NAR Genomics and Bioinformatics, 2026, **8**, lqag004
https://doi.org/10.1093/nargab/lqag004
Methods Article

OXFORD

A novel explainable deep-learning approach for network analysis of epistatic interactions

Andrea Mastropietro [1,2,*], Georgios Markopoulos [3], Evangelos Evangelou [4,5,6], Aris Anagnostopoulos [7,*]

[1]Department of Life Science Informatics and Data Science, B-IT, LIMES Program Unit Chemical Biology and Medicinal Chemistry, Rheinische Friedrich-Wilhelms-Universität, Friedrich-Hirzebruch-Allee 5/6, 53115 Bonn, Germany
[2]Lamarr Institute for Machine Learning and Artificial Intelligence, Friedrich-Hirzebruch-Allee 5/6, 53115 Bonn, Germany
[3]Neurosurgical Institute, University of Ioannina School of Medicine, 45110 Ioannina, Greece
[4]Department of Hygiene and Epidemiology, University of Ioannina Medical School, 45110, Ioannina, Greece
[5]Biomedical Research Institute, Foundation for Research and Technology–Hellas, 45110 Ioannina, Greece
[6]Department of Epidemiology & Biostatistics, School of Public Health, Imperial College London, W2 1PG London, United Kingdom
[7]Department of Computer, Control, and Management Engineering, Sapienza University of Rome, Via Ariosto 25, 00185 Rome, Italy

*To whom correspondence should be addressed. Email: aris@diag.uniroma1.it
Correspondence may also be addressed to Andrea Mastropietro. Email: mastropietro@diag.uniroma1.it

Abstract

Epistatic interactions of gene loci often determine complex trait phenotypes and may indicate the underlying molecular mechanisms of traits and diseases. Yet, the inference of epistatic interactions and gene–gene networks remains challenging. Neural networks have become successful in classifying complex data, revolutionizing various fields. However, their complexity does not reveal how they combine input features, and this lack of interpretability limits their use with genetic data. We thus introduce EpiDetect, a novel framework for discovering interactions between input features (single-nucleotide polymorphisms—SNPs—in our setting). EpiDetect neural-network–based classifiers detect interactions in systolic, diastolic, and pulse pressure genome-wide association datasets. Central to EpiDetect is EpiCID, a novel explainability algorithm for neural networks. Using EpiCID, we identified a network of highly interactive SNPs, performed centrality analysis to pinpoint central SNPs and genes, and outperformed established epistatic-interaction-detection algorithms. Moreover, pathway analysis uncovered well-known and novel pathways that could play a significant role in blood pressure traits, opening up new research directions.

Introduction

Recent advances in genome-wide association studies (GWAS) [1] and the significant increase of the available samples have helped identify thousands of robustly disease-associated loci (DNA positions where genes or genetic markers are located) and have provided novel insights about biological pathways underpinning complex diseases and traits. GWAS use the complete DNA set of individuals to determine whether the presence of a mutated gene correlates with a disease. Typically, this analysis is conducted with statistical instruments [2] using large-scale databases. Among genetic mutations, we find single-nucleotide polymorphisms (SNPs), which are substitutions of single nucleotides in a given locus. Nucleotides, which are adenine (A), cytosine (C), guanine (G), and thymine (T), are the basic building blocks of DNA. SNPs may be involved in the mechanisms of diseases or determine specific traits. Despite the increase of this kind of studies, the extent to which potential epistatic (i.e. SNP–SNP or gene–gene) interactions [3] contribute to complex networks that dictate the molecular mechanism underlying phenotypes remains an area of ongoing investigation. Epistatic interaction (or epistasis) is the phenomenon where different genetic loci contribute to a given phenotype in a cumulative nonadditive manner, and it refers to the modification of the effect of an allele exerted by other alleles. An allele is a version of a gene or DNA sequence at a given locus. An individual inherits two alleles, one from each parent, which will determine the genotype, which can be homozygous (same allele) or heterozygous (different alleles). In epistasis, interactions between genes influence the expression of a phenotype in a way that cannot be predicted solely based on the individual effects of each gene. It is a common occurrence in genetics and can have significant implications for understanding the genetic basis of complex traits and diseases that are not only governed by the mutation of one single gene. In this regard, resources such as OLIDA have catalogued hundreds of curated digenic and oligogenic variant combinations, mostly in the context of rare inherited diseases [4]. However, the extent to which such interactions contribute to the polygenic architecture of complex human traits remains unresolved, partly because studying epistatic interactions is particularly challenging given the large number of potential gene combinations and the complexity of the underlying genetic networks.

Received: August 23, 2024. **Revised:** December 15, 2025. **Accepted:** December 22, 2025

Researchers have developed several approaches to examine epistatic interactions in genetics [5]. One of the most known and effective interaction detection methods is Multifactor Dimensionality Reduction (MDR) [6]. It constructs a single attribute from variables, reducing data dimensionality. Genotypes at different loci are pooled into high- and low-risk groups. This reduces the dimensionality by creating a multilocus genotype variable, which can predict the disease status and assess the joint effect of the merged genotypes. Unfortunately, MDR suffers from a high computational cost caused by the elevated number of possible genotype combinations. Another approach, BOolean Operation-based Screening and Testing (BOOST) [7], relies on a logistic-regression model that considers both marginal and pairwise interaction effects. Relying on an efficiently processed Boolean representation of the data for computational efficiency, this method can detect epistatic pairs quickly. Other proposed models are based on random forests and Bayesian networks [8–10].

A caveat of these approaches is that they are influenced significantly by marginal (or main) effects. Those indicate the effect a single genetic variant alone has on the trait under consideration. Marginal effects can confuse the models and make them believe that a merely additive phenomenon caused by two high marginal effects can be a nonadditive (epistatic, so nonlinear) interaction. Marginal effects can be tricky because they may lure algorithms into believing that a nonlinear interaction exists, whereas it may just be the product of two additive effects that act independently. Indeed, algorithms for which the presence of additive effect is problematic will output interaction rankings in which a single SNP appears as interacting with a large number of other SNPs, which is not very likely. This is a significant limitation, as it can lead to the presence of false positives in the interactions. We argue that methodologies able to detect epistatic effects without being influenced by marginal effects will output a more diversified set of interactions, which is more biologically plausible; see the "Results" section. As we show, by analyzing the output results, our approach achieves exactly that.

To address the aforementioned problems, in this paper we develop a framework based on neural networks, enriched with an explainable artificial intelligence (XAI) component.

Despite their superior performance in various applications, neural networks have a significant drawback, which limits their use in genetics [11]: the complicated, highly nonlinear functions they learn make it difficult to understand how they operate. Thus, they are almost always used as black boxes. Whereas this usage is acceptable for most applications, such as image classification or text generation, it is unsuitable when it is needed to know *why* a particular output has been produced, as in medical and genetic applications. Even though neural networks, being used as black boxes, are not inherently interpretable (as simpler tree-based models, for instance), there is some line of work on explainability techniques that aim at unveiling the relationship between input and output. This can be achieved by looking for input features that have a heavier weight in the prediction [12, 13] (feature selection/importance problem) or by trying to assess whether there are groups of features whose joint interaction is responsible for the values of the output (feature interaction detection) [14]. Our method focuses on the latter, as this kind of explainability approach can help model and determine epistatic interaction effects.

