replicated the relationship rs479844→OVOL1→AP5B1 with both MRPC and findr in the GTEx LCL samples.

Construction of a causal network for frequently altered cancer genes. In the second application, we applied MRPC to the genomic data of breast cancer patients from the TCGA

297 consortium[38] (see "Analysis of the TCGA breast cancer data" in Methods), aiming to learn the
298 causal gene regulatory network for frequently altered genes in breast cancer; that is, genes with
299 point mutations, differential expression and different copy number in a large percentage of
300 patients[39]. Although these genes have been shown in our previous work to form a dense network
301 of epistasis and are known to be involved in many pathways known to play an important role in
302 cancer (e.g., DNA damage repair pathways[40], kinase signaling pathways[41], and so on)[39], how
303 these genes regulate one another transcriptionally remains unclear.
304
305 Copy number variation usually has a much stronger effect on gene expression than SNPs do in
306 breast cancer[42]. We therefore used the copy number variation as the genotype, and gene
307 expression as the molecular phenotype. Similar to an earlier investigation[39], we extracted 85
308 frequently altered genes (e.g., BRCA1, BRCA2, TP53, etc.) in breast cancer and their copy
309 number variation data. We calculated the Pearson correlation matrix (**Fig. 5a**), and applied
310 MRPC at FDR=0.05 (**Fig. 5b**), and subsequently at 0.01, 0.10 and 0.15. The inferred graphs
311 appeared reasonably stable: each graph contains around 200 edges; when the FDR changes by
312 0.05, the number of edges inferred differently tends to be around 10, which is roughly 1/20 of all
313 the edges (**Fig. 5c**; **Supplementary Figs. 12-14**); this is consistent with the change of 0.05 in
314 FDR, as this rate implies that on average roughly 5% of all the edges are likely to be false
315 positives and therefore would not be consistently inferred at another FDR. In other words, most
316 of the edges are inferred reliably across different FDR levels.
317
318 We searched through literature for evidence supporting the inferred directed edges here
319 (**Supplementary Table 6**). Existing literature is rich on the undirected relationships among

genes, but scarce on directed ones. Among 49 directed edges inferred here, we found literature support for only seven directed relationships inferred here (including one edge with the reference indicating an opposite direction), and for undirected relationships of 14 inferred gene pairs (**Supplementary Table 6**). Considering the thousands of pairwise relationships that may exist among these genes, the causal graph we have inferred here provides a list of plausible regulatory relationships and helps prioritize which genes to perform knockdown experiments on.

In the graph inferred at FDR=5%, one gene (MAML2) has three targets (NFIB, MET, and PIK3R1), followed by nine genes (ATM, CANT1, ELK4, ERCC4, IL6ST, KMT2C, KMT2E, MAP3K1, and MET) with two targets, 31 genes with one target, and 44 without targets (**Fig. 5b**). We then applied WGCNA[43] to help visualize the inferred graphs through grouping nodes into modules. We experimented with several module sizes, and in the end divided the graph into modules with at least seven nodes (including four genes) per module, such that all the visibly large clusters of nodes were represented (**Fig. 5b**; **Supplementary Fig. 15**; **Supplementary Table 7**). Genes have higher connectivity within the module than with other modules, although most modules have edges connecting one another, consistent with the notion that multiple biological pathways are involved, and possibly interacting in cancer[39] (**Fig. 5b**). We ran gene ontology (GO) enrichment analysis[44] on the genes in each module (excluding the grey nodes, which are not allocated to any module). Except for the green module, which contains only four genes, each module is significantly associated with distinct biological processes or PANTHER pathways[45], suggesting that the causal network we learned has a structure consistent with the underlying biological functions (**Fig. 5b**; **Supplementary Tables 8, 9**).

Additionally, our causal inference distinguishes "direct" from "indirect" correlation. For example, following hierarchical clustering, the correlation heatmap indicates that NF1, ERCC4, and TRIP11 have higher correlation with one another and are therefore clustered together (**Fig. 5a**). However, no edge connects NF1 and ERCC4 at any of the FDRs we examined (**Fig. 5b**; **Supplementary Figs. 12-14**). A closer look at the conditional independence tests showed that the strong correlation between these two genes is explained away by TRIP11, BRCA1 and the copy number variation of NF1 (p value: 3e-6; the significance threshold by LOND method: 5e-7). In other words, the correlation between NF1 and ERCC4 is indirect: it is induced by association with three other nodes. Indeed, there is no interaction between NF1 and ERCC4 in the literature to the best of our knowledge. Instead, NF1 has been shown to interact with the KMT2 family[40], also shown in our inferred network (**Fig. 5b**), whereas the DNA repair gene ERCC4 is recently shown to be involved in the translesion DNA synthesis together with TRIP11 and other genes[46], consistent with the edge between these two genes in our inferred network (**Fig. 5b**).

Discussion

In summary, we have developed MRPC to infer causal networks, which can be an NP-complete problem (see "Properties of MRPC" in Methods). Our MRPC method examines a variety of causal relationships implied by the PMR, and takes advantage of the development of machine learning algorithms for causal graph inference. MPRC integrates genotypes with molecular phenotypes, and can efficiently and accurately learn causal networks. Our method is flexible as it requires only the genotype data (SNPs or other types of variants; see "Multiple genetic variants of the same phenotype" in Methods) and the molecular phenotype measurements (gene expression, or other functional data, such as exon expression, RNA editing, DNA methylation,

366 etc.), and can be applied to a wide range of causal inference problems. Our method is also
367 nonparametric in that no explicit distributions are assumed for the underlying graph. MRPC uses
368 individual-level genomic data to learn plausible biological mechanisms from combining
369 genotype and molecular phenotypes.
370
371 The key improvements in MRPC over existing methods are i) implementation of the online FDR
372 control method (the LOND method), which helps reduce false positives. As our simulation
373 demonstrated, false positive edges are a severe problem in other methods, whether they are based
374 on the PMR or not; and ii) accounting for all possible causal relationships a triplet with a genetic
375 variant can have under the PMR. This extended interpretation of the PMR allows MRPC to go
376 beyond the typical "causal model" examined by other PMR-based methods and can deal with
377 networks of realistic causal relationships. Computationally, incorporation of the PMR puts
378 constraints on the space of possible graphs and allows for efficient search of graphs consistent
379 with the data (see "MRPC and other PMR-based methods" in Methods).
380
381 Here, we demonstrated the outstanding performance of MRPC on small to moderately-sized
382 graphs. Additional work is needed to extend the ability of MRPC to larger graphs. For example,
383 median precision of 90% in simulation (Fig. 3a) translates to the actual FDR being 10%, above
384 the expected level of 5%. We also applied MRPC to the directed networks of 1000 genes and
385 1000 genetic variants simulated in the DREAM5 Systems Genetics Challenge A[47], with FDR
386 being 30% (see "Analysis of the DREAM5 Systems Genetic Challenge A data" in Methods;
387 **Supplementary Table 10**). With a sample size of 999, our MRPC identified only 19% of 2048
388 true edges, although precision is 67%, meaning that the actual FDR is 33%, comparable to the

389 expected threshold of 30%. By contrast, the TRANSWESD method[48] (published after the
390 challenge and showed better performance than the best participating method for this challenge)
391 recovered 68% of the true edges, but had low precision of only 34%% (mmhc, the second best
392 performing method compared in this paper, has a similar issue). This suggests that although
393 controlling the FDR to certain extent, the LOND method can discard too many true positives.
394
395 Like most causal graph learning methods, a key assumption behind MRPC is that there are no
396 hidden nodes that are connected to the observed nodes in the graph. Whereas this assumption
397 may not hold in biology, we can take additional measures to alleviate the impact of hidden nodes.
398 For example, genes are often grouped in clusters that tend to have higher correlation within the
399 cluster. Our method can be applied to genes within a gene cluster to build the detailed causal
400 network. As the next step, we are working on extensions of MRPC to deal with hidden variables[7,
401 26]. The current version of MRPC has already demonstrated its power in tackling several
402 biological problems on causality and in integrating large amounts of genomic data.
403

404 1. Segal E*, et al.* Module networks: identifying regulatory modules and their condition-
405 specific regulators from gene expression data. *Nat Genet* **34**, 166-176 (2003).
406 2. Housden BE*, et al.* Transcriptional dynamics elicited by a short pulse of notch activation
407 involves feed-forward regulation by E(spl)/Hes genes. *PLoS Genet* **9**, e1003162 (2013).
408 3. Cheung VG, Spielman RS. Genetics of human gene expression: mapping DNA variants
409 that influence gene expression. *Nat Rev Genet* **10**, 595-604 (2009).
410 4. Davey Smith G, Hemani G. Mendelian randomization: genetic anchors for causal
411 inference in epidemiological studies. *Hum Mol Genet* **23**, R89-98 (2014).

