Learning causal biological networks with generalized 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. MRPC is the first machine learning algorithm that incorporates generalized Mendelian randomization (gMR) in classical algorithms for learning directed acyclic graphs in computer science. The causal network contains directed edges, which indicate causal directions. We demonstrate through simulation that MRPC outperforms existing large-scale network inference methods and MR-based methods. 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. 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. 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 Model 1 and Model 2 in **Fig. 1a**). This is the rationale behind Mendelian Randomization (MR). MR assumes that the alleles of a genetic variant are randomly assigned to individuals in a population, analogous to a natural perturbation experiment[4]. MR 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[5-18]. However, MR typically examines 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[5,13-18]. This graph is sensible when T_2 is a potential outcome of T_1. However, when we examine relationships among gene expression or other molecular phenotypes, the relationship between T_1 and T_2 is unclear, and Model 1 alone does not have the flexibility of examining other possibilities. As a result, existing methods for constructing gene regulatory networks (or networks of molecular phenotypes in general) based on MR are limited in the causal relationships they can recover. In this paper, we generalize MR to account for a variety of causal relationships.

Additionally, implementation of (standard) MR in genomics has not been efficient: existing methods generally work with a small number of nodes, may require temporal information, or are not easily generalizable to large graphs. There has been some recent effort to address the efficiency issue[14]. Meanwhile, in machine learning, a class of algorithms, such as those based on the classic PC algorithm[19-23], 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[22,23]), or to reduce the number of statistical tests needed for inferring the skeleton (e.g., the R package bnlearn[20,21]).

Here we develop a new method, namely MRPC, which incorporates generalized MR (gMR) 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 gMR with machine learning, our method is efficient and accurate. gMR can be thought of as a way of introducing useful constraints in graph learning and effectively reducing the search space of possible topologies. We generated simulation data from diverse topologies to assess the performance of MRPC and compare it with that of other large-scale network inference algorithms and MR-based methods. We use MRPC to distinguish direct and indirect targets among multiple genes associated with an eQTLs in the GEUVADIS data[24], and replicate the results in the GTEx data[25]. We use MRPC to construct a causal network among frequently altered cancer genes using the TCGA breast cancer patient data[26]. This network can be broken down into smaller modules of genes that are enriched for distinct biological functions. Therefore, MRPC efficiently uses individual-level genomic data to learn plausible biological mechanisms from combining genotype and gene expression data. Broadly, MRPC is applicable to other molecular phenotypes, organismal traits or phenotypes, in addition to gene expression.

Results

Generalized Mendelian Randomization (gMR). We generalized standard MR to allow for more possible causal relationships in a triplet of a genetic variant and two phenotypes (**Fig. 1a**). Under the assumption that genotype influences phenotype and not the other way around, the five models depicted in **Fig 1** 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 Mendelian randomization with T_1 being the exposure (e.g., gene expression) and T_2 being the outcome (e.g., clinical phenotype). In this setting, it is reasonable to expect the causal direction going from T_1 to T_2, making this model appropriate in epidemiological settings. cit[13,27] and findr[14], two existing MR-based methods, both focus on testing Model 1.

However, when T_1 and T_2 are of the same type of phenotypes (e.g., both being gene expression), Model 1 does not capture other possible causal relationships between T_1 and T_2. For example, 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 detected for T_2, or the best available variant does not provide association strong enough for network construction, likely due to noise in the data, a small sample size, etc. 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. Model 2 and Model 3 are also called the reactive model, and the independent model under standard MR, respectively[13]. Standard MR typically rejects these two models in search of the 'causal' model (Model 1). However, under our generalized MR, 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[28]): 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 graph inference method for genomic data (**Fig. 1b**; **Supplementary Fig. 1)**. 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 gMR in the second step of edge orientation (**Fig. 1b**), 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's correlation in all the marginal test and for the partial correlation in all the conditional test, consistent with the default test in pcalg[22]. 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 controlling both the FWER and FDR requires 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 edges being retained in the inferred graph, especially when the graph is large (see our simulation results below). We implemented in MRPC the LOND method for controlling the FDR in an online matter[29] (see Methods). The LOND method estimates the expected FDR conditioned on the number of tests performed so far and the number of rejections for these tests.

Furthermore, genomic data may contain outliers[30], which can greatly distort the inferred graph (see our simulation results below). Like pcalg, MRPC uses the correlation matrix, rather than the individual-feature matrix, as the input. We implemented in MRPC a method for calculating the robust correlation matrix[30] (see Methods) to alleviate the impact of outliers.

MRPC outperforms existing network inference algorithms on synthetic data in overall accuracy. We compared MRPC with the pc method (implemented in pcalg) and the mmhc method (implemented in bnlearn), and assessed their performance on synthetic data. We simulated data using linear models for the five basic topologies, three common topologies in biology (such as multi-parent, star, and layered graphs), as well as a complex topology with over 20 nodes (**Fig. 2a**). We varied the sample size, as well as the signal strength through the coefficients in the linear models (see Methods), learned the graph from the synthetic data, and calculated the distance between the inferred graph and the truth.

For each topology, we generated 1000 data sets with different combinations of signal strength and sample size, and compared the mean adjusted Structural Hamming Distance (aSHDs; see Methods) with standard deviation for each of the three methods (**Fig. 2b**; **Supplementary Table 1)**. Smaller aSHD corresponds to higher accuracy. MRPC performs better than pc and mmhc in most cases, and recovers the true graph particularly well at moderate or stronger signal with a medium or larger sample size. Take the star model for example. Even with a sample size of 1000, at different signal strengths, mmhc may infer an extra edge or wrongly orients an edge in a correct skeleton (**Fig. 2d**; **Supplementary Fig. 2**), whereas pc is unable to determine edge directions or wrongly identifies v-structures when the true model contains none (**Fig. 2d**; **Supplementary Figs. 3 and 2**). For the complex topology, MRPC performs consistently better than pc and mmhc. In the presence of outliers, MRPC also outperforms pc and mmhc (**Supplementary Fig. 4)**. In terms of runtime, MRPC is essentially the same as pc, both of which a bit slower than mmhc. Incidentally, mmhc does not handle missing values, which are common in genomic data.

Existing MR-based methods cannot deal with complex causal relationships. We next compare MRPC to two other MR-based methods, namely cit and findr; other MR-based methods cannot handle complex causal relationships, either[15]. 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. 2a**). These methods cannot handle scenarios such as M2 where no genetic variants are identified for T_2, or the association is too weak for the genetic variant to provide additional information.

Nor can they deal with scenarios like M4 where there is additional dependence between the genes, and where inference of the causal direction is not possible with the data. For comparison here, we limit 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 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. Since the total number of tests is not high, we applied Bonferroni correction with a familywise type I error rate of 0.05. For simple models (M0, M1 and M3), findr performs slightly better than MRPC when the signal is weak and the sample size is not large (**Fig. 2c; Supplementary Table 2**). Both outperform cit under the M1 model. With more complex models (star and layered graphs), however, MRPC outperforms the other two methods consistently. Take again the star model with a sample size of 1000 for example, where we varied the signal strengths in simulation. Although the conservative Bonferroni correction is applied to multiple testing, findr sometimes infers more edges than there are (**Supplementary Fig. 3**), whereas cit may infer a very dense graph or no edges at all (**Fig. 2d**; **Supplementary Figs. 3 and 2**), contributing to the large deviation from the truth (**Fig. 2c**).