With this spirit, this work introduces EpiDetect, a new framework for discovering potential epistatic interactions. The core element of EpiDetect is EpiCID (Epistatic Cosine Interaction Detection), a neural-network–based algorithm that opens the black box and leverages the power of neural networks to discover complicated and hidden interactions between input SNPs, actually explaining the neural network predictions. EpiCID is a method specialized at directly detecting purely interacting input features without relying on the marginal effect that a single SNP may have on the trait under consideration. Our method is designed to deliver global explanations. In contrast with most XAI methodologies, which are local, instance-based explainers that discover critical features for specific predictions [15], EpiCID determines the effect (in terms of interaction strength) that pairs of features have globally on the behavior of the model. Given the presence of consistent GWAS studies in the field, we selected blood pressure regulation as a case study [16].

Materials and methods

In this section, we will describe the data processing steps and the EpiDetect framework, which can be divided into three main components. The first one, EpiCID, takes a trained neural network as input and outputs interacting pairs of SNPs/genes. The second component, the network analysis, retrieves central genes in the epistatic network. The third component, the enrichment analysis, looks for pathways or ontologies associated with the disease or trait under study.

Data sources and processing

For our study, we retrieved individuals' genotype and phenotype information from UK Biobank [17, 18], a large population-based cohort in the UK that includes around half a million volunteers aged 40 to 69 years. For our analysis, we selected SNPs that were found to be robustly associated with systolic (SBP), diastolic (DBP), and pulse pressure (PP), taking into account genome-wide significant signals for those traits derived from the largest GWAS on blood pressure [16] (more details on SNP selection in Supplementary data). This allows us to narrow down the pool of possible epistatic SNPs, reducing the complexity associated with those studies. Thus, we ended up having 264 SNPs for SBP, 342 for DBP, and 283 for PP, with no overlap among the three sets. We created three datasets, one for each trait, in which an array of SNPs describes each individual, and the target variable is the blood pressure value of interest. For each patient, we calculated the mean SBP, DBP, and PP values from two automated or, wherever needed, manual measurements, as reported in the UK Biobank dataset. We further performed quality control, excluding individuals with >10% of missing genotype. After the quality control, we ended up with 456 057 individuals. Every patient in the three datasets is described by an array of SNPs. The SNPs were encoded in a one-hot encoding fashion to be suitable to be fed to a neural network model. Each SNP is a three-element vector in which the one-valued element indicates the genotype:

- [1,0,0]: homozygous for the major allele.
- [0,1,0]: heterozygous.
- [0,0,1]: homozygous for the minor allele.

The major allele is the most frequent in the population, whereas the minor is the least occurring. For details on feature vector distributions, see Supplementary Fig. S7.

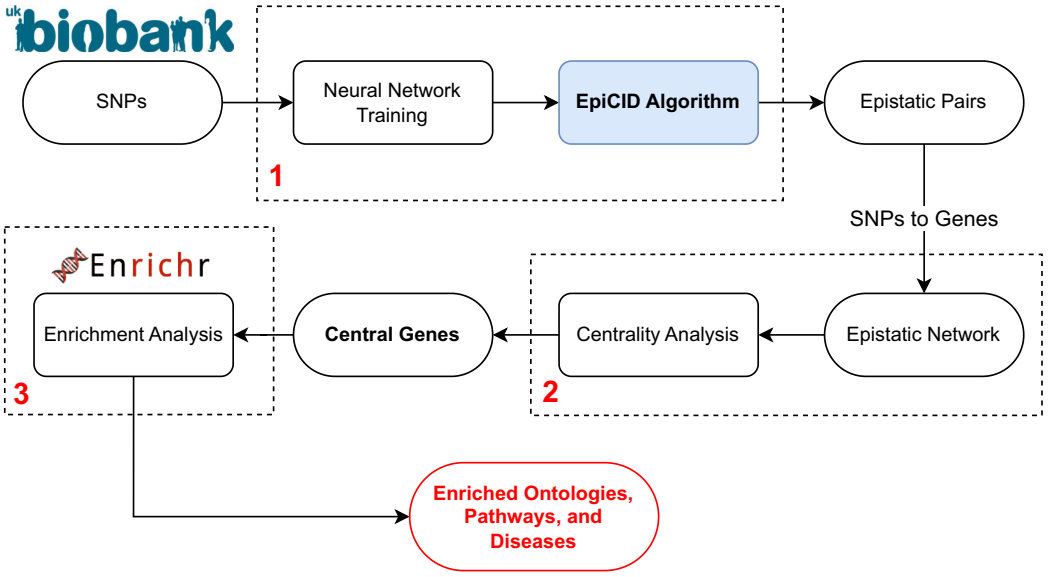

Figure 1. The workflow of the EpiDetect framework. The numbers in the blocks indicate the three components of the pipeline. A neural network is trained and explained with EpiCID (block 1) to find interacting pairs that are mapped to genes and used to build an epistatic network and perform centrality analysis (block 2). The central genes are used for enrichment analysis (block 3), which determines associated ontologies, pathways, and diseases.

EpiDetect framework

Given the genotype (SNPs) and phenotype (blood pressure value) of the subjects (patients), our system identifies SNPs correlated with the phenotype expression based on SNP–SNP epistatic interactions. To do that, we designed three regression models based on a multilayer perceptron (MLP) [19] for detecting systolic SBP, DBP, and PP values, and we evaluated the weights of the layers inside the neural networks and the effect that these weights have on the final output (the blood pressure value). The use of this simple type of neural network is aligned with the necessity in genetics to develop simple and explainable models rather than complex and obscure ones [20]. For each pair of SNPs (our method can also be extended to subsets of more than two SNPs), we obtained an *interaction score*, which captures the strength of the interaction between the two SNPs of the pair, as well as the effect that this interaction has on the final regression function. Then, each SNP is mapped to the corresponding gene. Subsequently, we identified the genes/SNPs with the most interactions by creating a network and connecting those with a high interaction score. From this network, we identified the most central genes that significantly affect SBP, DBP, or PP. This allowed us to discover potentially novel pathways that affect blood pressure and cardiovascular risk.

As mentioned, EpiDetect consists of three main components: (i) the EpiCID neural-network explainability module that detects candidate epistatic pairs, (ii) the design of a gene–gene network and centrality analysis, and (iii) enrichment analysis. The last is a method that aims at identifying genes that are over-represented in biological pathways, gene ontologies, or disease networks, and that can have associations with specific phenotypes. We will now describe each component in detail, and a graphical representation of the workflow of our approach is shown in Fig. 1.

Neural-network model and training

Our methodology is based on the use of an MLP. Our network is composed of two fully connected layers with 200 and 50 neurons, followed by a dropout layer (probability of dropout set to 0.3). For each blood-pressure trait (SBP, DBP, and PP), the network was trained for 40 epochs using the Adam optimizer [21] with a learning rate of 10^{-3} and a batch size of 16, generating three different models. The three blood-pressure datasets were split into training (70%) and test sets (30%). The average mean absolute error with standard deviation over five independent runs, measured in mmHg (the unit of measurement for blood pressure), was 14.83 ± 0.03, 8.04 ± 0.02, and 8.79 ± 0.03 for SBP, DBP, and PP, respectively (details in Supplementary Fig. S8). The trained neural network is the base for the explanation phase. As a note, even though the regression error may seem high, it is important to point out that, with this study, we are capturing only the genetic components of blood pressure regulation without taking into account other important factors, such as patients' habits and lifestyles, which significantly affect the traits.