5. Hill SM, *et al.* Inferring causal molecular networks: empirical assessment through a community-based effort. *Nat Methods* **13**, 310-318 (2016).

6. Zhang B, *et al.* Integrated systems approach identifies genetic nodes and networks in late-onset Alzheimer's disease. *Cell* **153**, 707-720 (2013).

7. Chaibub Neto E, *et al.* Modeling causality for pairs of phenotypes in system genetics. *Genetics* **193**, 1003-1013 (2013).

8. Gutierrez-Arcelus M, *et al.* Passive and active DNA methylation and the interplay with genetic variation in gene regulation. *Elife* **2**, e00523 (2013).

9. Oren Y, Nachshon A, Frishberg A, Wilentzik R, Gat-Viks I. Linking traits based on their shared molecular mechanisms. *Elife* **4**, e04346 (2015).

10. Stojnic R, Fu AQ, Adryan B. A graphical modelling approach to the dissection of highly correlated transcription factor binding site profiles. *PLoS Comput Biol* **8**, e1002725 (2012).

11. Franzen O, *et al.* Cardiometabolic risk loci share downstream cis- and trans-gene regulation across tissues and diseases. *Science* **353**, 827-830 (2016).

12. Millstein J, Zhang B, Zhu J, Schadt EE. Disentangling molecular relationships with a causal inference test. *BMC genetics* **10**, 23 (2009).

13. Wang L, Michoel T. Efficient and accurate causal inference with hidden confounders from genome-transcriptome variation data. *PLoS Comput Biol* **13**, e1005703 (2017).

14. Zhu J, *et al.* Stitching together multiple data dimensions reveals interacting metabolomic and transcriptomic networks that modulate cell regulation. *PLoS Biol* **10**, e1001301 (2012).

15. Huang YT, Liang L, Moffatt MF, Cookson WO, Lin X. iGWAS: Integrative genome-wide association studies of genetic and genomic data for disease susceptibility using mediation analysis. *Genet Epidemiol* **39**, 347- 356 (2015).

16. Yang F, Wang J, The GTEx Consortium, Pierce BL, Chen LS. Identifying cis-mediators for trans-eQTLs across many human tissues using genomic mediation analysis. *Genome Res* **27**, 1859-1871 (2017).

17. Connor LJ, Price AL. Distinguishing genetic correlation from causation across 52 diseases and complex traits. *bioRxiv*, 205435 (2017).

18. Lappalainen T, *et al.* Transcriptome and genome sequencing uncovers functional variation in humans. *Nature* **501**, 506-511 (2013).

19. The GTEx Consortium. Genetic effects on gene expression across human tissues. *Nature* **550**, 204-213 (2017).

20. MacArthur J, *et al.* The new NHGRI-EBI Catalog of published genome-wide association studies (GWAS Catalog). *Nucleic Acids Res* **45**, D896-d901 (2017).

21. Nicolae DL, Gamazon E, Zhang W, Duan S, Dolan ME, Cox NJ. Trait-associated SNPs are more likely to be eQTLs: annotation to enhance discovery from GWAS. *PLoS Genet* **6**, e1000888 (2010).

22. Hemani G, Tilling K, Davey Smith G. Orienting the causal relationship between imprecisely measured traits using GWAS summary data. *PLoS Genet* **13**, e1007081 (2017).

23. Spirtes P, Glymour C, Scheines R. *Causation, Prediction, and Search*. The MIT Press, (2000).

24. Tsamardinos I, Brown LE, Aliferis CF. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning* **65**, 31-78 (2006).

25. Scutari M. Learning Bayesian networks with the bnlearn R package. *J. Stat. Softw.* **35**, 22 (2010).

26. Kalisch M, Mächler M, Colombo D, Maathuis MH, Bühlmann P. Causal inference using graphical models with the R package pcalg. *J. Stat. Softw.* **47**, 26 (2012).

27. Colombo D, Maathuis MH. Order-independent constraint-based causal structure learning. *J. Mach. Learn. Res.* **15**, 3741-3782 (2014).

28. Millstein J, Chen GK, Breton CV. cit: hypothesis testing software for mediation analysis in genomic applications. *Bioinformatics* **32**, 2364-2365 (2016).

29. Richardson T. A characterization of Markov equivalence for directed cyclic graphs. *Int J of Approx Reason* **17**, 107-162 (1997).

30. Javanmard A, Montanari A. On online control of false discovery rate. *arXiv* 150206197v2, (2015).

31. Badsha MB, Mollah MN, Jahan N, Kurata H. Robust complementary hierarchical clustering for gene expression data analysis by beta-divergence. *J Biosci Bioeng* **116**, 397-407 (2013).

32. Wang H, van Eeuwijk FA. A new method to infer causal phenotype networks using QTL and phenotypic information. *PLoS One* **9**, e103997 (2014).

33. Alon U. Network motifs: theory and experimental approaches. *Nat Rev Genet* **8**, 450-461 (2007).

34. Hunter T. Signaling--2000 and beyond. *Cell* **100**, 113-127 (2000).

35. Marenholz, I., Esparza-Gordillo, J., Rüschendorf, F., Bauerfeind, A., Strachan, D.P., Spycher, B.D., Baurecht, H., Margaritte-Jeannin, P., Sääf, A., Kerkhof, M., Ege, M. Meta-analysis identifies seven susceptibility loci involved in the atopic march. *Nat Commun* **6**, 8804 (2015).

36. Paternoster L, *et al.* Multi-ancestry genome-wide association study of 21,000 cases and 95,000 controls identifies new risk loci for atopic dermatitis. *Nat Genet* **47**, 1449-1456 (2015).

37. Paternoster L, *et al.* Meta-analysis of genome-wide association studies identifies three new risk loci for atopic dermatitis. *Nat Genet* **44**, 187-192 (2011).

38. The Cancer Genome Atlas Network. Comprehensive molecular portraits of human breast tumours. *Nature* **490**, 61-70 (2012).

39. Wang X, Fu AQ, McNerney ME, White KP. Widespread genetic epistasis among cancer genes. *Nat Commun* **5**, 4828 (2014).

40. Rao RC, Dou Y. Hijacked in cancer: the KMT2 (MLL) family of methyltransferases. *Nat Rev Cancer* **15**, 334-346 (2015).

41. Martini M, De Santis MC, Braccini L, Gulluni F, Hirsch E. PI3K/AKT signaling pathway and cancer: an updated review. *Ann Med* **46**, 372-383 (2014).

42. Li, Q., Seo, J.H., Stranger, B., McKenna, A., Pe'er, I., LaFramboise, T., Brown, M., Tyekucheva, S., Freedman, M.L. Integrative eQTL-based analyses reveal the biology of breast cancer risk loci. *Cell* **152**, 633-641 (2013).

43. Langfelder P, Horvath S. WGCNA: an R package for weighted correlation network analysis. *BMC Bioinformatics* **9**, 559 (2008).

500 44. Gene Ontology Consortium. Gene Ontology Consortium: going forward. *Nucleic Acids*
501 *Res* **43**, D1049-1056 (2015).
502 45. Mi H, Muruganujan A, Thomas PD. PANTHER in 2013: modeling the evolution of gene
503 function, and other gene attributes, in the context of phylogenetic trees. *Nucleic Acids Res*
504 **41**, D377-386 (2013).
505 46. Ziv O*, et al.* Identification of novel DNA-damage tolerance genes reveals regulation of
506 translesion DNA synthesis by nucleophosmin. *Nat Commun* **5**, 5437 (2014).
507 47. Cokelaer T*, et al.* DREAMTools: a Python package for scoring collaborative challenges
508 [version 1; referees: 3 approved with reservations]. *F1000Research*, 4:1030 (2016).
509 48. Flassig RJ, Heise S, Sundmacher K, Klamt S. An effective framework for reconstructing
510 gene regulatory networks from genetical genomics data. *Bioinformatics* **29**, 246-254
511 (2013).