Next, we examined 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. 5**). In other words, potential regulators and targets are known to findr and cit. For MRPC we continued to assume that the skeleton was unknown. However, even with such an advantage in inference, findr still did not outperform MRPC on simple models, and consistently produced a biased graph even with strong signals and large sample sizes. cit's performance is greatly improved compared to when the skeleton is unknown, although it still did not perform well on simple or complex models at weak signals or lower sample sizes.

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 the dependence structure among the genes. Indeed, among the eQTLs identified from the GEUVADIS data[24] (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 Methods). We applied MRPC to each of these eQTLs and their associated genes in the 373 Europeans, and identified 11 types of topologies (**Fig. 3; Supplementary Table 3;** also see comparison with mmhc and pc for some of the eQTL-gene sets in **Supplementary Fig. 6)**. 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 3)**.

Although the multiple associated genes of the same eQTL are physically near one another, our method managed to tease apart the different dependence (or regulatory relationships) among these genes. For example, the SNP rs479844 (chr11:65,784,486), one of the 62 eQTLs, turns out to be significant in at least three genome-wide association studies (GWASs) for atopic march and more specifically, atopic dermatitis (p values ranging from 10^{-10} to 10^{-18})[31-34]. 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. 3a**), 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 mis-specify edges or their directions (**Supplementary Fig. 6**). 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 Methods).

Since the GTEx consortium[25] 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 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 4**). This result is consistent with results in simulation studies (**Fig. 2c, d**): 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 cohort[26], with the aim of learning the causal gene regulatory network for frequently altered genes in breast cancer. Although frequently altered genes are involved in many pathways known to play an important role in cancer (e.g., DNA damage repair pathways[35], kinase signaling pathways[36], and so on)[37], how these genes regulate one another transcriptionally remains unclear.

Copy number variation usually has a much stronger effect on gene expression than SNPs do in breast cancer. We therefore used the copy number variation as the genotype, and gene expression as the molecular phenotype. Similar to an earlier investigation[37], we extracted 85 frequently altered genes (e.g., BRCA1, BRCA2, TP53, etc.) in breast cancer and their copy number variation data. We calculated the Pearson correlation matrix (**Fig. 4a**), and applied MRPC at the FDR of 5% (**Fig. 4b**), and subsequently at 0.01, 0.10 and 0.15. The inferred graphs appeared reasonably stable: each graph contains around 200 edges; when the FDR changes by 0.05, the number of edges inferred differently tends to be around 10, which is roughly 1/20 of all the edges

(**Fig. 4c**; **Supplementary Figs. 7-9**); this is consistent with the change of 0.05 in FDR, as this rate implies that on average roughly 5% of all the edges are likely to be false positives and therefore would not be consistently inferred at another FDR. In other words, most of the edges are inferred reliably across different FDR levels.

We searched through literature for evidence supporting the inferred directed edges here (**Supplementary Table 5**). 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 5**). 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 to prioritize which gene pairs to perform knockdown experiments on.

In the graph inferred at the FDR of 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. 4b**). For better visualization of the result, we then applied WGCNA[38] to the inferred graphs. 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. 4b**; **Supplementary Fig. 10**; **Supplementary Table 6**). 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[37] (**Fig. 4b**). We ran gene ontology (GO) enrichment analysis[39] 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[40], suggesting that the causal network we learned has a structure consistent with the underlying biological functions (**Fig. 4b**; **Supplementary Tables 7, 8**).

Additionally, our results also illustrate that 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. 4a**). However, no edge connects NF1 and ERCC4 at any of the FDRs we examined (**Fig. 4b**; **Supplementary Figs. 7-9**). A closer look at the conditional independence tests showed that the relatively strong correlation between the two genes is explained away by TRIP11, BRCA1 and the copy number variation of NF1 (p value: 3e-6; the significance level after accounting for the FDR: 5e-7). In other words, the correlation between NF1 and ERCC4 is indirect: it is induced by their 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[41], also shown in our inferred network (**Fig. 4b**), whereas the DNA repair gene ERCC4 is recently shown to be involved in the translesion DNA synthesis together with TRIP11 and other genes[42], consistent with the edge between these two genes in our inferred network (**Fig. 4b**).

Discussion

In summary, our MRPC method extends standard MR 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) and the molecular phenotype measurements (gene expression, or other functional data, such as exon expression, RNA editing, DNA methylation, etc.), and can be applied to a variety of causal inference problems. By incorporating gMR, MRPC makes it possible to examine how genotypes induce changes in phenotypes. Our method is nonparametric in that no explicit distributions are assumed for the underlying graph. Our method resolves several issues present in popular implementations of causal graph inference algorithms, which have led to large biases in inferred graphs. Like most causal graph learning methods, a key assumption behind MRPC is that there are no hidden nodes that are connected to the observed nodes in the graph. Whereas this assumption may not hold in biology, we can take additional measures to alleviate the impact of hidden nodes. For example, genes are often grouped in clusters that tend to have higher correlation within the cluster. Our method can be applied to genes within a gene cluster to build the detailed causal network. As the next step, we are working on extensions of MRPC to deal with hidden variables[8,22]. The current version of MRPC has already demonstrated its power in tackling several biological problems on causality and in effectively using existing large amounts of genomic data.

1. Segal, E. *et al*. Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data. *Nat Genet* **34**, 166-76 (2003).

2. Housden, B.E. *et al*. Transcriptional dynamics elicited by a short pulse of notch activation involves feed-forward regulation by E(spl)/Hes genes. *PLoS Genet* **9**, e1003162 (2013).

3. Cheung, V.G. & Spielman, R.S. Genetics of human gene expression: mapping DNA variants that influence gene expression. *Nat Rev Genet* **10**, 595-604 (2009).

4. Davey Smith, G. & Hemani, G. Mendelian randomization: genetic anchors for causal inference in epidemiological studies. *Hum Mol Genet* **23**, R89-98 (2014).

5. Zhu, Z. *et al*. Integration of summary data from GWAS and eQTL studies predicts complex trait gene targets. *Nat Genet* **48**, 481-7 (2016).

6. Hill, S.M. *et al.* Inferring causal molecular networks: empirical assessment through a community-based effort. *Nat Methods* **13**, 310-8 (2016).

7. Zhang, B. *et al.* Integrated systems approach identifies genetic nodes and networks in late-onset Alzheimer's disease. *Cell* **153**, 707-20 (2013).

8. Chaibub Neto, E. *et al.* Modeling causality for pairs of phenotypes in system genetics. *Genetics* **193**, 1003-13 (2013).

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

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

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

12. Franzen, O. *et al.* Cardiometabolic risk loci share downstream cis- and trans-gene regulation across tissues and diseases. *Science* **353**, 827-30 (2016).

13. Millstein, J., Zhang, B., Zhu, J. & Schadt, E.E. Disentangling molecular relationships with a causal inference test. *BMC Genet* **10**, 23 (2009).