Epistatic Cosine Interaction Detection

The core of EpiDetect is the Epistatic Cosine Interaction Detection algorithm (EpiCID). Our idea is that to find interacting pairs of input features, it is possible to consider their correlation in the vector space. Thus, one possible metric to use for measuring feature correlation is cosine similarity. From this follows the need to represent a feature as a vector, using a sort of "feature vector of a feature." This vector should describe how the feature is represented and seen internally by the neural network. It can be thought of as a fingerprint or embedding of the feature itself. We can build this vector by using the weights that the neural network learned during its training. To construct this vector, we consider the first-layer weights (the ones connecting the input with the first-layer neurons) and the influence that hidden units exert on the output. The latter can be taken into account using the concept of *aggregated weight* introduced in Neural Interaction Detection (NID) [14]. This is computed by cumulative matrix multiplications of the absolute values of weight matrices and indicates the influence that hidden units have on the output. The aggregated weights of

the units at layer ℓ are defined as

$$\mathbf{z}^{(\ell)} = \left|\mathbf{w}^{y}\right|^{\mathrm{T}} \cdot \left|\mathbf{W}^{(L)}\right| \cdot \left|\mathbf{W}^{(L-1)}\right| \dots \left|\mathbf{W}^{(\ell+1)}\right|, \tag{1}$$

where $\mathbf{w}^{y}$ are the weights at the output layer, $\mathbf{W}^{(j)}$ is the weight matrix at the j^{th} hidden layer, L is the network depth (i.e. the number of hidden layers), and $|\cdot|$ denotes elementwise absolute values. Thus, each element $z_k^{(\ell)}$ indicates the aggregated weight at unit k of layer ℓ. As demonstrated [14], this definition is an upper bound of the gradient magnitude. It can be used as an approximation of the importance of the hidden units, as gradients have been commonly used as importance measures in neural networks [12, 22, 23]. Now that we have all the elements needed, we can define the *neural feature vector* $\boldsymbol{\xi}^{(i)}$ for input feature i as follows:

$$\boldsymbol{\xi}^{(i)} = \left|\mathbf{W}_i^{(1)}\right| \odot \mathbf{z}^{(1)}, \tag{2}$$

where $\odot$ indicates the Hadamard (elementwise) product of two vectors, $\mathbf{W}_i^{(1)}$ is the vector of the weights between feature i in the input and the first hidden layer units, and $\mathbf{z}^{(1)}$ is the aggregated weight vector at the first hidden layer, as defined in Equation (1). Every element $\xi_k^{(i)}$ is given by the product between the connection weight from feature i to unit k of the first hidden layer and the aggregated weight at the unit. The vector defined above is an embedding of the feature, a hidden representation created by the network during the learning process, implicitly considering its contribution to the output throughout the network.

Cosine Interaction Strength

Now, we know how to represent a feature in the vector space as embedded by the neural network. We can exploit this to compute the interaction strength between two or more features. We decided to consider the correlation of the feature in the vector space as a measure of interaction importance using the cosine similarity measure, defining the *cosine interaction strength* between features i and j as

$$\phi(i, j) = \cos(\boldsymbol{\xi}^{(i)}, \boldsymbol{\xi}^{(j)}) = \frac{\boldsymbol{\xi}^{(i)} \cdot \boldsymbol{\xi}^{(j)}}{\left\|\boldsymbol{\xi}^{(i)}\right\| \left\|\boldsymbol{\xi}^{(j)}\right\|}, \tag{3}$$

where $\boldsymbol{\xi}^{(i)}$ and $\boldsymbol{\xi}^{(j)}$ are the neural feature vectors describing features i and j. Notice that $\phi(i, j) \in [0, 1]$ since $\xi_k^{(i)} \geq 0$ and $\xi_k^{(j)} \geq 0$ for every k, as per Equations (1) and (2). The rationale behind this measure is that a high cosine similarity between two neural feature vectors indicates that two features interact in the network. However, this measure alone does not capture the real magnitude of the interaction, but it is only able to detect if an interaction is present. Early experiments showed that this measure could find interacting pairs of features but failed to determine the interaction strength properly. For this reason, we improved the definition of cosine interaction strength, defining a more appropriate one that was able to differentiate between weak and strong interactions.

To do that, we take into consideration where the interaction is formed, namely, between the input and the first hidden layer. Inspired by NID, given features i and j, we consider their connection weights to the first hidden layer units. Differently from Equation (2), in which the effect of a single feature on the network is considered, we now analyze the joint effect of the two features on the first hidden layer. For each first-layer unit k, we apply the min averaging function to the absolute values of the weights connecting k to i and j. The choice of this function is justified by previous work [14], which demonstrated its effectiveness in detecting feature influence in neural networks. Then, we sum over all the units to obtain what we define as the *first-layer interaction influence*:

$$\eta(i, j) = \sum\nolimits_k \min(|W_{i,k}^{(1)}|, |W_{j,k}^{(1)}|). \tag{4}$$

The rationale behind the choice of the minimum is that an interaction is strong (at the first hidden layer) when its η is large. When the minimum of the two weights is large (i.e. when both weights have a high value) for a consistent number of hidden units, η will be high, indicating the rise of strong interaction. Conversely, the previously defined ϕ in Equation (3) helps understand how the interaction evolves when passing throughout the network until the output: an interaction may start strong at the first hidden layer but lose its importance toward the output layer, or a mild interaction may preserve its strength, having a considerable effect on the prediction. Given those observations and merging Equations (3) and (4), we obtain the new definition of cosine interaction strength:

$$\Phi(i, j) = \eta(i, j)\phi(i, j). \tag{5}$$

The measure defined in Equation (5) can describe the interaction completely and coherently, from its rise from the input to the first hidden layer and considering its evolution toward the output through the network.

Finally, in our application, given that an SNP is represented by a three-element one-hot–encoded vector, it follows that it is described by three different neural feature vectors (one for each feature), which need to be aggregated before computing the interaction strength. We merge them into a single one by performing an element-wise sum of the vectors. The same aggregation is used for the first-layer weights to compute the first-layer interaction influence of a SNP. Once we have the vectors, it is possible to obtain the cosine interaction strength for a pair of SNPs as in Equation (5). Notably, EpiCID can be extended to sets of more than two SNPs. A visual representation explaining the EpiCID algorithm is shown in Fig. 2.

Network and centrality analysis

After training the MLP three times (once for each trait: SBP, DBP, and PP) and applying EpiCID, we obtain a ranked list of highly interacting (as measured by EpiCID) pairs of SNPs. We next map these SNPs to genes, obtaining a ranking of interacting genes. The mapping is performed via the following steps:

(1) We map an SNP to a protein-coding gene according to the Consensus Coding Sequence Project (CCDS) [24], which provides very high-quality annotations for protein-coding genes.
(2) If no reference in CCDS is found, we map it using Genecode [25].

In cases where an SNP is annotated in a genome region with more than one gene present, we follow the process:

(1) We map to a coding gene if present.
(2) If more than one coding gene is present, we choose a mapping already appearing in the literature.
(3) If no coding genes are present, we look up different sources, such as dbSNP [26] and Ensembl [27], to choose the most suitable mapping.