512

513

514 **METHODS**

515 **Conditional independence tests based on partial correlations.** We use the same method (and
516 R functions) as that used in the R package pcalg[26] for conducting conditional independence tests
517 based on partial correlations. Consider testing conditional independence between variables x
518 and y conditioned on a set of variables S. From the correlation matrix, one may estimate the
519 partial correlations using an iterative approach[26]. Then application of Fisher's z transformation
520 gives the test statistic

$$T = \frac{\sqrt{n - |S| - 3}}{2} \log\left(\frac{1 + \hat{r}_{x,y|S}}{1 - \hat{r}_{x,y|S}}\right),$$

521 which follows N(0,1) under the null hypothesis of conditional independence[49]. In the expression
522 above, $\hat{r}_{x,y|S}$ is the estimated partial correlation, n the sample size, and $|S|$ the number of
523 variables in the set S.
524
525 **Calculation of robust correlation.** We implemented the method in Badsha et al.[31] to calculate
526 the robust correlation matrix as the input to the MRPC inference. Specifically, for data that are
527 approximately normal (usually after preprocessing of the data), we calculated iteratively the
528 robust mean vector $\boldsymbol{\mu}$ and the robust covariance matrix $\boldsymbol{v}$ until convergence. At the t+1st iteration,

529 $$\boldsymbol{\mu}_{t+1} = \frac{\sum_{i=1}^{n}[\varphi_\beta(\boldsymbol{x}_i;\boldsymbol{\mu}_t,\boldsymbol{v}_t)\boldsymbol{x}_i]}{\sum_{i=1}^{n}\varphi_\beta(\boldsymbol{x}_i;\boldsymbol{\mu}_t,\boldsymbol{v}_t)} \quad (1)$$

530 and

531 $$\boldsymbol{v}_{t+1} = \frac{\sum_{i=1}^{n}[\varphi_\beta(\boldsymbol{x}_i;\boldsymbol{\mu}_t,\boldsymbol{v}_t)(\boldsymbol{x}_i-\boldsymbol{\mu}_t)(\boldsymbol{x}_i-\boldsymbol{\mu}_t)^T]}{(1+\beta)^{-1}\sum_{i=1}^{n}\varphi_\beta(\boldsymbol{x}_i;\boldsymbol{\mu}_t,\boldsymbol{v}_t)}, \quad (2)$$

532 where,

533 $$\varphi_\beta(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{v}) = exp\left(-\frac{\beta}{2}(\boldsymbol{x}-\boldsymbol{\mu})^T\boldsymbol{v}^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right). \quad (3)$$

534 In the equations above, $\boldsymbol{x}_i$ is the vector of gene expression in the ith sample, n the sample size,
535 and β the tuning parameter. Equation (3) downweighs the outliers through β, which takes values
536 in [0,1]. Larger β leads to smaller weights on the outliers. When $\beta = 0$, equation (2) is similar to
537 the standard definition of the variance, except that the scalar is 1/n, whereas the unbiased
538 estimator of the variance has a scalar of 1/(n-1). When the data matrix contains missing values,
539 we perform imputation using the R package mice[50]. Alternatively, one may impute the data using
540 other appropriate methods, and calculate the correlation matrix as the input for MRPC.
541

When analyzing simulated data with no outliers, we set $\beta = 0$, which is close to Pearson correlation. We set $\beta = 0.005$ if outliers were included in simulation. On real data, we would usually perform two analyses with $\beta = 0$ and $\beta = 0.005$. These two values did not lead to different results in most cases. See details in "Analysis of the GEUVADIS data" in Methods, which refers to **Supplementary Figures 16-18**.

Sequential FDR control. We implemented the LOND algorithm that control FDR in an online manner, as we did not know the number of tests beforehand in learning the causal graph. Specifically, consider a sequence of null hypotheses (marginal or conditional independence between two molecular phenotypes) $H(m) = H_1, H_2, H_3, \dots, H_m$, with corresponding p-values $p(m) = p_1, p_2, p_3, \dots, p_m$. The LOND algorithm aims to determine a sequence of significance level α_i, such that the decision for the ith test is

$$R_i = \begin{cases} 1, & \text{if } p_i \leq \alpha_i \quad \text{(reject } H_i\text{)} \\ 0, & \text{if } p_i > \alpha_i \quad \text{(accept } H_i\text{)} \end{cases}.$$

The number of rejections over m tests is then

$$D_{(m)} = \sum_{i=1}^{m} R_i.$$

For the overall FDR to be δ, the significance level α_i is set to be

$$\alpha_i = \delta_i \left[D_{(i-1)} + 1 \right],$$

where the FDR for the ith test is

$$\delta_i = \frac{C}{i^a},$$

such that

$$\sum_{i=1}^{\infty} \delta_i = \delta,$$

for integer a>1 and a constant c. We choose a nonnegative sequence δ_i, such that $\sum_{i=1}^{\infty} \delta_i = FDR$. The default value for a is set to 2 in MRPC. At an FDR of 0.05 and $a = 2$, we have

$$\sum_{i=1}^{\infty} \delta_i = \sum_{i=1}^{\infty} \frac{c}{i^2} = c \sum_{i=1}^{\infty} \frac{1}{i^2} = \frac{c\pi^2}{6} = 0.05.$$

Then

$$c = \frac{6*0.05}{\pi^2} = 0.0304.$$

Values of δ_i and α_i for the first 18 tests of analysis of a simulated data set are listed in an example given in **Supplementary Table 11.** The larger a is, the more conservative the LOND method, which means that fewer rejections will be made. We therefore used $a = 2$ throughout simulation and real data analyses. Simulation results in the Results section show that this choice of a works reasonably well for small and moderately-sized networks, although it can lead to exclusion of many true edges in large networks (see "Analysis of the DREAM5 Systems Genetic Challenge A data" in Methods).

Multiple genetic variants of the same phenotype. MRPC currently does not directly deal with multiple genetic variants associated with the same molecular phenotype. For network inference, we recommend using the variant with the strongest association, or merging the multiple variants to create a haplotype variant with the haplotypes being the new genotypes (e.g., two SNPs in linkage disequilibrium, each having three genotypes, can be merged into one variant with genotypes 00, 01, 02, 10, 11, 12, 20, 21, and 22).

Generating simulated data. We generated synthetic data for a variety of graphs, which fall into three categories depending on the complexity (**Fig. 2a**): i) basic topologies of a triplet; ii) topologies common in biological networks: star (i.e., one molecular phenotype has multiple

586 targets); multi-parent (i.e., one molecular phenotype has multiple regulators apart from the
587 genetic variants); and layered; and iii) a complex topology.
588
589 In each topology, we simulated the data first for the nodes without parents, and then for other
590 nodes. Genetic variants are nodes without parents, and we assume them to be biallelic SNPs with
591 three genotypes 0, 1, and 2. Denote the minor allele frequency by q and assume Hardy-Weinberg
592 equilibrium. Then the genotype of the ith variant, V_i , follows a multinomial distribution:
593 $$\Pr(V_i = 0) = (1-q)^2, \Pr(V_i = 1) = 2q(1-q), \Pr(V_i = 2) = q^2.$$
594 Denote the jth molecular phenotype by T_j and the set of its parent nodes by P, which may be
595 empty, or may include variant nodes or nodes of other molecular phenotypes. We assume that
596 the molecular phenotype T_j follows a normal distribution
597 $$T_j \sim N\left(\gamma_0 + \sum_{k\in \mathrm{P}} \gamma_k V_k + \sum_{l\in \mathrm{P}} \gamma_l T_l\,, \sigma_j^2\right).$$
598 The variance may be different for different nodes. For simplicity, we use the same value for all
599 the nodes.
600
601 We treat undirected edges as bidirected edges and interpret such an edge as an average of the two
602 directions with equal weights. For example, for the undirected edge in Model 4 in **Fig. 1a**, we
603 generate data for $T_1 \rightarrow T_2$:
604 $$T_1 \sim N\left(\gamma_0 + \gamma_1 V, \sigma_1^2\right);\ T_2 \sim N\left(\gamma_0 + \gamma_1 V + \gamma_2 T_1, \sigma_2^2\right),$$
605 and separately for $T1 \leftarrow T2$:
606 $$T_1 \sim N\left(\gamma_0 + \gamma_1 V + \gamma_2 T_2,\ \sigma_1^2\right);\ T_2 \sim N\left(\gamma_0 + \gamma_1 V,\ \sigma_2^2\right).$$
607 We then randomly choose a pair of values with 50:50 probability for each sample.
608

609 For simplicity in simulation, we set $\gamma_0 = 0$ and all the other γ's to take the same value, which
610 reflects the strength of the association signal. We considered three values for the slopes: 0.2
611 (weak signal), 0.5 (moderate signal), and 1.0 (strong signal). We also varied the sample size: 50
612 (very small), 200 (small), 500 (medium), and 1000 (large). Thus, we considered twelve
613 combinations of signal strength and sample size (**Supplementary Figs. 3, 4; Supplementary**
614 **Tables 1, 2**).
615
616 Under each combination, we generated 1000 data sets for each topology. For each data set, we
617 shuffled the columns corresponding to gene expression to generate one data set with those
618 columns reordered; if an inference method is sensitive to the ordering of the columns, the
619 inferred graph would have a large variance across data sets. We then applied each method to a
620 data set with permuted columns. To summarize the results, we computed the mean and standard
621 deviation of recall and precision (see "Recall and precision" in Methods) across 1000 data sets
622 for each method, and displayed the mean as the bar and the standard deviation as the error bar in
623 the horizontal bar plots (**Supplementary Figs. 3, 4**). The median of recall and precision of each
624 method across all topologies and all parameter settings are displayed in **Fig. 3a**. We also
625 summarize the median standard deviation of recall and precision in **Supplementary Fig. 16**.
626 Note that the standard deviation in recall and precision reflects variation due to both different
627 data sets and different node orderings. Except for QPSO, the methods under comparison do not
628 differ much in variation. QPSO had a larger variation because only 20 data sets were used for
629 assessing the performance (due to long runtime).
630