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

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

16. Huang, Y.T., Liang, L., Moffatt, M.F., Cookson, W.O. & Lin, X. iGWAS: integrative genome-wide association studies of genetic and genomic data for disease susceptibility using mediation analysis. *Genet Epidemiol* **39**, 347-56 (2015).

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

18. Connor, L. & Price, A.L. Distinguishing genetic correlation from causation across 52 diseases and complex traits. *bioRxiv* 205435 (2017).

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

20. Tsamardinos, I., Brown, L.E. & Aliferis, C.F. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning* **65**, 31-78 (2006).

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

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

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

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

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

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

27. Millstein, J., Chen, G.K. & Breton, C.V. cit: hypothesis testing software for mediation analysis in genomic applications. *Bioinformatics* **32**, 2364-5 (2016).

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

29. Javanmard, A. & Montanari, A. On online control of false discovery rate. *arXiv* 1502.06197v2 (2015).

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

31. Marenholz, I. *et al.* Meta-analysis identifies seven susceptibility loci involved in the atopic march. *Nat Commun* **6**, 8804 (2015).

32. 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**, 1449-56 (2015).

33. Paternoster, L. *et al.* Meta-analysis of genome-wide association studies identifies three new risk loci for atopic dermatitis. *Nat Genet* **44**, 187-92 (2011).

34. MacArthur, J. *et al.* The new NHGRI-EBI catalog of published genome-wide association studies (GWAS Catalog). *Nucleic Acids Res* **45**, D896-901 (2017).

35. Roos, W.P., Thomas, A.D. & Kaina, B. DNA damage and the balance between survival and death in cancer biology. *Nat Rev Cancer* **16**, 20-33 (2016).

36. Martini, M., De Santis, M.C., Braccini, L., Gulluni, F. & Hirsch, E. PI3K/AKT signaling pathway and cancer: an updated review. *Ann Med* **46**, 372-83 (2014).

37. Wang, X., Fu, A.Q., McNerney, M.E. & White, K.P. Widespread genetic epistasis among cancer genes. *Nat Commun* **5**, 4828 (2014).

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

39. Gene Ontology Consortium: going forward. *Nucleic Acids Res* **43**, D1049-56 (2015).

40. Mi, H., Muruganujan, A. & Thomas, P.D. PANTHER in 2013: modeling the evolution of gene function, and other gene attributes, in the context of phylogenetic trees. *Nucleic Acids Res* **41**, D377-86 (2013).

41. Rao, R.C. & Dou, Y. Hijacked in cancer: the MLL/KMT2 family of methyltransferases. *Nat Rev Cancer* **15**, 334-46 (2015).

42. Ziv, O. *et al.* Identification of novel DNA-damage tolerance genes reveals regulation of translesion DNA synthesis by nucleophosmin. *Nat Commun* **5**, 5437 (2014).

METHODS

Calculation of robust correlation. We implemented the method in Badsha et al.[30] to calculate the robust correlation matrix as the input to the MRPC inference. Specifically, for data that are approximately normal (usually after preprocessing of the data), we calculated iteratively the robust mean vector $\boldsymbol{\mu}$ and the robust covariance matrix $\boldsymbol{v}$ until convergence. At the t+1st iteration,

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

and

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

where,

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

In the equations above, $\boldsymbol{x}_i$ is the vector of gene expression in the ith sample, n the sample size, and β the tuning parameter. Equation (3) downweighs the outliers through β, which takes values in [0,1]. Larger β leads to smaller weights on the outliers. When $\beta = 0$, equation (2) is similar to the standard definition of the variance, except that the scalar is 1/n, whereas the unbiased estimator of the variance has a scalar of 1/(n-1). In our applications, we generally use $\beta = 0.005$ if the robust correlation matrix is calculated, consistent with Badsha et al.[30]. When the data matrix contains missing values, we perform imputation using the R package mice[43]. Alternative, one may impute the data using other appropriate methods, and calculate the correlation matrix as the input for MRPC.

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 *i*th test is

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

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 *i*th 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 are listed in an example given in **Supplementary Table 9.**

Simulation. 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 targets); multi-parent (i.e., one molecular phenotype has multiple regulators apart from the genetic variants); and layered; and iii) a complex topology.

In each topology, we simulated the data first for the nodes without parents, and then for other nodes. Genetic variants are nodes without parents, and we assume them to be biallelic SNPs with three genotypes 0, 1, and 2. Denote the minor allele frequency by q and assume Hardy-Weinberg equilibrium. Then the genotype of the *i*th variant, V_i , follows a multinomial distribution:

$$\Pr(V_i = 0) = (1-q)^2, \Pr(V_i = 1) = 2q(1-q), \Pr(V_i = 2) = q^2.$$

Denote the *j*th molecular phenotype by T_j and the set of its parent nodes by P, which may be empty, or may include variant nodes or nodes of other molecular phenotypes. We assume that the molecular phenotype T_j follows a normal distribution

$$T_j \sim N(\gamma_0 + \sum_{k \in \mathrm{P}} \gamma_k V_k + \sum_{l \in \mathrm{P}} \gamma_l T_l \,,\, \sigma^2).$$

The variance may be different for different nodes. For simplicity, we use the same value for all the nodes.

We treat undirected edges as bidirected edges and interpret such an edge as an average of the two directions with equal weights. For example, for the undirected edge in Model 4 in **Fig. 1a**, we generate data for $T_1 \rightarrow T_2$:

$$T_1 \sim N\left(\gamma_0 + \gamma_1 V,\ \sigma^2\right);\ T_2 \sim N\left(\gamma_0 + \gamma_1 V + \gamma_2 T_1,\ \sigma^2\right),$$

and separately for T1←T2:

$$T_1 \sim N\left(\gamma_0 + \gamma_1 V + \gamma_2 T_2,\ \sigma^2\right);\ \ T_2 \sim N\left(\gamma_0 + \gamma_1 V,\ \sigma^2\right).$$

We then randomly choose a pair of values with 50:50 probability for each sample.

For simplicity in simulation, we set $\gamma_0 = 0$ and all the other γs to take the same value, which reflects the strength of the association signal. We considered three values for the slopes: 0.2 (weak signal), 0.5 (moderate signal), and 1.0 (strong signal). We also varied the sample size: 200 (small), 500 (medium), and 1000 (large). Thus, we considered nine combinations of signal strength and sample size (**Fig. 2b, c**).

Under each combination, we generated 1000 data sets for each topology. For each data set, we shuffled the columns corresponding to gene expression; if an inference method is sensitive to the ordering of the columns, the inferred graph would have a large variance across data sets. We then applied each method to a data set with permuted columns and calculated the distance (described below) between the inferred graph and the truth. To summarize the results, we computed the mean and standard deviation of the distance across 1000 data sets for each method, and displayed the mean as the bar and the standard deviation as the error bar in the horizontal bar plots (**Fig. 2b, c**).

Calculation of adjusted Structural Hamming Distance (aSHD). The SHD, as is implemented in the R package pcalg[22] and bnlearn[21], counts how many differences exist between two directed graphs. This distance is 1 if an edge exists in one graph but missing in the other, or if the direction of an edge is different in the two graphs. The larger this distance, the more different the

two graphs are. We adjusted the SHD to reduce the penalty on the wrong direction of an edge to 0.5. For example, 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, 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), both the SHD and aSHD are 1. Therefore, our adjustment penalizes the wrong direction less than the wrong inference of the edge.