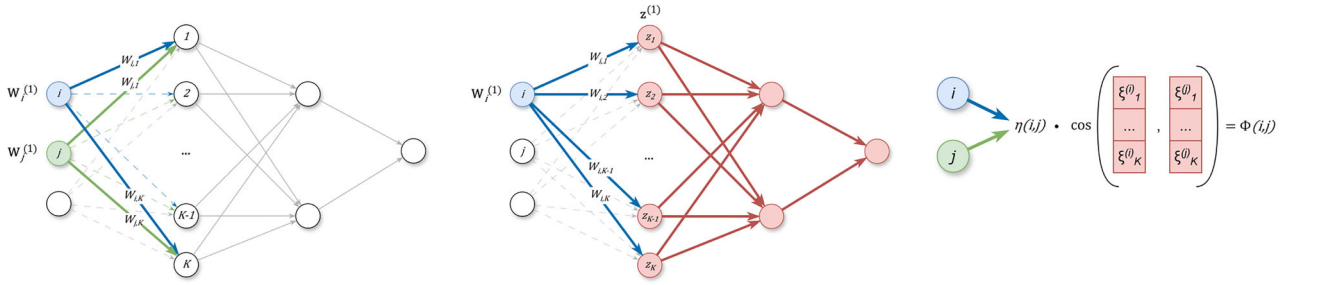

Figure 2. Graphical representation of the EpiCID algorithm steps. (**a**) Computation of the first-layer interaction influence $\eta(i, j)$, which evaluates the joint effect of two features on the first hidden layer by applying the min function to their connection weights and summing over all units. (**b**) Construction of the neural feature vector $\boldsymbol{\xi}^{(i)}$, which represents feature i in the vector space by combining the first-layer weights and aggregated weights through an element-wise product. (**c**) Calculation of the cosine interaction strength $\Phi(i, j)$, which combines the first-layer interaction influence $\eta(i, j)$ and the cosine similarity $\phi(i, j)$ of the neural feature vectors to quantify the interaction strength between two features. For clarity, the superscript (1) indicating the first hidden layer is omitted in the figure.

In both cases, whenever we do not find a mapping, we assign the SNP to the closest gene according to Genecode, giving a preference for coding genes.

This mapping ensures that we do not discard any SNPs to avoid losing important information for the traits and miss possible interactions between coding and noncoding genes, as the latter may contain regulatory elements that are of paramount importance in complex traits [28].

We use the interactions identified by EpiCID to create an epistatic interaction gene–gene network for each of the three blood pressure traits and discover the most central genes by studying the top 1000 interactions [29, 30]. For comparison purposes, to investigate the effectiveness of the proposed methodology, we also analyzed the central networks obtained by using other algorithms for epistatic interaction detection. We choose the central genes, that is, those with a degree higher than the average degree in the top-1000 network, for enrichment analysis. The degree, corresponding to the number of direct neighbors to a given node, allows an immediate evaluation of the regulatory relevance of a node and can be used to validate centrality in different kinds of networks, such as signaling and metabolic networks [31].

Enrichment analysis

We perform enrichment analysis for central genes of the epistatic networks using the Enrichr online tool [32–34]. This is among the most reliable, robust, and popular tools for this task, thanks to the large amount of data sources it uses for its analyses. In our study on blood pressure, we compared gene sets from the different algorithmic approaches to the datasets available in the databases of the Gene Ontology Consortium [35], DisGeNET discovery platform [36–38], GWAS Catalog [39], and UK Biobank [17, 18], to understand the biological relevance of the epistatic networks obtained for the three blood pressure traits.

Even though Enrichr models monogenic interactions, its results remain highly relevant in interpreting the functional implications of genes involved in epistatic effects. Performing pathway analysis allows us to determine whether these genes converge to common biological processes relevant to the complex traits under study. Moreover, enrichment analysis helps us understand potential collective biological roles, providing an important biological context for interpreting our findings. By leveraging Enrichr, we aim to determine whether our identified genes are overrepresented in key biological functions relevant to complex traits.

Point penalty score

To quantify and evaluate the coherence of the retrieved epistatic networks with the traits of interest and compare them with one another, we defined a *point penalty score*, which penalizes the approaches that perform the worst, considering the ranking among the proposed strategies. The method ranking the trait higher than the others is awarded 0 penalty points, the second method receives 1 point, and so on (ties are awarded the same penalty points). This is done for all the traits and the data sources considered. Eventually, the scores are summed to obtain the final point penalty score. A lower score is better. No penalty is assigned if no method ranks the trait among the top 10 enriched terms. We introduced this metric to be able to evaluate the epistatic networks retrieved by different algorithms in a comparative manner.

Results

We hereby present the application of the EpiDetect pipeline. First, we show the application of EpiCID on synthetically generated data to evaluate its performance in the presence of a ground truth. Second, we apply our framework to the three blood pressure traits (SBP, DBP, and PP). To compare our approach, we evaluated it against other widely used epistasis detection frameworks such as MDR and BOOST. We also compared it with NID, a methodology specific to neural networks that we adapted to work with genetic inputs. Our results indicate that our strategy outperformed these methods and minimized the influence of marginal effects. Regarding real-world blood-pressure data, given that validating epistatic interactions is rather challenging because of the complex nature of the phenomenon and the absence of consistent ground truth, we proceed with a three-fold analysis to have a thorough evaluation of our methodology. At first, we analyze the EpiCID output to investigate its reliability and robustness, mainly focusing on marginal effects and interaction distributions. Then, we describe the network and centrality analysis of the interacting pairs found by EpiCID and compare the central genes obtained against the results of other algorithms. Finally, we

Table 1. Top 20 interactions for the GAMETES additive model (reporting also marginal effects—ME—for BOOST and MDR)

Rank	BOOST		BOOST ME	MDR		MDR ME	NID		EpiCID	
1	M0P1	M0P2	M0P2	M0P1	M0P2	M0P2	M0P1	M0P2	M0P1	M0P2
2			M0P1	N6	M0P2	M0P1	N32	M0P1	N207	N236
3			N24	N40	M0P2	N24	N162	M0P1	N144	N182
4			N68	N101	M0P2	N222	N101	M0P1	N85	N236
5			N163	N0	M0P2	N111	N37	M0P1	N206	N236
6			N6	N1	M0P2	N70	N40	M0P1	N142	N162
7			N222	N2	M0P2	N214	N209	M0P1	N153	N207
8			N103	N3	M0P2	N5	N207	M0P1	N162	N188
9			N142	N4	M0P2	N99	N236	M0P1	N32	N110
10			N111	N5	M0P2	N89	N71	M0P1	N162	N236
11			N5	N7	M0P2	N161	N55	M0P1	N162	N207
12			N70	N8	M0P2	N243	N58	M0P1	N236	N251
13			N89	N9	M0P2	N68	N59	M0P1	N44	N236
14			N102	N10	M0P2	N196	N130	M0P1	N22	N236
15			N243	N11	M0P2	N119	N32	M0P2	N22	N209
16			N166	N12	M0P2	N228	N125	M0P1	N30	N162
17			N214	N13	M0P2	N103	N37	M0P2	N162	N197
18			N55	N14	M0P2	N179	N70	M0P1	N40	N207
19			N175	N15	M0P2	N181	N246	M0P1	N117	N236
20			N196	N16	M0P2	N110	N188	M0P1	N151	N242

The names of the SNPs are the ones generated by GAMETES; SNPs named M0P1 and M0P2, marked in yellow for convenience, are the interacting pair, and SNPs named Nxxx are randomly generated with no effect on the trait.

analyze the output of the entire framework, which is represented by the ontologies, pathways, and traits found with enrichment analysis on the central genes obtained from the EpiCID-derived epistatic pairs.