631 **Simulation under the complex topology with heterogeneous signal strengths.** The simulation
632 strategy described above assumes the same signal strength (value of γ, the coefficient of the
633 parent node) across the network, which allows us to examine the performance of the methods in
634 simple and well-controlled settings. For the complex strategy, we further allowed the values of γ
635 to vary when generating data for each node. Each γ has equal probability of taking on one of
636 three values: 0.2, 0.5 and 1.0. Similar to the procedure described above, we also generated 1000
637 data sets with this strategy, applied relevant methods, and computed recall and precision.
638
639 **Recall and precision.** Under the standard definition:
640 Recall = (# edges correctly identified in inferred graph) / (# edges in true graph);
641 Precision = (# edges correctly identified in inferred graph) / (# edges in inferred graph).
642 However, we consider it more important to be able to identify the presence of an edge than to
643 also get the direct correct. Therefore, we assign 1 to an edge with the correct direction and 0.5 to
644 an edge with the wrong direction or no direction. For example, when the true graph is
645 $V \rightarrow T_1 \rightarrow T_2$ with 2 true edges, and the inferred graphs are i) $V \rightarrow T_1 \rightarrow T_2$, and $V \rightarrow T_2$; ii)
646 $V \rightarrow T_1 - T_2$; and iii) $V \rightarrow T_1 \leftarrow T_2$, the number of correctly identified edges is then 2, 1.5 and 1.5,
647 respectively. Recall is calculated to be 2/2=100%, 1.5/2=75%, and 1.5/2=75%, respectively,
648 whereas precision is 2/3=67%, 1.5/2=75%, and 1.5/2=75%, respectively.
649
650 When analyzing the complex topology in simulation, which involves correlated genetic variants,
651 we ignored the edges among genetic variants in the calculation of recall and precision, since
652 mmhc and pc are not designed to infer the relationships among genetic variants correctly. findr
653 and cit are not applicable to this topology, and QPSO requires the graph skeleton as the input,

654 with the graph skeleton already specifying the relationship among genetic variants and between
655 variants and their associated genes.
656
657 **Application of findr and cit.** Unlike mmhc and pc that learn the graph skeleton first and orient
658 the edges next, findr and cit test for directed edges in a single step for a triplet of nodes (the
659 genetic variant and two gene expression nodes). This means that in order to learn the topology,
660 we needed to examine all possible gene pairs (e.g., T_1 and T_2; and T_2 and T_1) and then apply
661 findr or cit to the triplet of each of the gene pairs and the genetic variant. Based on the hypothesis
662 testing result from findr or cit, if there was evidence for a directed edge between two nodes, we
663 added 1 to the current value in the adjacency matrix for those two nodes. Otherwise we left the
664 value unchanged. After examining all gene pairs, we converted all positive values in the
665 adjacency matrix to 1 to represent a directed edge. This way, no edges inferred would be
666 eliminated in later tests. We then calculated the aSHD between the inferred adjacency matrix and
667 that of the true graph, and averaged the aSHDs across simulated data sets.
668
669 Although findr aims to compute a causality probability for a triplet, its current implementation
670 for this calculation cannot be applied to small graphs, or cases where multiple genes share the
671 same eQTL and where some of the genes do not have eQTLs. We therefore used the function
672 findr.pijs_gassist_pv() from the R package findr to conduct five hypothesis tests (the p values
673 from these five tests are then converted to a causality probability) for each ordered gene pair with
674 the genetic variant. Consider a triplet V_1, T_1 and T_2. The null (H_0) and alternative (H_a)
675 hypotheses of these five tests are:
676 Test #1: H_0: V_1 and T_1 independent; H_a: $V_1 \rightarrow T_1$;

677 Test #2: H_0: V_1 and T_2 independent; H_a: $V_1 \rightarrow T_2$;

678 Test #3: H_0 (M1): $V_1 \rightarrow T_1 \rightarrow T_2$; H_a: $V_1 \rightarrow T_1$, $V_1 \rightarrow T_2$, $T_1 \rightarrow T_2$;

679 Test #4: H_0 (M0): $V_1 \rightarrow T_1$, both independent of T_2; H_a: $V_1 \rightarrow T_1$, $V_1 \rightarrow T_2$, $T_1 \rightarrow T_2$;

680 Test #5: H_0 (M3): $V_1 \rightarrow T_1$, $V_1 \rightarrow T_2$; H_a: $V_1 \rightarrow T_1$, $V_1 \rightarrow T_2$, $T_1 \rightarrow T_2$.

681 We extract the p values (i.e., $p_i, i = 1, \dots, 5$) for the five tests. The data supports M0, if p_1 is less
682 than, and p_2 and p_4 greater than a certain threshold. The data supports M1, if p_1 is less than, and
683 p_3 greater than a certain threshold. The data supports M3, if p_1 and p_2 are less than, and p_5
684 greater than a certain threshold. We determine the p value threshold with Bonferroni correction,
685 dividing the unadjusted p value 0.05 by $5m$, where m is the total number of genes pairs, because
686 each findr test contains five tests.

687

688 cit generates an omnibus p value for testing whether the triplet follows M1. We used the function
689 cit.cp() from the R package cit for calculation of the omnibus p value. Similarly, we determine
690 the p value threshold also with Bonferroni correction (unadjusted p value 0.05 divided by the
691 total number of genes pairs).

692

693 **Analysis of the GEUVADIS data.** The GEUVADIS project
694 (http://www.ebi.ac.uk/Tools/geuvadis-das/) performed RNA-seq (gene expression) on 373
695 Europeans and 89 Africans from the 1000 Genomes Project. The GEUVADIS project combined
696 the gene expression data with the genotype data, and identified eQTLs across the human genome.
697 Among the most stringent set of eQTLs, 62 have more than one target gene. We extracted the
698 genotypes of these eQTLs and the expression of the target genes in the 373 Europeans, and
699 applied MRPC to each eQTL with its target genes.

The SNP rs479844, which has GWAS significance, is identified in the European sample to be the best eQTL for genes OVOL1 and AP5B1. However, this SNP is not identified to be the eQTL of any gene in the African sample. No eQTLs are reported for these two genes in the African sample. When we further examined the correlation matrices (**Supplementary Figure 17**) between the SNP genotype and expression of the two genes from the two samples, they have qualitative differences: whereas the eQTL has a much stronger correlation with OVOL1 than with AP5B1 in Europeans, it is the reverse in Africans. However, these differences are likely due to the small sample size of the African sample. We therefore do not seek to replicate with the African sample the topology we identified in the European sample.

Also because of the small sample size of the African sample, eQTLs and genes identified to have eQTLs are very different in the two populations. In order to examine whether it is possible at all to replicate the causal network inference from the European sample, we focused on the five top eQTLs identified in both samples: namely, esv2658282, esv2676246, rs11305802, rs230326, and rs7663027. The pairwise correlation matrices (**Supplementary Figure 18**) for each eQTL in the two samples are largely similar. However, due to the difference in the sample size, the topology inferred from the African sample is usually part (**Supplementary Figure 19**) of that from the European sample.