In our comparison with other PC algorithms, since these PC algorithms are not designed to infer the relationships among genetic variants correctly, we ignored the edges among genetic variants in the aSHD calculation for mmhc and pc.

Comparison of MRPC, findr and cit. For MRPC, we followed the same procedure as that in the comparison with other PC algorithms: we simulated 1000 data sets for each combination of the parameters (signal strength and sample size), permuted the columns for each data set, applied MRPC algorithm, and calculated the mean aSHD across simulated data sets.

As explained in the main text, findr can infer M0, M1, M3, star and layered graphs, whereas cit M1, star and layered graphs, as star and layered graphs involve only M1 relationships. We therefore generated 1000 data sets from each of these topologies with different parameter combinations. We next applied findr and cit to the simulated data in order to infer the causal graph. The causal graph is represented by an asymmetric adjacency matrix, where 1 at position (i, j) indicates a directed edge from node i to node j, and 0 no edge.

Unlike PC algorithms that learn the graph skeleton first and orient the edges next, findr and cit test for directed edges in a single step for a triplet of nodes (the genetic variant and two gene expression nodes). This means that in order to learn the topology, we needed to examine all possible gene pairs (e.g., T_1 and T_2; and T_2 and T_1) and then apply findr or cit to the triplet of each of the gene pairs and the genetic variant. Based on the hypothesis testing result from findr or cit, if there was evidence for a directed edge between two nodes, we added 1 to the current value in the adjacency matrix for those two nodes. Otherwise we left the value unchanged. After examining all gene pairs, we converted all positive values in the adjacency matrix to 1 to represent a directed edge. This way, no edges inferred would be eliminated in later tests. We then calculated the aSHD between the inferred adjacency matrix and that of the true graph, and averaged the aSHDs across simulated data sets.

Although findr aims to compute a causality probability for a triplet, its current implementation for this calculation cannot be applied to small graphs, or cases where multiple genes share the same eQTL and where some of the genes do not have eQTLs. We therefore used the function findr.pijs_gassist_pv() from the R package findr to conduct five hypothesis tests (the p values from these five tests are then converted to a causality probability) for each ordered gene pair with the genetic variant. Consider a triplet V_1, T_1 and T_2. The null (H_0) and alternative (H_a) hypotheses of these five tests are:

Test #1: H_0: V_1 and T_1 independent; H_a: $V_1 \rightarrow T_1$;

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

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

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

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

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

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

Analysis of the GEUVADIS data. The GEUVADIS project (http://www.ebi.ac.uk/Tools/geuvadis-das/) performed RNA-seq (gene expression) on 373 Europeans and 89 Africans from the 1000 Genomes Project. The GEUVADIS project combined the gene expression data with the genotype data, and identified eQTLs across the human genome. Among the most stringent set of eQTLs, 62 have more than one target gene. We extracted the genotypes of these eQTLs and the expression of the target genes in the 373 Europeans, and 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 11**) 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 12**) 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 13**) 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 genotype data of the eQTLs used in the GEUVADIS analysis. Since GTEx uses chromosome locations to identify genetic variants, we extracted the coordinates of the GEUVADIS eQTLs in Ensemble (GRCh 37; https://grch37.ensembl.org/index.html) using the rs IDs. Not all

GEUVADIS eQTLs can be found in the GTEx samples. Among eQTLs that can be found in the GTEx samples, not all their associated genes have expression measurements. In the end, we found 40 eQTL-gene sets with data available in both GEUVADIS and GTEx LCLs (**Supplementary Table 3)**. For each of these sets, we ran MRPC with an FDR of 0.05, and summarized the results in **Supplementary Table 3**. For those sets that were inferred to have an M1 model by MRPC in GEUVADIS, we also ran function findr.pijs_gassist_pv() from the R package findr, and function cit.cp() from the R package cit on each set to test whether there is a causal model as in $V_1 \rightarrow T_1 \rightarrow T_2$ or $V_1 \rightarrow T_2 \rightarrow T_1$ (**Supplementary Table 4**).

The Genotype-Tissue Expression (GTEx) Project was supported by the Common Fund of the Office of the Director of the National Institutes of Health, and by NCI, NHGRI, NHLBI, NIDA, NIMH, and NINDS. The gene expression data used for the analyses described here were obtained from the GTEx Portal (https://www.gtexportal.org/home/datasets; GTEx_Analysis_2016-01-15_v7_RNASeQCv1.1.8_gene_tpm.gct.gz) on 10/24/2017, the genotype data were available through dbGaP accession number phs000424.v7.

Analysis of the TCGA breast cancer data. We used the frequently altered genes identified earlier[37] for this analysis. We downloaded the copy number variation, expression and methylation data for these genes from cBioPortal (http://www.cbioportal.org/), which provides processed and normalized data. We also downloaded the clinical data of the breast cancer patients from the TCGA website (https://cancergenome.nih.gov/). Using the clinical data, we selected 566 patients that were ER+, were identified as white, and also had genetic and molecular data, for our analysis.

Software. MRPC is implemented in an R package at https://github.com/audreyqyfu/mrpc/releases.

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

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. 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.