Application to synthetic data

We generated datasets composed of 500 000 individuals (half cases, labeled 1, and half controls, labeled 0) having 256 SNPs each. We created multiple datasets by using the GAMETES software [40] so as to take into account multiple interaction patterns, including (i) additive, (ii) threshold, (iii) epistatic, and (iv) purely epistatic interactions, based on previous work [10]. (i) The additive pattern is a simple linear phenomenon in which the risk of disease (label 1) increases with the number of minor alleles; this is not an interaction but an additive sum of effects, characterized by strong marginal effects of both SNPs in the pair. (ii) In the threshold model the risk is given by the presence of at least one minor allele; this is a step-function effect in which one SNP of the pair has a strong marginal effect. (iii) In the epistatic pattern the risk of disease is given by a nonlinear interaction of the SNPs involved. (iv) Finally, the purely epistatic effect is an extreme epistatic pattern in which all marginal effects of SNPs are minimized, in favor only of the nonlinear phenomenon. In the synthetic datasets, only one pair of SNPs is generated using those patterns, and the rest of the SNPs are noncausal for the trait and are generated randomly. The use of this type of data allows us to evaluate the output in the presence of a ground truth. More details on those patterns and the dataset generation are available as Supplementary data (see Supplementary Table S3).

Additive, threshold, and epistatic models

The application of EpiCID to the additive, threshold, and epistatic synthetic dataset models yielded analogous results. Therefore, we report only the rankings obtained on the additive-pattern dataset in Table 1 and discuss the corresponding findings; results for the threshold and epistatic models are provided in the Supplementary data.

In the additive scenario, BOOST shows no significant problems; the additivity is properly found as well as the marginal effects of the two SNPs. No additional SNPs are reported in the table for BOOST, because this method does not output pairs with low interaction scores.

MDR shows a behavior that is expected in the presence of additive effects. Even though the interaction is properly detected, the subsequent pairs feature one of the interacting SNPs. This is because the interaction is created by the additivity of the marginal effects of SNPs M0P1 and M0P2. MDR relies on those marginal effects when looking for interactions. Actually, the interaction scores of pairs ranked from positions 2 to 20 (and even beyond) featuring SNP M0P2 (M0P2–Nxxx) are the same score we have for the marginal effect of SNP M0P2: 0.5685 (see Supplementary Tables S26 and S27 in file `Supplementary_Tables_S6-S47.xlsx` for more details); this means that the marginal effect determines the whole interaction score. On the contrary, EpiCID is able to filter out marginal effects in additive phenomena. With EpiCID we note that the interaction is properly detected, but the ranking is not governed by the high marginal effect SNPs. Indeed, from the second position on, we do not notice any of the two interacting SNPs M0P1 or M0P2, despite their marginal effect being high. The SNP pairs returned from position two and below are apparently random/nonimportance-related (no effect on the trait) with a lower and well-separated cosine interaction strength score: pair M0P1–M0P2 has a score of 46.02, whereas the pairs Nxxx–Nxxx have an average score of 17.17 (see Supplementary Table S28 for the complete ranking with scores). The NID approach also suffers from the presence of marginal effects. It first detects the correct pair M0P1–M0P2, but in the next pairs we see the presence of an M0Px SNP, suggesting the influence of these SNPs on the interaction score. Even though the scores denote a separation with noninteracting pairs (126.65 versus 54.71), the average score of the pairs involving M0Px SNPs is 65.30, whereas the rest of the pairs have an average of 54.55, suggesting clear influence of marginal effects on NID, even if small in magnitude. Analogous results and observations were obtained for the threshold and epistatic models (see Supplementary Tables S4 and S5). Overall, we can conclude that EpiCID manages to filter out marginal effects and properly retrieve the interacting pair of SNPs. Check Supplementary Tables S24–S41 for the complete rankings with scores for the additive, threshold, and epistatic models.

Purely epistatic model

With the purely epistatic model (Table 2), which is the most complex one, we observe a different behavior.

Both MDR and BOOST are able to correctly identify the epistatic pair, although MDR seems to have a bias given by SNP N8, which repeatedly appears in the interaction ranking and is the top-ranked SNP in the marginal effect ranking. EpiCID also detects the interacting pair, with no marginal effect bias (cosine interaction strength of 40.94 against an average of 16.79 of the noninteracting pairs). NID is also able to capture the epistatic interaction but suffers from marginal effects. Actually, the SNPs M0Px involved in the interaction are reported as interacting with all the rest of the SNPs, retrieving many spurious interactions, showing the sensitivity of NID to even small marginal effects. The complete rankings are available in Supplementary Tables S42–S47.

The analysis on GAMETES data indicates that EpiCID was conceived in the correct way to both filter out marginal effects and correctly detect epistatic pairs. Despite having similar architecture to NID, our method addresses the main issue of NID, which is the sensitivity to marginal effects.

Table 2. Top 20 interactions for the GAMETES purely epistatic model (reporting also marginal effects—ME—for BOOST and MDR)

Rank	BOOST		BOOST ME		MDR		MDR ME		NID		EpiCID	
1	M0P1	M0P2	N8	M0P1	M0P2	N8	M0P1	M0P2	M0P1	M0P2		
2			N151	N8	N171	N151	N217	M0P1	N82	N217		
3			N171	N83	N166	N173	N161	M0P1	N175	N217		
4			N73	N8	N62	N83	N161	M0P2	N54	N175		
5			N33	N8	N98	N239	N70	M0P1	N5	N217		
6			N173	N8	N228	N163	N175	M0P1	N70	N241		
7			N228	N8	N97	N171	N1	M0P1	N187	N217		
8			N83	N8	N249	N201	N235	M0P1	N5	N209		
9			N239	N8	N252	N205	N17	M0P1	N175	N209		
10			N163	N8	N196	N243	N15	M0P1	N17	N46		
11			N36	N8	N129	N57	N194	M0P1	N55	N217		
12			N205	N8	N222	N33	N249	M0P1	N5	N175		
13			N72	N8	N12	N9	N55	M0P1	N161	N217		
14			N243	N8	N188	N204	N70	M0P2	N82	N144		
15			N201	N8	N30	N153	N19	M0P2	N5	N226		
16			N210	N8	N198	N188	N17	M0P2	N55	N175		
17			N57	N188	N253	N28	N68	M0P1	N70	N217		
18			N9	N8	N246	N73	N240	M0P1	N17	N70		
19			N153	N8	N240	N89	N133	M0P1	N99	N217		
20			N204	N8	N151	N58	N172	M0P1	N54	N217		

The names of the SNPs are the ones generated by GAMETES; SNPs named M0P1 and M0P2, marked in yellow for convenience, are the interacting pair, and SNPs named Nxxx are randomly generated with no effect on the trait. We have also marked with green the SNP N8, which appears in multiple SNP pairs, indicating some false signal.

Overall, the analysis on synthetic data showed the effectiveness of our approach in discovering known interactions in different marginal effect conditions, rendering it suitable for usage with real data. Regarding BOOST (and MDR with the exception of the additional interactions in output), we notice that they also perform correctly on synthetic data; yet, as we will see later, they appear to be less effective on real data. This is not surprising, as the GAMETES simulator encodes a single genetic effect in discrete penetrance tables, where each possible genotype combination is assigned a fixed probability of disease. Even the GAMETES additive model is expressed in this way, with risk increasing stepwise across genotype categories rather than smoothly as in real additive effects. This discrete, single-interaction setup makes it easy for MDR to bin genotypes into "high" and "low" risk groups and for BOOST to detect strong pairwise signals, so both methods appear powerful in simulation. In real Biobank data, however, effects are typically weaker, more continuous, and influenced by many variants at once, which breaks this neat alignment.

EpiCID output analysis

A complete evaluation of our results on real-world blood-pressure data involves the entire pipeline of the EpiDetect framework. We begin the analysis by investigating the EpiCID output and detect possible bias in the rankings of the various algorithms for epistasis detection caused by the presence of marginal effects. We can notice this bias by looking at the rankings and examining whether they are dominated by the influence of a small number of SNPs that appear to interact with an exceedingly high number of other SNPs; such a result would be quite unusual in a ranking with nonlinearly interacting pairs [29]. Indeed, epistatic interactions are typically considered to be distributed across multiple loci rather than being dominated by a small subset of highly connected nodes. Whereas certain genes may exhibit higher connectivity, epistatic networks tend to be more distributed. This assumption is also supported by prior studies [41]. By inspecting the degree of the 5 most interacting SNPs in the top-1000 interaction network for SBP (Table 3), we notice how in the BOOST and MDR output only a small number of SNPs are involved in almost all the top 1000 interactions.