Analysis of the GTEx data. The GTEx consortium has profiled genotypes and gene expression levels in 53 tissues across 714 donors (Release V7, dbGaP Accession phs000424.v7.p2; https://www.gtexportal.org/home/). We extracted the gene expression data of the LCLs, and the

723 genotype data of the eQTLs used in the GEUVADIS analysis. Since GTEx uses chromosome
724 locations to identify genetic variants, we extracted the coordinates of the GEUVADIS eQTLs in
725 Ensemble (GRCh 37; https://grch37.ensembl.org/index.html) using the rs IDs. Not all
726 GEUVADIS eQTLs can be found in the GTEx samples. Among eQTLs that can be found in the
727 GTEx samples, not all their associated genes have expression measurements. In the end, we
728 found 40 eQTL-gene sets with data available in both GEUVADIS and GTEx LCLs
729 (**Supplementary Table 4)**. For each of these sets, we ran MRPC with an FDR of 0.05, and
730 summarized the results in **Supplementary Table 4**. For those sets that were inferred to have an
731 M1 model by MRPC in GEUVADIS, we also ran function findr.pijs_gassist_pv() from the R
732 package findr, and function cit.cp() from the R package cit on each set to test whether there is a
733 causal model as in $V_1 \rightarrow T_1 \rightarrow T_2$ or $V_1 \rightarrow T_2 \rightarrow T_1$ (**Supplementary Table 5**).
734
735 The Genotype-Tissue Expression (GTEx) Project was supported by the Common Fund of the
736 Office of the Director of the National Institutes of Health, and by NCI, NHGRI, NHLBI, NIDA,
737 NIMH, and NINDS. The gene expression data used for the analyses described here were
738 obtained from the GTEx Portal (https://www.gtexportal.org/home/datasets;
739 GTEx_Analysis_2016-01-15_v7_RNASeQCv1.1.8_gene_tpm.gct.gz) on 10/24/2017, the
740 genotype data were available through dbGaP accession number phs000424.v7.
741
742 **Analysis of the TCGA breast cancer data.** We used the frequently altered genes identified
743 earlier[39] for this analysis. We downloaded the copy number variation, expression and
744 methylation data for these genes from cBioPortal (http://www.cbioportal.org/), which provides
745 processed and normalized data. We also downloaded the clinical data of the breast cancer

746 patients from the TCGA website (https://cancergenome.nih.gov/). Using the clinical data, we
747 selected 566 patients that were ER+, were identified as white, and also had genetic and
748 molecular data, for our analysis.
749
750 **Adjusted Structural Hamming Distance (aSHD).** The SHD, as is implemented in the R
751 package pcalg[26] and bnlearn[25], counts how many differences exist between two directed graphs.
752 This distance is 1 if an edge exists in one graph but missing in the other, or if the direction of an
753 edge is different in the two graphs. The larger this distance, the more different the two graphs are.
754 Similar to our approach to recall and precision (see "Recall and precision" in Methods), we
755 adjusted the SHD to reduce the penalty on the wrong direction of an edge to 0.5. For example,
756 between two graphs $V \rightarrow T_1 \leftarrow T_2$ and $V \rightarrow T_1 \rightarrow T_2$, the SHD is 1 and our aSHD is 0.5. By contrast,
757 between graphs $V \rightarrow T_1 \leftarrow T_2$ and $V \rightarrow T_1$, T_2 (no edge between T_1 and T_2, or between V and T_2),
758 both the SHD and aSHD are 1. Therefore, our adjustment penalizes the wrong direction less than
759 the wrong inference of the edge.
760
761 **Properties of MRPC.** A causal graph with a mixture of directed and undirected edges is
762 essentially an equivalent class of directed acyclic graphs (DAGs) that have the same likelihood.
763 However, the search problem of learning the DAG with the highest likelihood when the number
764 of parent nodes is greater than 1 has been proven to be NP-complete[51], the hardest computational
765 problem. Learning even just the equivalent classes of a DAG with the number of parent nodes
766 being greater than 1 is also NP-complete[52], as the space of equivalent classes of DAGs is super-
767 exponential[49] in the number of nodes. Therefore, the PC algorithm and similar algorithms get
768 around the computational issue with local searches. Although it is not known theoretically that

769 these PC algorithms achieve the global optimality defined by, for example, the likelihood, it has
770 been shown that the PC algorithm is consistent[49]: with a large sample size, the PC algorithm is
771 expected to recover the true graph. In particular, consistency of the PC algorithm is essentially
772 consistency of the step of graph skeleton inference, as this step contains all the statistical
773 inference[53]. Since MRPC uses essentially the same procedure for skeleton inference as the PC
774 algorithm, MRPC is also consistent.
775
776 **MRPC and other PMR-based methods.** Although our MRPC employs the PMR, it is
777 fundamentally different from other PMR-based methods. Most of the methods incorporating the
778 PMR fall into two classes. One class, including cit and findr, is called *mediation-based methods*
779 that require individual-level data, generally do not estimate the causal effect sizes, and can infer
780 networks of multiple phenotypes (e.g., a network of gene expression). The other class of
781 methods are called *MR methods*[22] that can be applied to individual-level data as well as summary
782 statistics, estimate the causal effect sizes, and generally focus on three-node graphs with one
783 node being the genetic variant, and the other two nodes being phenotypes of interest. Both
784 classes of methods employ the PMR and focus on the "causal model", in which exposure acts as
785 the mediator. Although our MRPC method is closer to the mediation-based methods according
786 to the characteristics described above, the notion of "mediation" is less relevant to our method;
787 only Model 1 considers the "causal model", and therefore one of the two genes acts as the
788 mediator (**Fig. 1a**). More importantly, with our method we consider the PMR as a way to define
789 plausible causal relationships and to put constraints on the space of possible graphs. As a result,
790 our method can recover a variety of causal relationships, instead of the few that other PMR-
791 based methods can identify (**Fig. 2b**).

Analysis of the DREAM5 Systems Genetic Challenge A data. This challenge (https://www.synapse.org/#!Synapse:syn2820440/wiki/) provided gene expression and genotype data simulated for 1000 genes in five different directed networks for each of three sample sizes (100, 300 and 999). We focused on networks labeled 'net1' of each sample size. These three networks each contain around 2000 directed edges between gene expression. We applied MRPC to these three data sets. We ran MRPC with FDR=0.3 and TRANSWESD[48] with default parameters. True positives (1 for an edge with the correct direction, and 0.5 for an edge not with the correct direction), false positives, recall and precision are summarized in **Supplementary Table 10**. We set the FDR to be 0.3 for MRPC because results from TRANSWESD had a high FDR of 64-84%. Setting a low FDR in MRPC also led to fewer edges to be recovered. For the data set with the largest sample size (999), we also ran mmhc, since mmhc has the second best performance in our simulation, and summarized the statistics in this table.

Code availability. MRPC is implemented in an R package (v1.0.0) at https://github.com/audreyqyfu/mrpc/releases.

49. Kalisch M, Bühlmann P. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *J Mach Learn Res* **8**, 613-636 (2007).

50. van Buuren S, Groothuis-Oudshoorn K. mice: Multivariate imputation by chained equations in R. *J. Stat. Softw.* **45**, 67 (2011).

51. Hoffgen K. Learning and robust learning of product distributions. *Technical Report 464, Fachbereich Informatik, Universitat Dortmund* (1993).

52. Chickering DM. Learning Bayesian networks is NP-complete. In*: Fisher D, Lenz HJ (eds) Learning from Data Lecture Notes in Statistics, vol 112 Springer, New York, NY*, 121-130 (1996).

53. Meek C. Causal inference and causal explanation with background knowledge. In P.Besnard and S.Hanks, editors. *Uncertainty in Artificial Intelligence* **11**, 403–410 (1995).

ACKNOWLEDGMENTS

We thank Jonathan Pritchard, Anand Bhaskar, Towfique Raj, Boxiang Liu, Deigo Calderon, and Xiaoyue Wang for helpful discussions, Boxiang Liu for processing the GTEx genotype data, and Diego Calderon for detailed and constructive comments on an earlier version of the paper. We also thank Lingfei Wang for help with the findr package, Huange Wang for providing the QPSO code, and Thomas Cokelaer for help with the DREAM5 Challenge data and scoring strategy. This research is supported by NIH R00HG007368 (to A.Q.F.) and partially by the NIH/NIGMS grant P20GM104420 to the Center for Modeling Complex Interactions at the University of Idaho.

AUTHOR CONTRIBUTIONS

A.Q.F. conceived project. M.B.B. and A.Q.F. developed the method. M.B.B. implemented the software. M.B.B. and A.Q.F. performed all the analyses and wrote the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