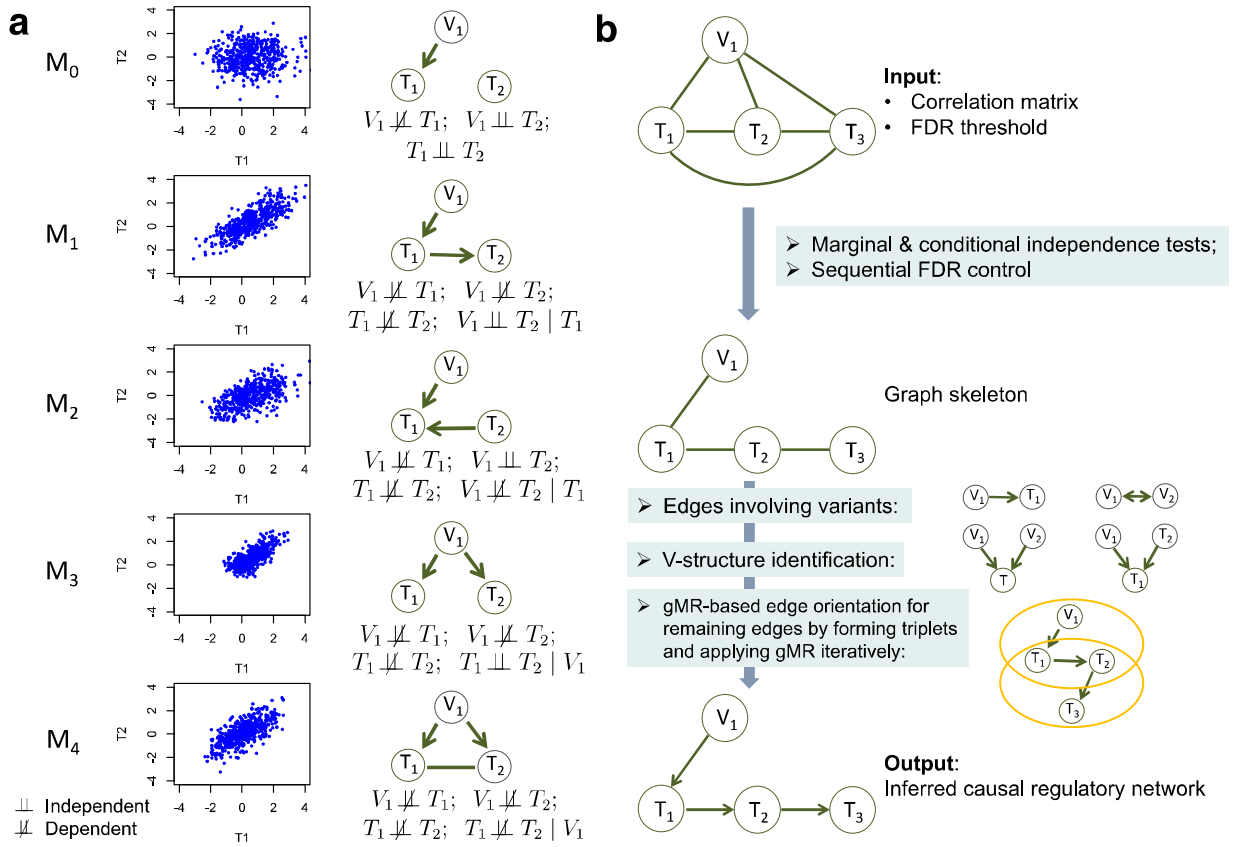

Figure 1: Five basic topologies under generalized Mendelian randomization and the MRPC algorithm. **(a)** Each topology involves three nodes: a genetic variant (V_1), and two molecular phenotypes (T_1 and T_2). Directed edges indicate direction of causality, and undirected edges indicate that the direction 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 takes as input the correlation matrix and FDR threshold. It starts with a fully connected graph, first learns a graph skeleton, whose edges are present in the final graph but are undirected, and then orients these edges. MRPC implements sequential FDR control for skeleton inference and generalized MR for edge orientation.

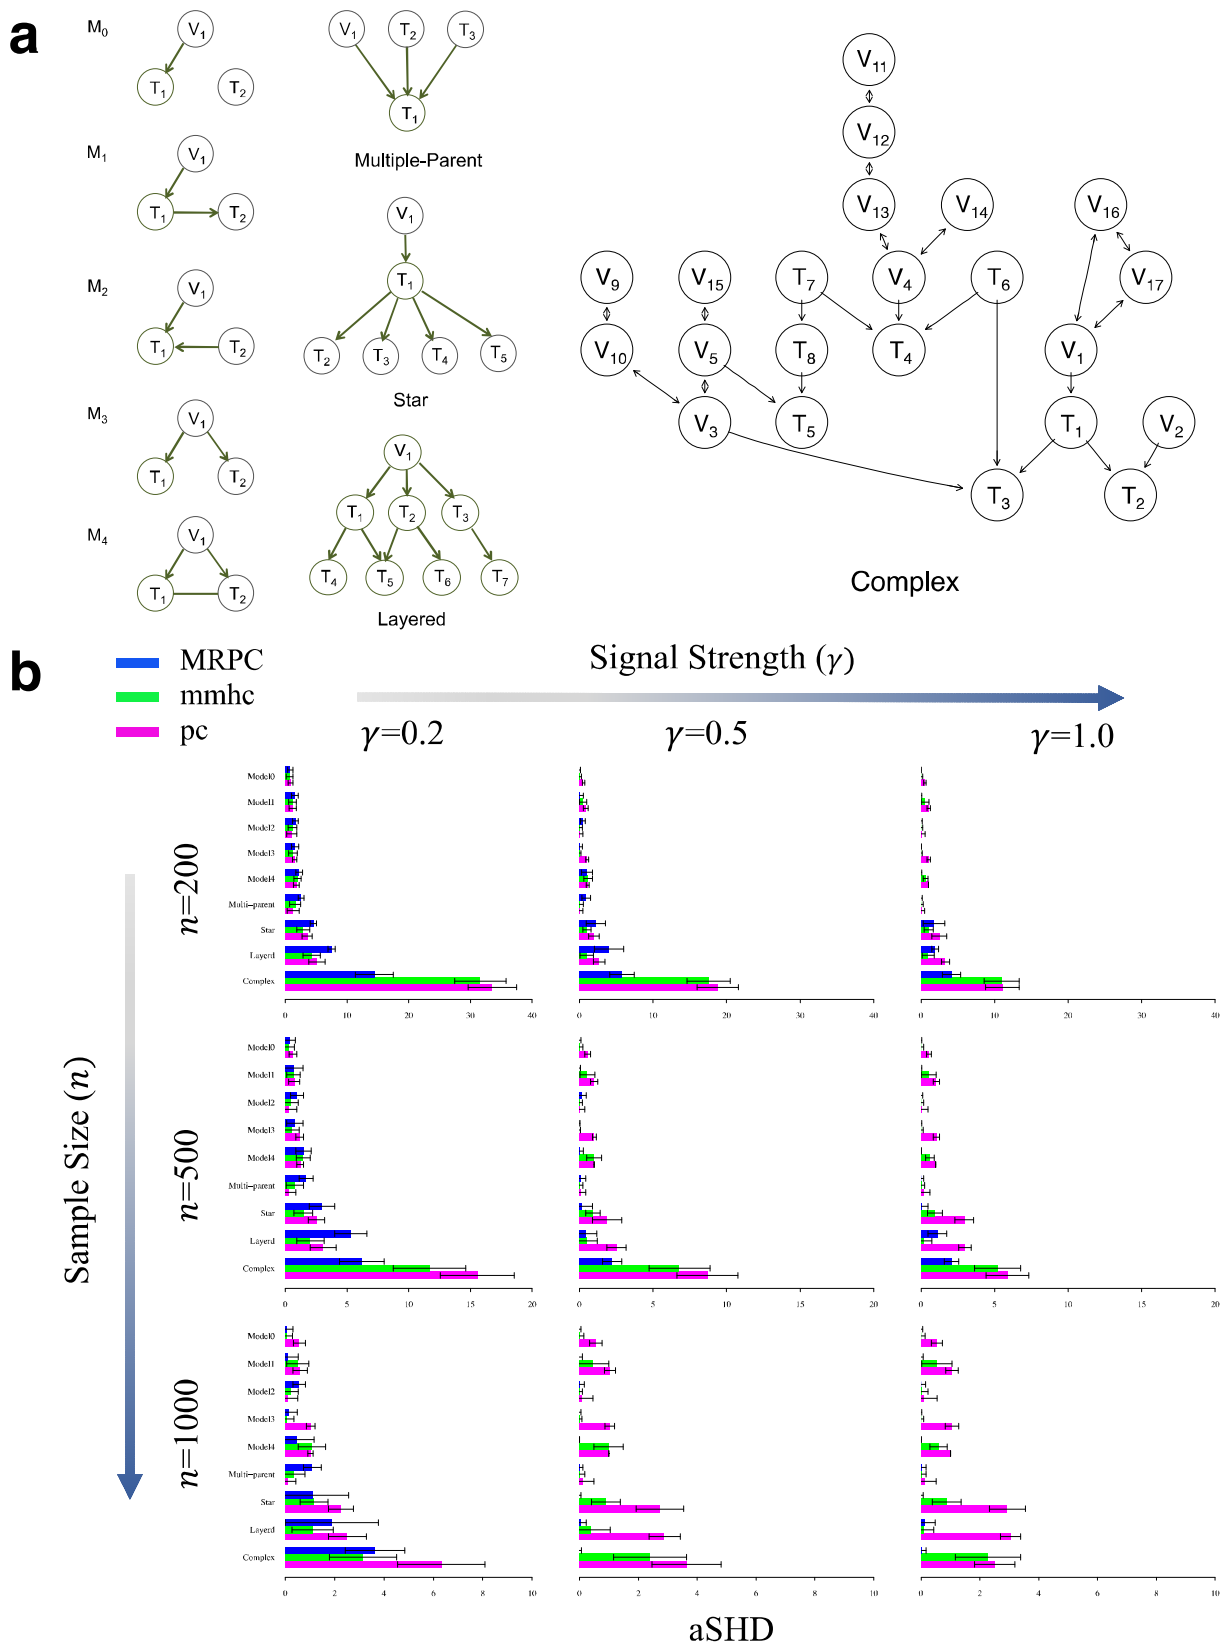
a
M0
M1
M2
M3
M4
Multiple-Parent
Star
Layered
Complex
b
MRPC
mmhc
pc
Signal Strength (γ)
γ=0.2
γ=0.5
γ=1.0
Sample Size (n)
n=200
n=500
n=1000
Model0
Model1
Model2
Model3
Model4
Multi-parent
Star
Layerd
Complex
aSHD