Table 3. Distributions of the 5 most interacting SNPs in first 1000 interactions for SBP (the number of interactions corresponds to the degree in the top-1000 interaction network) and the total number of interactions in which they are involved

Rank	SNP	Gene	Interactions
BOOST			
1	rs1012089	AC138627.1	226
2	rs9667596	OR4A44P	180
3	rs7023828	AL358074.1	157
4	rs10224002	PRKAG2	130
5	rs1694068	ARL15	89
Total number of interactions			782
MDR			
1	rs17477177	AC004917.1	263
2	rs17249754	ATP2B1	263
3	rs11191548	CNNM2	237
4	rs1173771	NPR3	163
5	rs17367504	MTHFR	46
Total number of interactions			972
NID			
1	rs77413490	PTEN	163
2	rs1126930	PRKAG1	115
3	rs7331680	CDC16	114
4	rs28621435	GRIN2B	92
5	rs139354822	FARP2	91
Total number of interactions			575
EpiCID			
1	rs77413490	PTEN	150
2	rs1126930	PRKAG1	118
3	rs75961402	HNF4GP1	68
4	rs7331680	CDC16	66
5	rs10437954	ARHGEF25	64
Total number of interactions			466

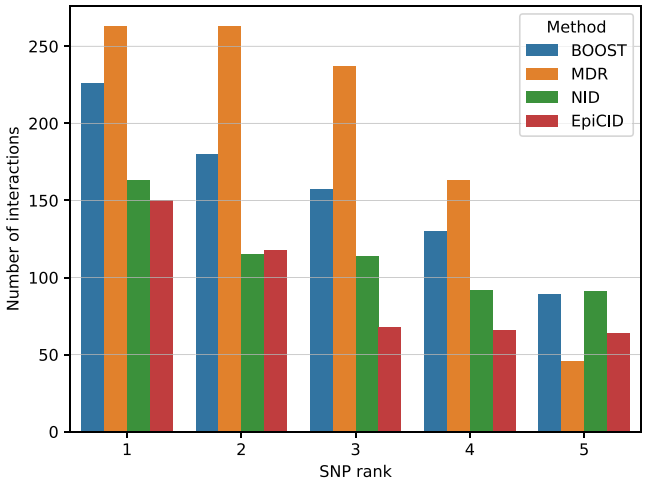

Figure 3. Graphical representation of the distribution of the 5 highest-degree SNPs in the first 1000 interactions for SBP.

For BOOST, 5 SNPs are involved in 782 of the top 1000 interactions. This concentration is much more evident with MDR, with 972 interactions. In contrast, the top 5 SNPs are involved in 575 interactions for NID and even less for EpiCID (466). Thus, neural-network–based methods, especially EpiCID, exhibit more variability in the distribution of the interacting SNPs. This suggests that these methodologies are less affected by the marginal effects of single SNPs, in line with the results obtained using synthetic data. We can observe this behavior also in Fig. 3.

Next, we compared the different algorithms with respect to the distribution of the SNPs with the highest number of interactions at different levels. Figure 4 shows stacked bar charts visualizing the most interacting SNPs in the top-100, top-500, and top-1000 interaction networks. In the first 100 interactions, BOOST is characterized by only one SNP interacting with all the rest of the SNPs, suggesting a presence of bias toward this particular SNP, presumably caused by a strong marginal effect (Fig. 4a). MDR has just 3 SNPs dominating the top 100 interactions, and almost 80 interactions show the presence of the same SNP. On the contrary, for NID and EpiCID, 5 and 6 different SNPs are involved in the top 100 interactions, respectively. We obtained similar insights from the top 500 interactions (Fig. 4b). For BOOST and MDR, we found 4 top-interacting SNPs, 10 for NID and 14 for EpiCID. This is further confirmed by the top-1000 chart (Fig. 4c), which shows 9 top SNPs for BOOST, only 6 for MDR, 14 for NID, and 20 for EpiCID. Given the tendency of standard algorithms to be more strongly affected by marginal effects, they may wrongly detect spurious interactions as pure ones, impinging on the output and resulting in unusual rankings with small numbers of SNPs, also reflected in the subsequent centrality analysis. In particular, it is remarkable how BOOST, despite performing well on synthetic datasets, seems to fail to capture more complex interactions inherently present in real data. This may be attributed to the fact that this methodology struggles when looking for multiple interacting pairs whose correlation with the output follows unknown and complicated nonlinear patterns, as opposed to looking for single pairs with predefined or mathematically generated penetrance functions. On the contrary, neural-network–based models, especially EpiCID, returned a more variable output in the number of interacting SNPs, which is what we would expect in nature, probably providing a more realistic depiction of several interactions that underlie complex traits [42]. Analogous results were obtained for DBP and PP, available as Supplementary data (Supplementary Tables S1 and S2 and Supplementary Figs S1–S4). Supplementary Tables S6–S23 contain the complete rankings with scores for SBP, DBD, and PP for the four algorithms.

We explore more the phenomenon of the complex interactions of real data, by studying the effect of the number of patients on the performance of EpiCID. In particular we study to what extent we are able to recover the interactions discovered with the full dataset, when using an increasing number of patients. Whereas even with a much smaller number of patients (50K), EpiCID is able to detect the real interaction using the GAMETES data, in the case of the UK Biobank data, a small number of patients is able to recover only a small percentage of the top interactions detected by the full dataset, and, not surprisingly, increasing the number of patients leads to more accurate results (Supplementary Fig. S5). Therefore, the number of patients is of crucial importance, and we conjecture that, as in the future we will have larger patient databases, approaches such as EpiCID will return new results. For more details refer to the Supplementary data (Section 4).

Centrality analysis

After training the neural network on the three traits and applying EpiCID to detect interacting SNPs, the top 1000 interactions provided by EpiCID involved 195, 168, and 185 interacting genes, and subsequent centrality analysis resulted in 49, 48, and 53 central genes for the SBP, DBP, and PP traits, respectively. Next, we looked for common genes among the EpiCID-derived network and the networks obtained from BOOST, MDR, and NID (Fig. 5).

Thirty-one out of 49, 33 out of 48, and 37 out of 53 central genes obtained by EpiCID for SBP, DBP, and PP, respectively, were common to the central genes obtained by at least one or more other approaches. EpiCID had the most common genes with NID, compared to other methods, namely 30, 29, and 35 for SBP, DBP, and PP, respectively. MDR contained an intermediate but low number of common central genes with EpiCID and the other strategies: 3 out of 7, 10 out of 33, and 3 out of 5. BOOST had the least common genes to the other three approaches; we found 5, none, and 2 common genes for SBP, DBP, and PP, respectively. EpiCID led to a high percentage of common central genes with methodologies adopting similar architecture (i.e. NID), whereas BOOST and MDR led to more divergent sets of central genes. The central genes identified in this phase are used for the subsequent enrichment analysis.

Enrichment analysis

The final output of our system is provided by the enrichment analysis carried out on the central genes obtained from the epistatic pairs. Centrality analysis led to a gene set that contained only a fraction of the total genes of the epistatic network, which are assumed to play a central role in the trait/disease network under consideration. A critical question is which epistatic network may more accurately represent the affected intracellular pathways that alter organismal physiology and lead to a trait/disease. To support the role of central genes identified through each epistatic network on blood pressure, we performed enrichment analysis using the Enrichr analysis tool [32–34].