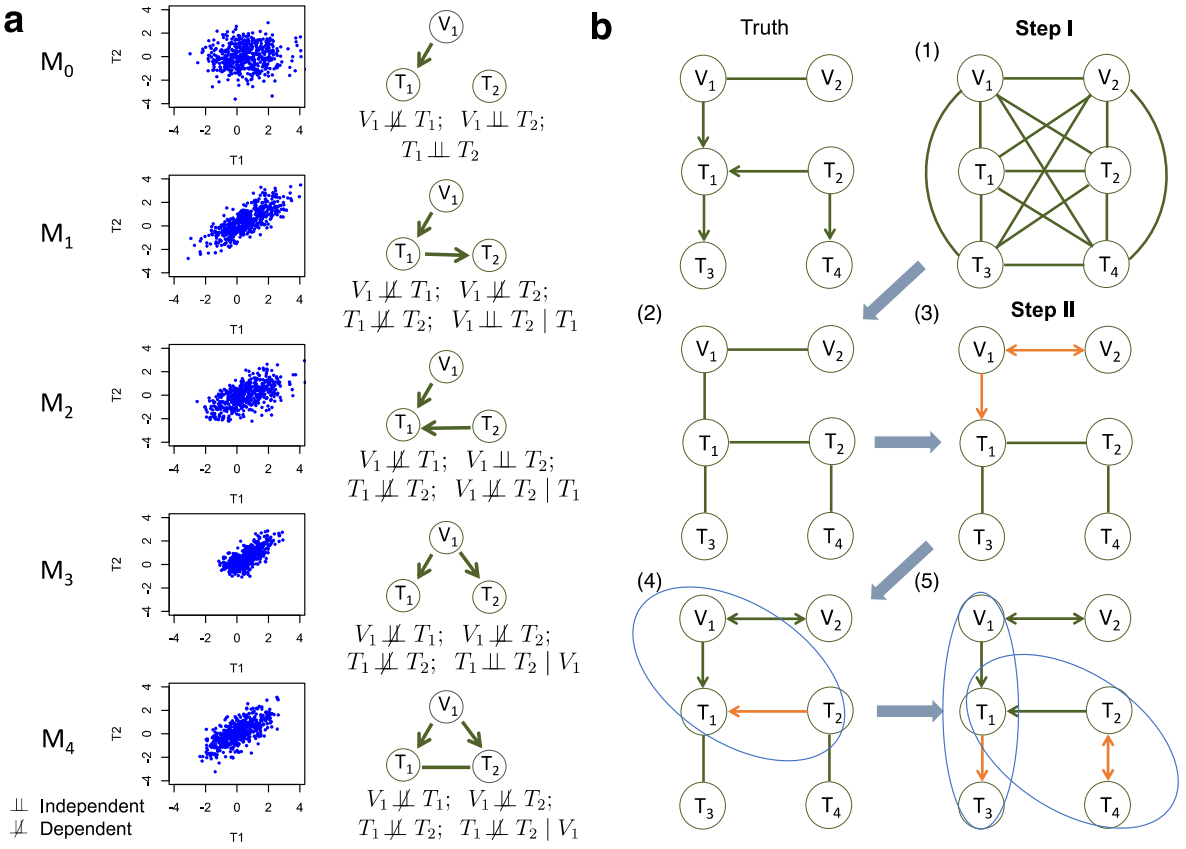

Figure 1: Five basic topologies under the principal of Mendelian randomization and the MRPC algorithm. **(a)** Each topology involves three nodes: a genetic variant (V_1), and two molecular phenotypes (T_1 and T_2). Directed edges indicate direction of causality, and undirected edges indicate that the direction is undetermined (or equivalently, both directions are equally likely). For each topology (or model), a scatterplot between the two phenotypes is generated using simulated data, the topology is shown, and the marginal and conditional dependence relationships are given. M_0 is the null model where T_1 and T_2 are marginally independent, and therefore the scatterplot does not show correlation. All the other models show scatterplots with similar levels of correlation. Our MRPC can distinguish the non-null models despite similar correlation. **(b)** The MRPC algorithm consists of two steps (see details in **Supplementary Figs 1**

854 **and 2**). In Step I, it starts with a fully connected graph shown in (1), and learns a graph skeleton
855 shown in (2), whose edges are present in the final graph but are undirected. In Step II, it orients
856 the edges in the skeleton in the following order: edges involving at least one genetic variants (3),
857 edges in a v-structure (if v-structures exist) (4), and remaining edges, for which MRPC
858 iteratively forms a triplet and checks which of the five basic models under the PMR is consistent
859 with the triplet (5). If none of the basic models matches the triplet, the edge is left unoriented
860 (shown as bidirected).

861

862

863

864

a

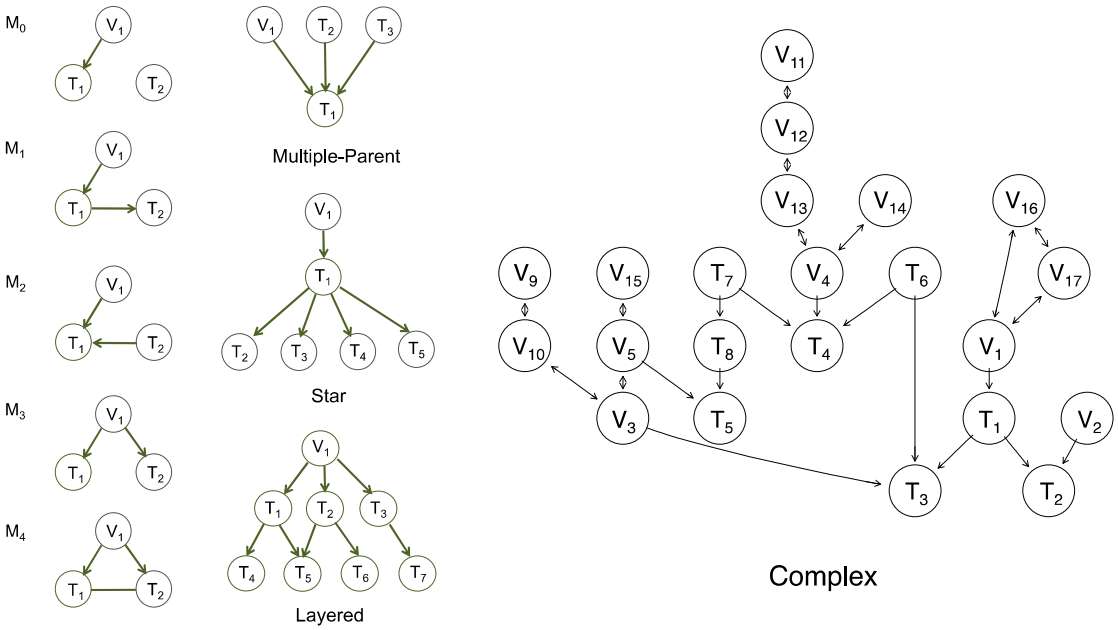

b

	MRPC	mmhc	pc	findr	cit	QPSO*
M_0	√	√	√	√	X	X
M_1	√	√	√	√	√	√
M_2	√	√	√	X	X	√
M_3	√	√	√	√	X	X
M_4	√	√	√	X	X	√
Multi-parent	√	√	√	X	X	√
Star	√	√	√	√	√	√
Layered	√	√	√	√	√	√
Complex	√	√	√	X	X	√

865

866

867 **Figure 2: Comparison of MRPC with other methods on simulated data**. **(a)** Topologies used
868 to generate synthetic data (see "Generating simulated data" in Methods). **(b)** Table summarizing
869 graphs to which each method under comparison is applicable. *Note that QPSO does not learn
870 the causal graph from scratch. Instead, it takes a graph skeleton as the input and seeks the
871 optimal orientation of the edges in this undirected network. Edges involving genetic variants
872 need to be already oriented in the skeleton. Therefore, QPSO does not identify M_0 or M_3.