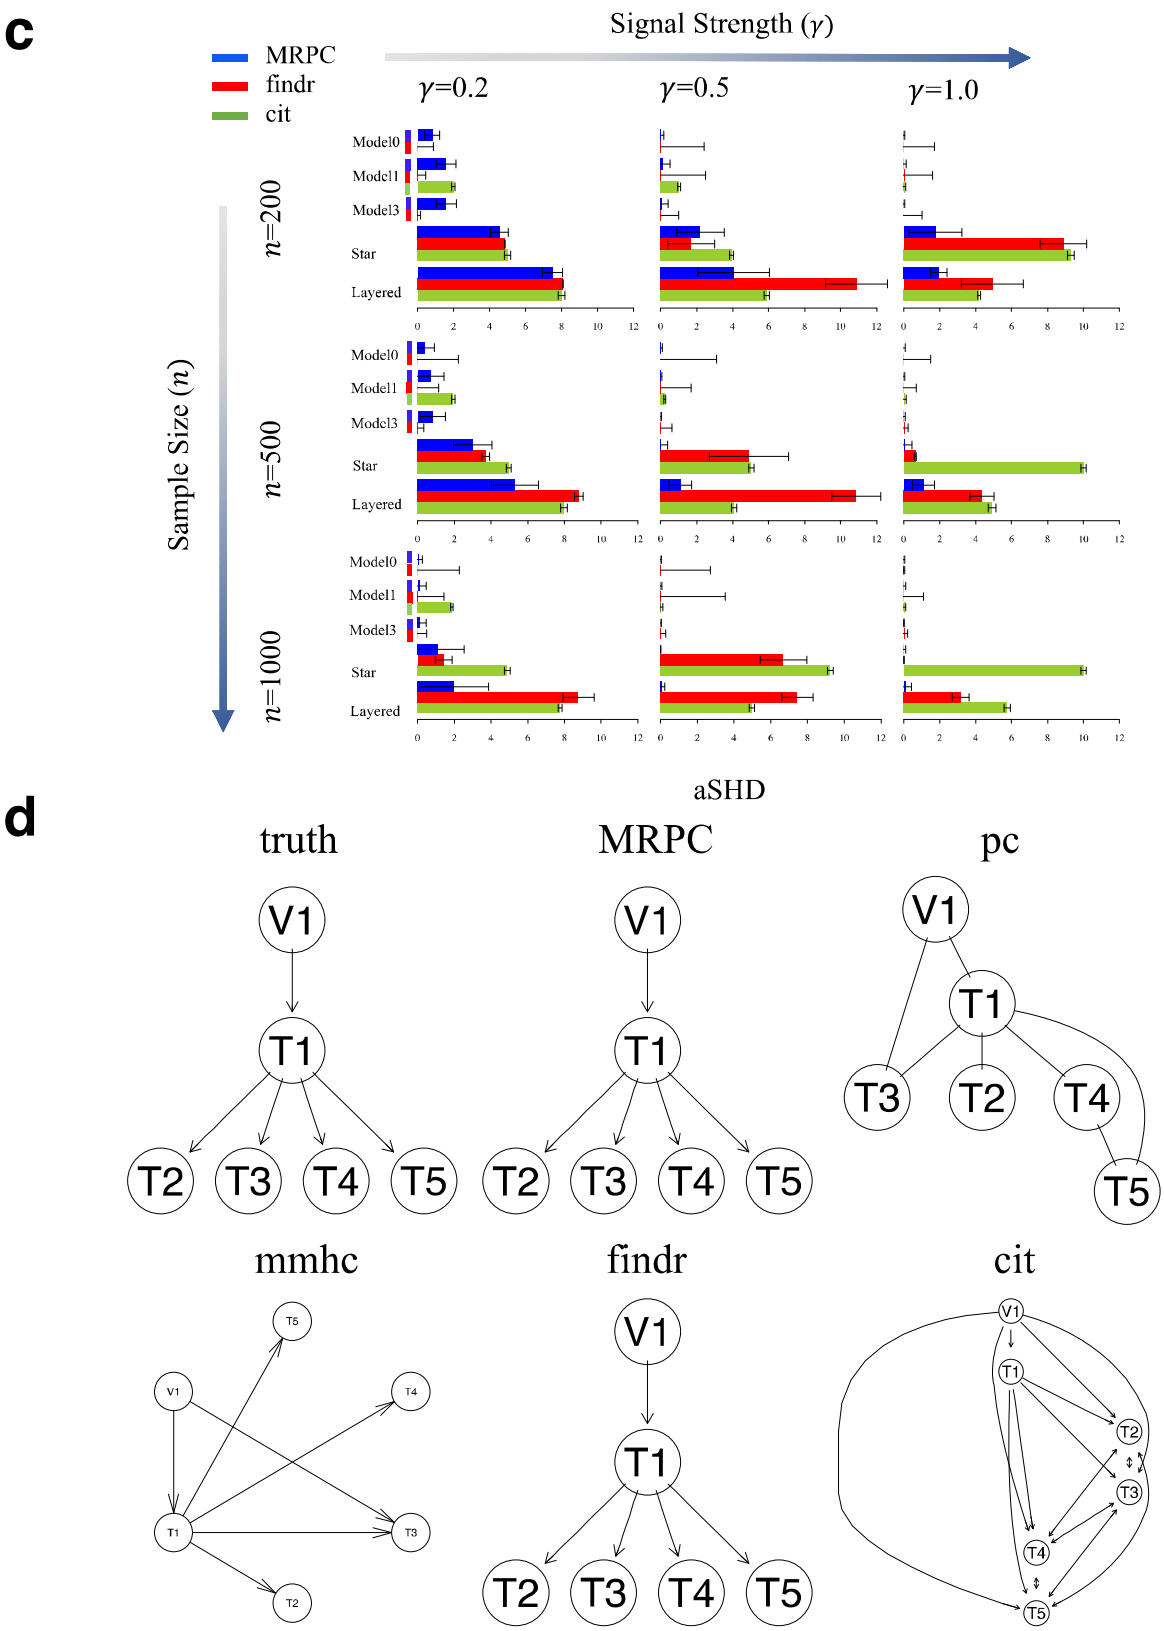
c
Signal Strength (γ)
MRPC
findr
cit
γ=0.2
γ=0.5
γ=1.0
Sample Size (n)
n=200
n=500
n=1000
Model0
Model1
Model3
Star
Layered
aSHD
d
truth
MRPC
pc
mmhc
findr
cit
V1
T1
T2
T3
T4
T5