In the enrichment analysis, we queried for the traits that are most significantly associated with each of the sets of central genes returned by the four approaches (EpiDetect, NID, MDR, and BOOST), based on data from DisGeNET [36–38], the GWAS Catalog 2023 [39], and UK Biobank [17, 18], as a means to discover which method returns a core network for each trait. So, we examined whether SBP, DBP, and PP traits were found as associated traits and, if yes, in which position in the enrichment ranking; we report the position for the different approaches in Table 4. From the ranking of the Enrichr results, we can observe that the networks obtained by EpiDetect can represent core networks on each associated trait more coherently compared to the other algorithms. With the exception of the PP trait in DisGeNET and UK Biobank, where none of the algorithms predicted the trait, EpiDetect-associated traits were ranked first in all cases. NID had slightly lower accuracy, with DBP in UK Biobank ranking third. Finally, MDR and BOOST had the lowest accuracy among the methods tested, as the relevant traits were ranked in lower positions. This difference can be quantified by the Point Penalty Score that we defined. EpiDetect scores 0, being always the top-performing approach able to retrieve gene networks representative of the analyzed disease.

Having verified the specificity of EpiDetect, we next analyzed for enriched gene ontologies (GO Biological Process databases) found using the central genes of the epistatic networks. We performed a ranking of the pathways found, based

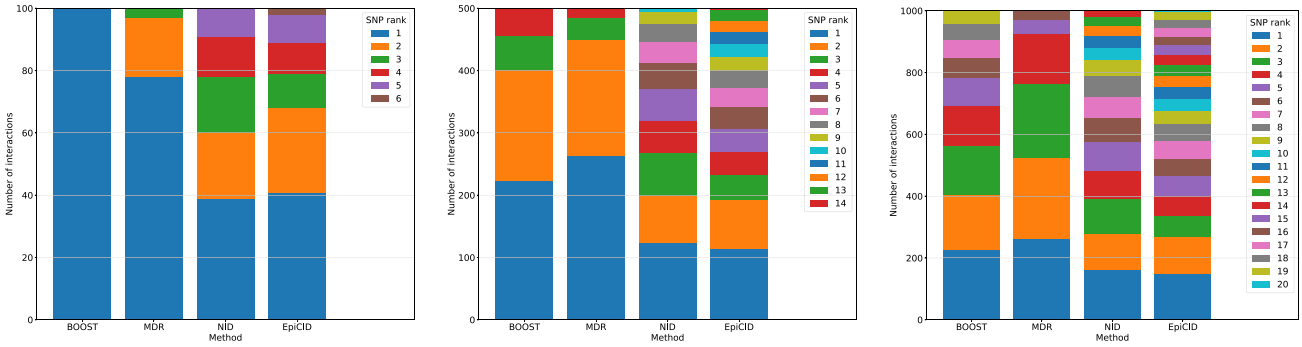

(a) Top 100 interactions. (b) Top 500 interactions. (c) Top 1000 interactions.

Figure 4. Highest-degree SNPs in the top 100 (**a**), top 500 (**b**), and top 1000 (**c**) interactions for SBP. On the *x*-axis, we have the method for detecting interactions, and on the *y*-axis, the number of interactions in which the most popular SNPs are involved (corresponding to the degree in each top-n interaction network).

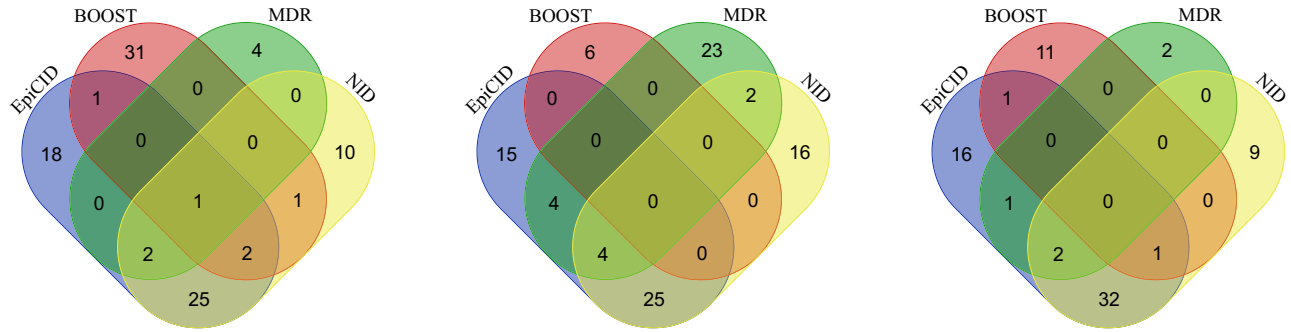

(a) SBP (b) DBP (c) PP

Figure 5. Common central genes from SNPs associated with SBP (**a**), DBP (**b**), and PP (**c**), following analysis with different algorithms.

Table 4. Association of central genes to the respective trait (N/A: not associated among the top 10 terms

Algorithm	Rank									Point Penalty Score
	DisGeNET			GWAS Catalog			UK Biobank			
	SBP	DBP	PP	SBP	DBP	PP	SBP	DBP	PP	
EpiDetect	1	1	N/A	1	1	1	1	1	N/A	0
BOOST	1	1	N/A	1	1	1	1	N/A	N/A	2
MDR	5	1	N/A	3	1	N/A	3	1	N/A	4
NID	1	1	N/A	1	1	1	1	3	N/A	1

We also present the results of our framework when we substitute the first component (EpiCID) with BOOST, MDR, and NID for comparison. The rank indicates the position of the trait in the retrieved ranking of traits/diseases associated with the central genes of each algorithmic approach in each of the three databases. Each method is presented with its point penalty score. The complete enrichment rankings are available in Supplementary Figs S9–S11.

on the obtained *P*-value, following Fisher's exact test; the top results are shown in Fig. 6. The enrichment of the SBP network revealed associations with GOs related to the regulation of synaptic transmission (GO:0050806) and other nervous system-related ontologies (GO:0014020 and GO:0001843, among others), in line with the evidence of its connection with blood pressure regulation [43, 44], as well as established associations between hypertension and wound healing processes (GO:0061045) [45] and plasma membrane abnormalities (GO:0120035) [46]. DBP-central genes were enriched in MAPKs and MAPK signaling (GO:0051403 and GO:0031098) pathways known to be associated with blood-pressure traits [16]. A new finding is that cellular sodium ion homeostasis (GO:0006883) and similar GOs (GO:0055078 and GO:0030004) were found enriched based on EpiDetect epistatic network genes. In the PP network of central genes, among the known associations were that of regulation from RNA polymerase II (GO:0045944) [47] and cardiac muscle fiber development (GO:0048739) [16]. Also, sarcomere organization (GO:0045214) was found to be enriched and in line with previous studies [48]. Among others, an interesting association is that of the Notch signaling pathway (GO:0008593). In summary, EpiDetect-derived epistatic networks provide supporting evidence for known pathways associated with blood pressure traits, as well as the potential for discovering new pathways.