873

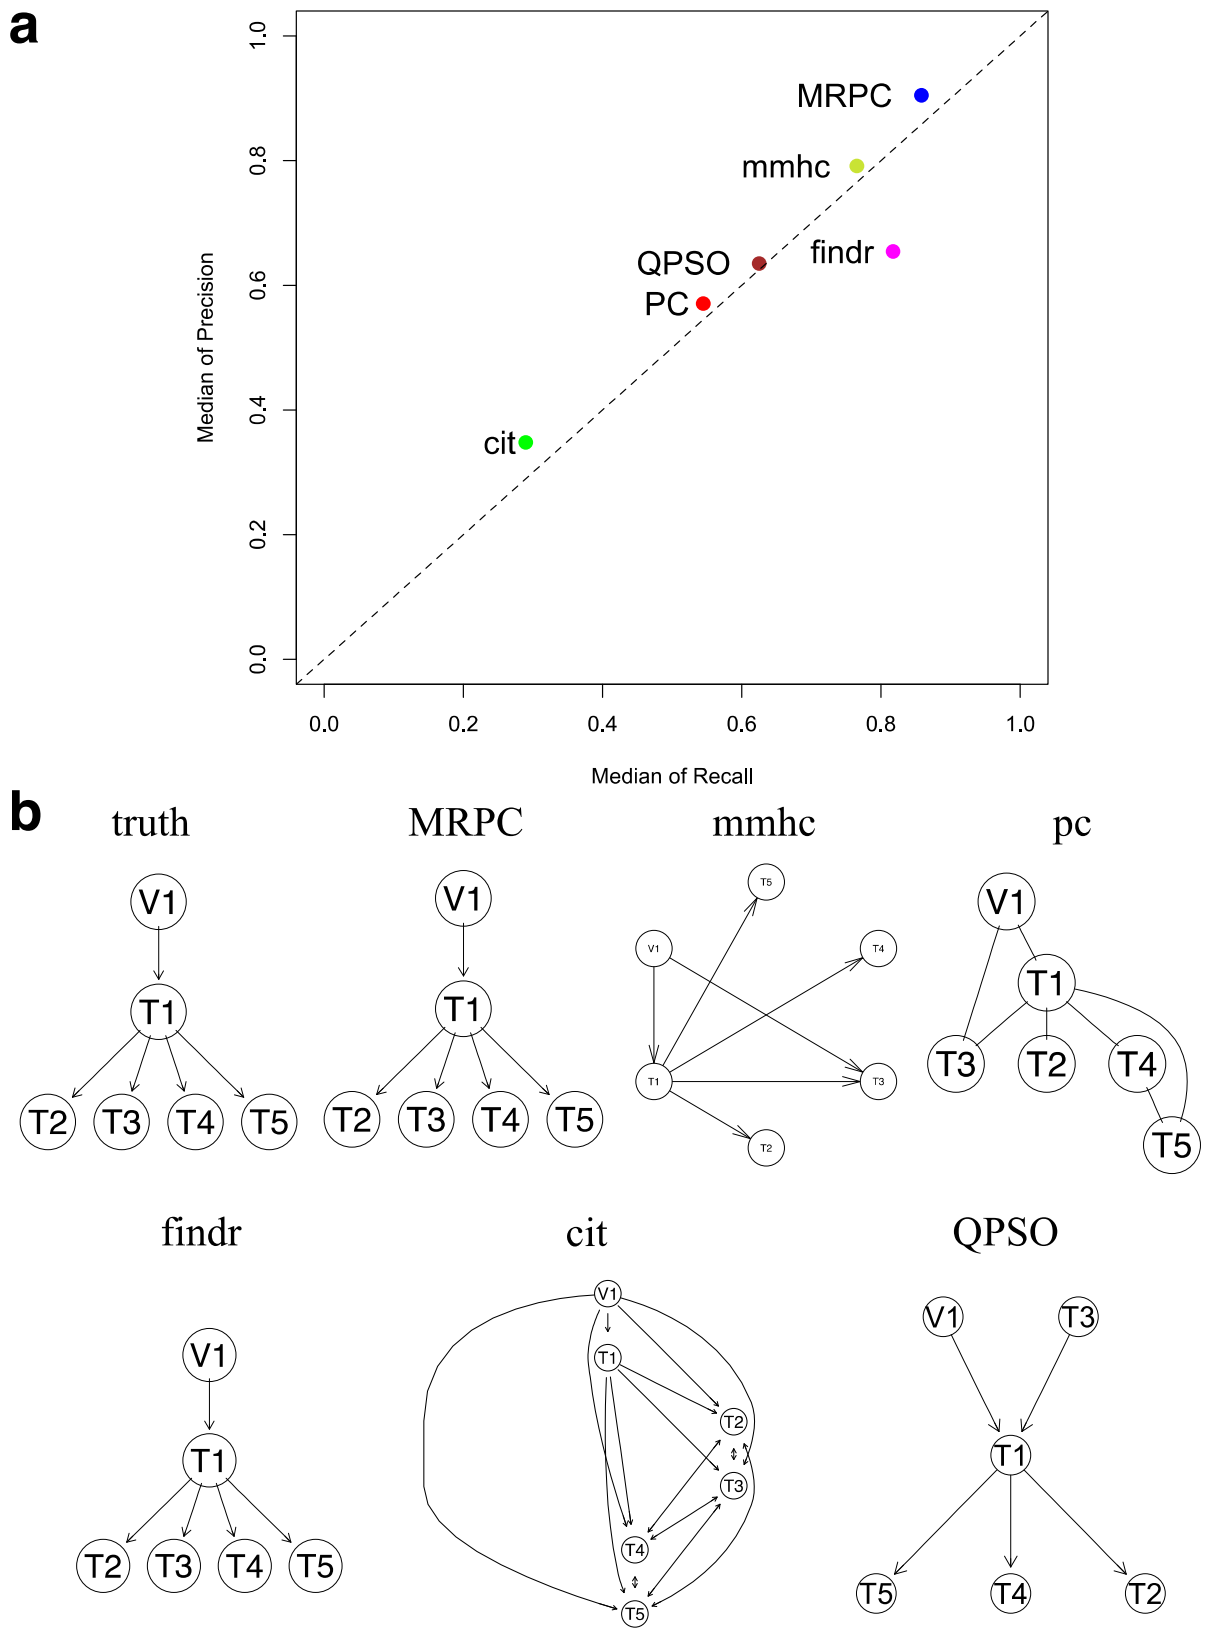

874

875 **Figure 3: Results of method comparison on simulated data**. **(a)** Median recall and precision
876 over all parameter settings. For each of the topologies in **Fig. 2a**, 1000 datasets were generated
877 for three different signal strengths (γ, which is the coefficient of parent nodes in the linear model;
878 see Methods for simulation details) and four different sample sizes (n). Each of the six methods
879 was applied where possible and the recall and precision were calculated for the inferred graph
880 relative to the truth. The median of all the mean recall (or precision) is used as a metric of the
881 overall performance of the method. Note that only 20 datasets were used for QPSO in each
882 parameter setting due to long runtime. **(b)** An example of inferred graphs from all six methods
883 on data simulated under a star model with a large sample size (n =1000) and strong signal
884 ($\gamma = 1.0$).