Figure 2: Performance of MRPC on synthetic data compared with other methods. **(a)** Topologies used to generate synthetic data (see Methods for details of simulation). **(b)** Mean adjusted Structural Hamming Distances (aSHDs) for MRPC and two other network inference algorithms, namely, mmhc from bnlearn and pc from pcalg. For each of the topologies in (a), 1000 datasets were generated for different signal strengths (γ, which is the coefficient of parent nodes in the linear model; see Methods for simulation details) and different sample sizes (n). Each of the three methods was applied and the aSHD was calculated for the inferred graph relative to the truth. The mean distance over 1000 datasets is represented by the horizontal bar, and the standard deviation the black segment on the bar. Smaller aSHD indicates higher accuracy. **(c)** aSHDs for MRPC and two other MR-based methods, findr and cit. Since findr and cit can learn only some of the topologies in (a), only these topologies were used to simulate the data. Simulation and aSHD calculation are similar to that described in (b). **(d)** An example of inferred graphs from all five methods on data simulated under a star model with a large sample size (n =1000) and strong signal ($\gamma = 1.0$).

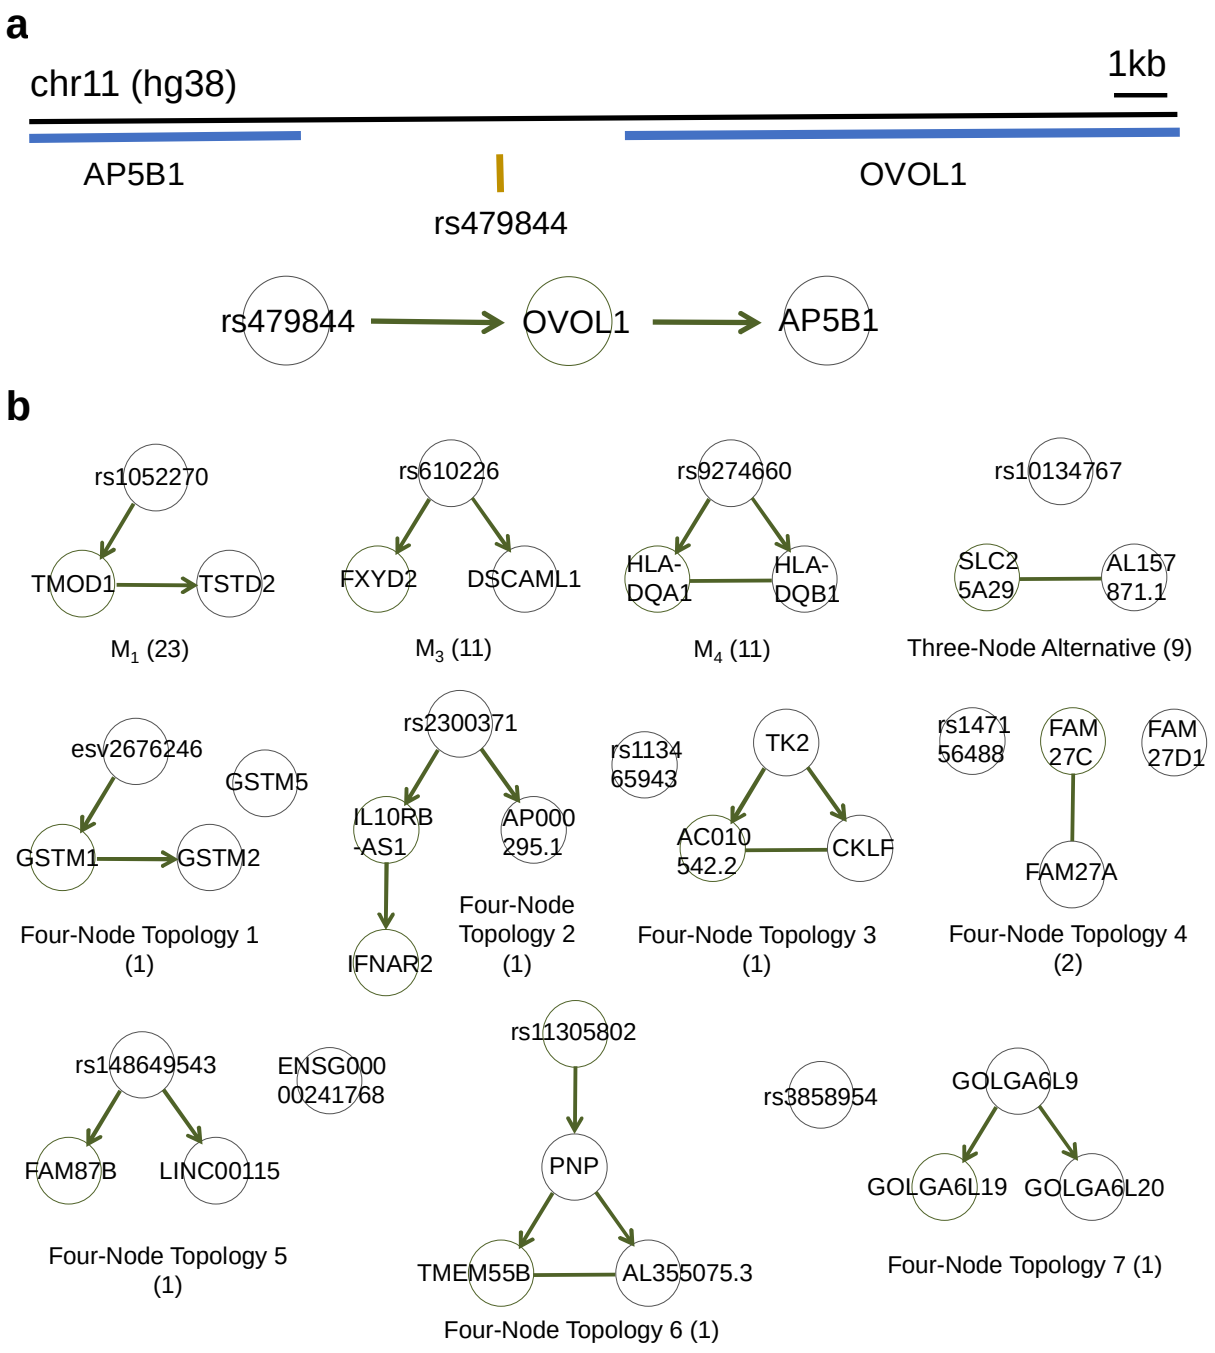

Figure 3: MRPC distinguishes direct and indirect target genes of eQTLs in the GEUVADIS data for the European cohort. **(a)** rs479844 is a GWAS significant SNP for atopic march in the GWAS Catalog, and an eQTL identified in GEUVADIS for two genes. (**b**) MRPC learns 11 distinct topologies among associated genes for eQTLs. The number in parentheses are the