885

886

887

888

889

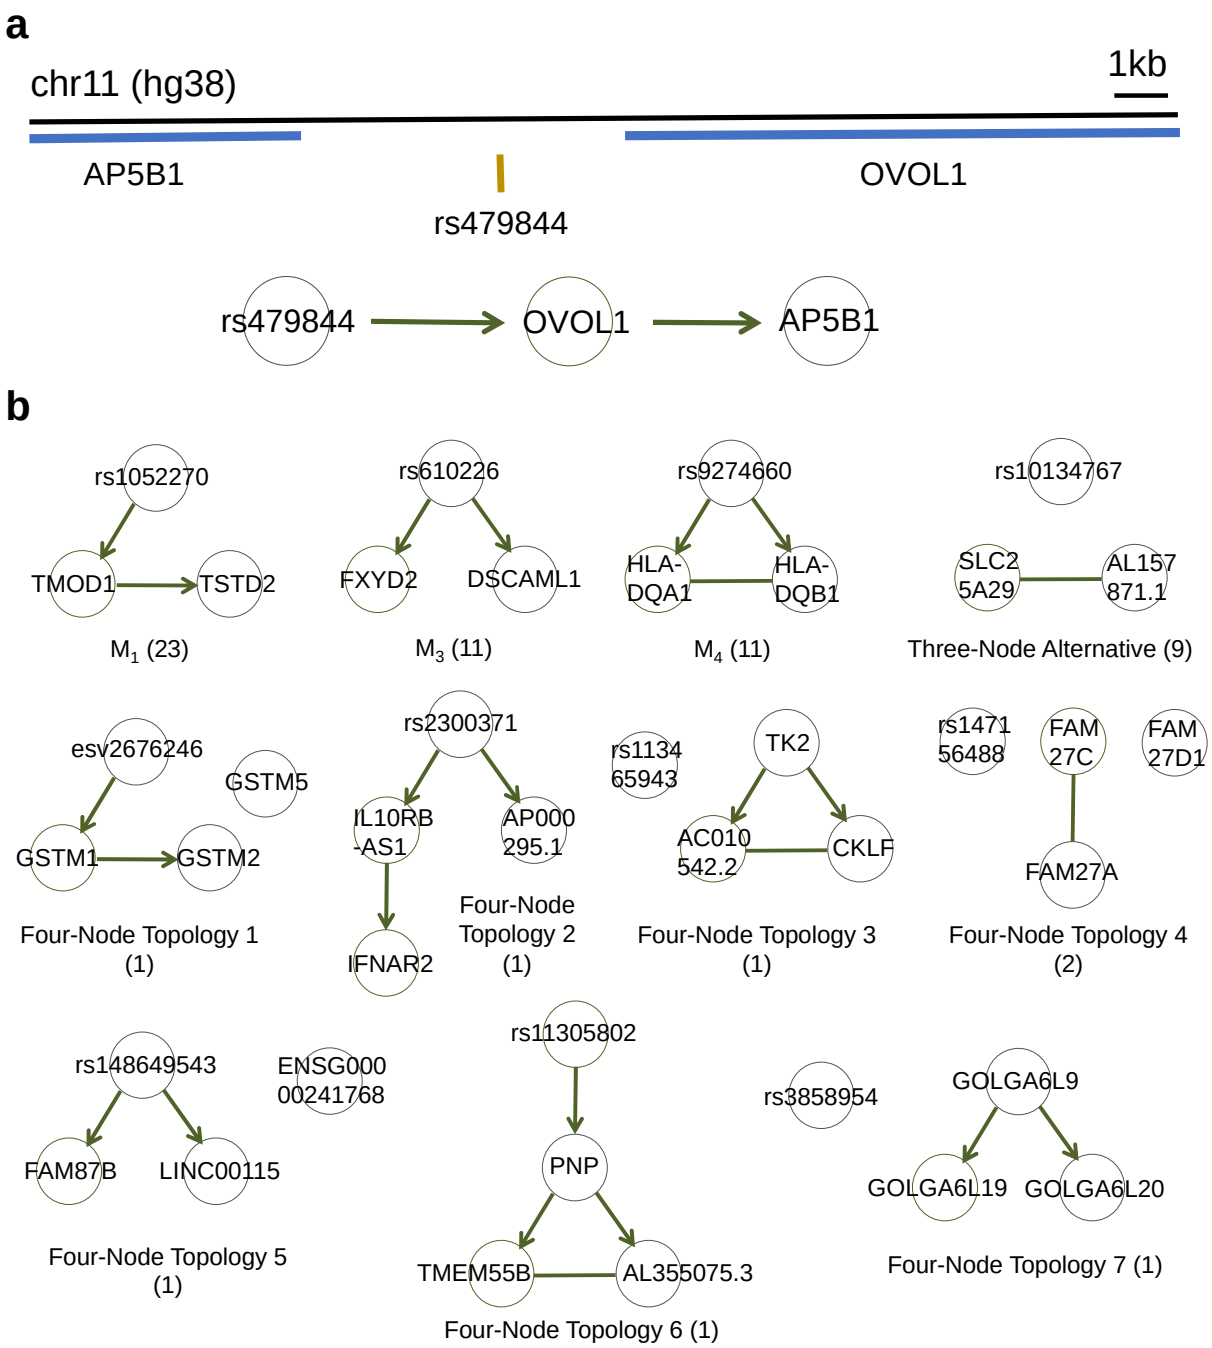

Figure 4: MRPC distinguishes direct and indirect target genes of eQTLs in the GEUVADIS data for the European cohort. **(a)** rs479844 is a GWAS significant SNP for atopic march in the GWAS Catalog, and an eQTL identified in GEUVADIS for two genes. (**b**) MRPC learns 11 distinct topologies among associated genes for eQTLs. The number in parentheses are the

895 number of eQTL-gene sets with the corresponding inferred topology.

896

897

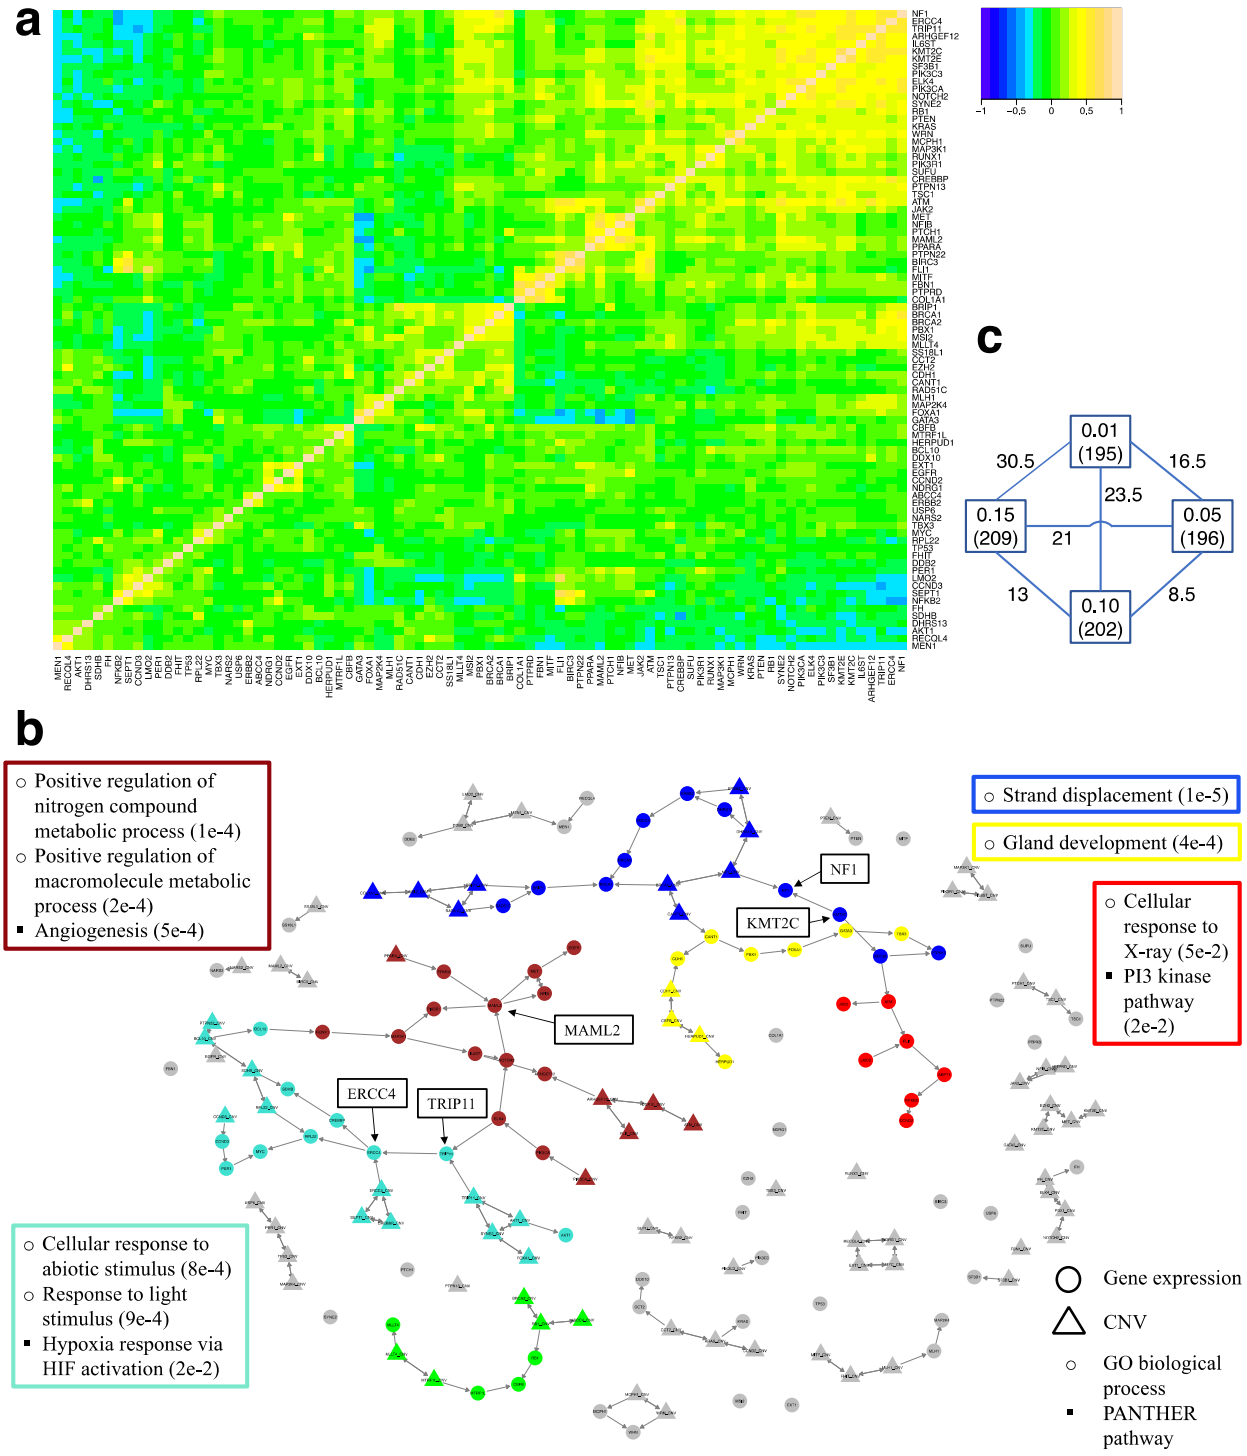

898

899 **Figure 5: MRPC learns a causal regulatory network for frequently altered cancer genes**

900 **using the TCGA breast cancer data**. **(a)** Pearson correlation heatmap for the 85 genes with

901 hierarchical clustering in rows and columns. **(b)** The causal network inferred at FDR of 5% by

MRPC. Modules were identified by WGCNA, such that each non-grey module contains at least seven nodes and four genes. Grey nodes were not assigned to any module. For each module, the box with the corresponding color contains the top GO biological processes and PANTHER pathways (if exist) enriched for the module, with p values in parentheses (complete results in **Supplementary Tables 8, 9**). **(c)** Distances between networks inferred by MRPC at different values of FDR. The square indicates the FDR with the total number of edges in parentheses, and the numbers on the lines are the adjusted Structural Hamming Distance (aSHD) between two graphs (see "Adjusted Structural Hamming Distance (aSHD)" in Methods). These numbers demonstrate the stability of the MRPC inference; see main text for detail.