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

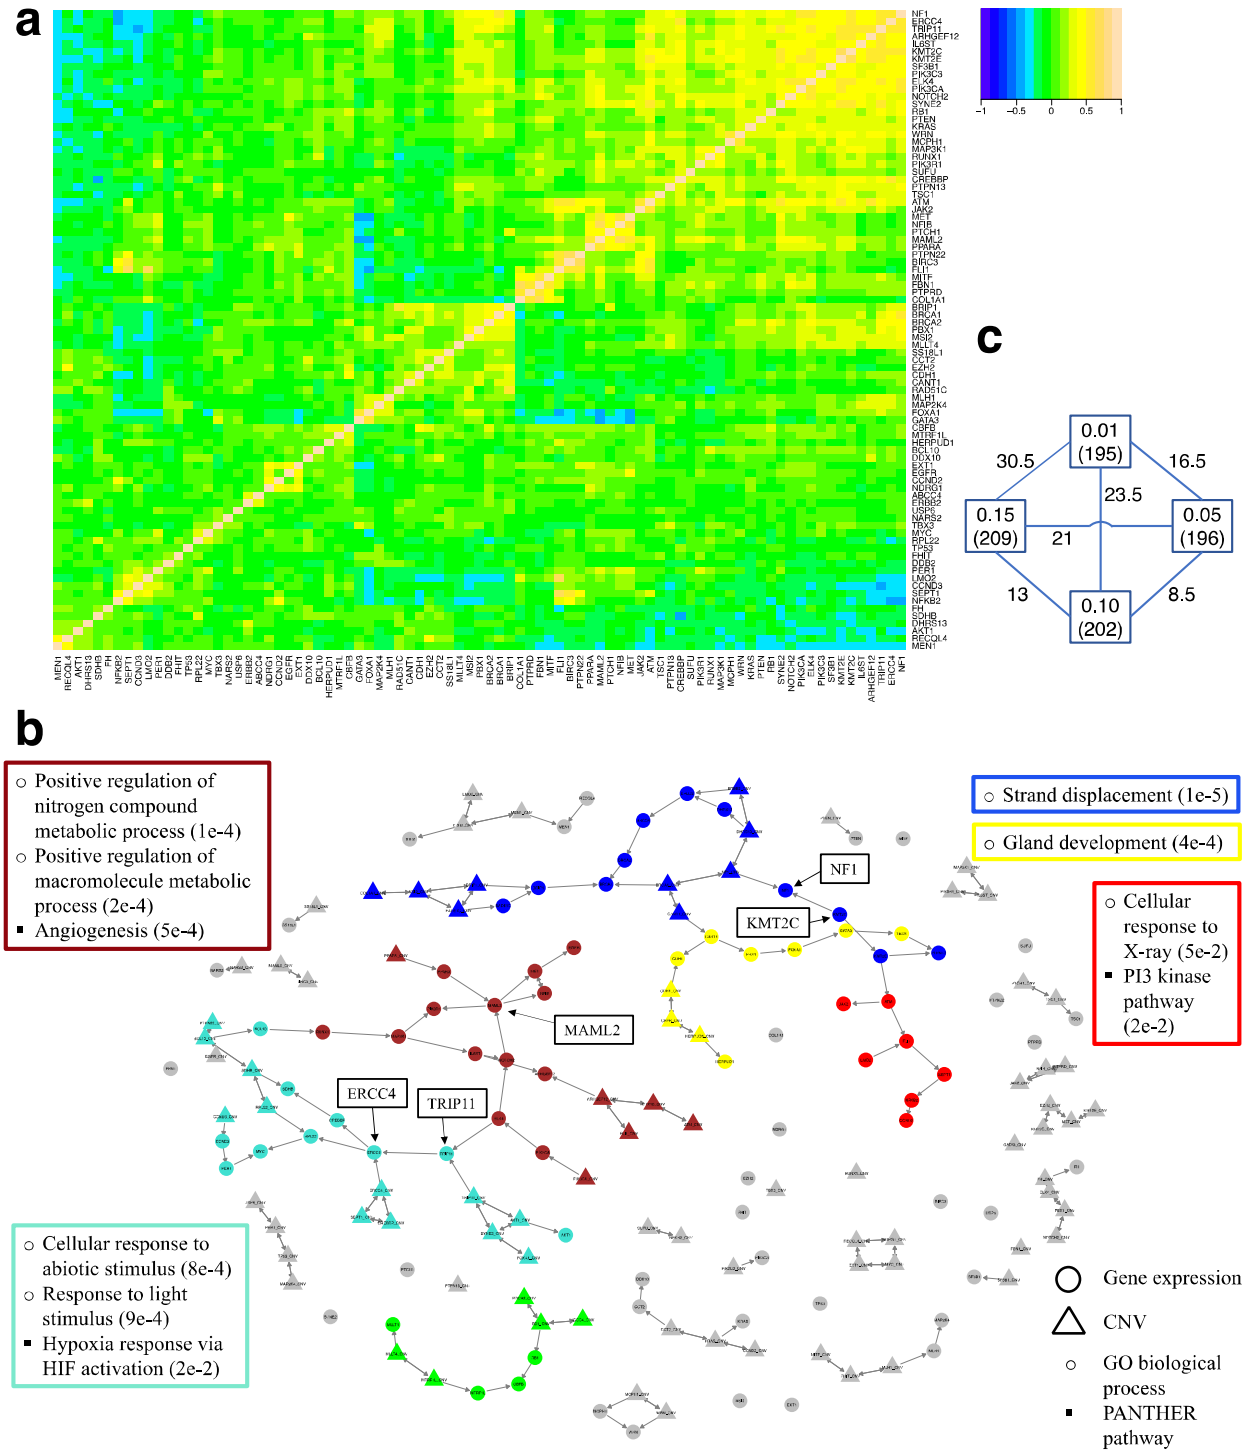

Figure 4: MRPC learns a causal regulatory network for frequently altered cancer genes using the TCGA breast cancer data. **(a)** Pearson correlation heatmap for the 85 genes with hierarchical clustering in rows and columns. **(b)** The causal network inferred at FDR of 5% by

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 7, 8**). **(c)** The aSHD 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 aSHD. These numbers demonstrate the stability of the MRPC inference; see main text for detail.