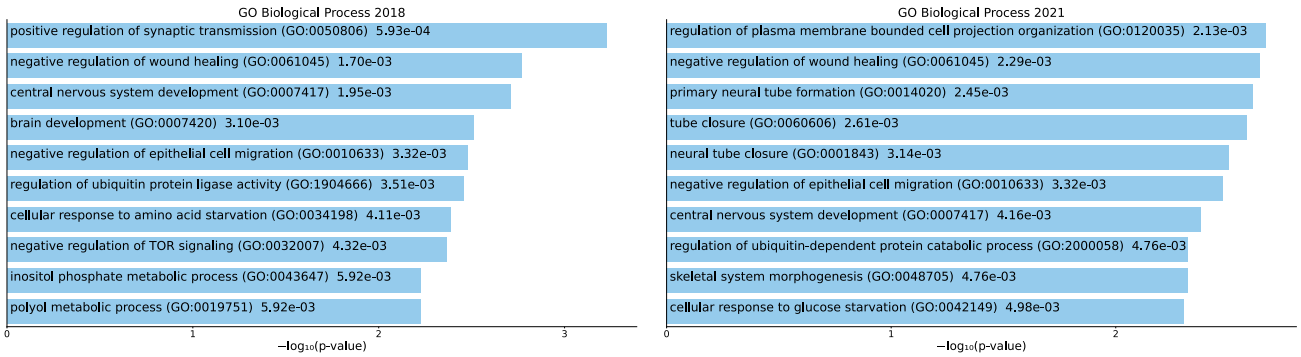

(a) SBP

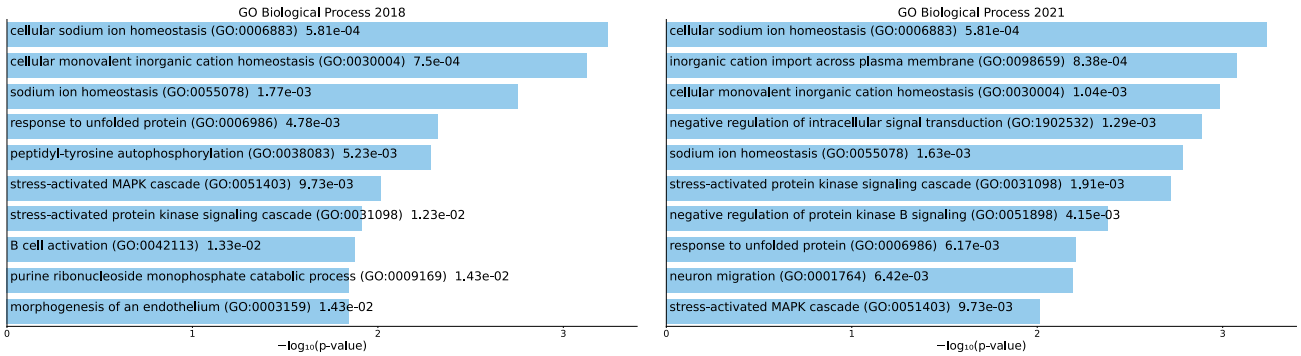

(b) DBP

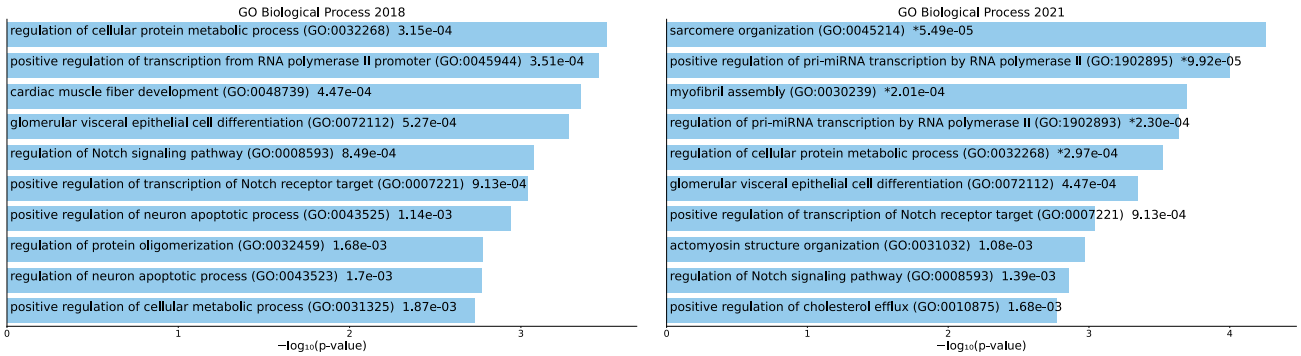

(c) PP

Figure 6. Gene Ontology Biological Process enriched terms (databases 2018 and 2021) in EpiDetect central genes from SBP (**a**), DBP (**b**), and PP (**c**) epistatic networks. Every term is statistically significant and reported with its *P*-value (∗significant adjusted *P*-value).

Discussion

Complex traits are typically influenced by multiple genetic factors, making their analysis challenging. We thereby developed EpiDetect, a novel explainable deep-learning–based method, and we provided a step-by-step protocol to detect epistatic effects and create epistatic gene networks. Our approach detected epistatic interactions, delivering results that are more reasonable than the ones of the other approaches that we compared with (MDR, BOOST, and NID). Clearly, the observations reported on blood pressure are mostly qualitative, as there is no clear way to quantify this apart from comparing the output of the methods and the results of the enrichment analysis.

EpiDetect consists of three major components. First, EpiCID, a novel algorithm that calculates epistatic interactions relying on the inner mechanics of neural networks, properly unveiled by our explainability scheme. Next, a network analysis component is used to extract central genes. Finally, we apply an enrichment analysis process to those central genes. The framework provides a three-level output: ranked interactions of SNPs, central genes, and pathways/ontologies. This multilevel output helps provide a more differentiated view of the results that would be otherwise hard to evaluate given the complexity of the epistasis phenomenon and the lack of a consistent ground truth.

The first component of the framework identifies purely interacting features, attenuating marginal effects of single SNPs, which can lead algorithms to detect spurious interactions; this represents a novelty even in the field of neural-network explainability. Interaction detection approaches usually rely on importance metrics of single features to compute interaction scores [49]; this may leave out interactions for which the main effects are negligible. EpiCID exploits the neural-

network weights learned and optimized during the training process and, by extracting representative feature embeddings, is able to determine how an interaction is created and evolves in the neural network, and how it affects the output. The result of the EpiCID component is a list of ranked epistatic interactions.

The second component is a network analysis workflow, which calculates a gene–gene network based on the top-ranked SNP–SNP interactions previously obtained. This process is in agreement with the concept that a holistic, systems-biology approach is needed to explain gene–gene interactions, which explain epistasis and complex phenotypes [50–53]. It uses centrality analysis to calculate the most significant/central genes in the EpiCID-derived networks [54]. Central genes were identified using the degree measure, which was proven to be effective in the analysis of biological networks. A protein with a high degree interacts with several other proteins, suggesting a possible central regulatory role. There is convincing evidence that proteins with a central regulatory role, or hubs, have been elucidated to be central oncogenes or tumor suppressor genes with clinical significance in cancer-related networks [55–58]. However, our workflow can be perfectly adapted to use other notions of centrality measures [59] (see Supplementary Fig. S6 for additional details). We validated the proposed approach by repeating the workflow using a number of epistatic analysis algorithms, namely MDR, BOOST, and NID. The accuracy of our strategy was assessed by comparing it with the final central gene networks from the other methods. In all three traits (SBP, DBP, and PP), EpiCID had more common genes with NID, whereas MDR and BOOST networks featured more unique genes. EpiCID resulted in a large set of shared central genes with the compared strategies. The fact that EpiCID identified genes that were also found by other methods highlights its robustness. At the same time, the discovery of unique genes found using EpiCID may be more reliable and likely to represent the actual gene networks underlying these traits (confirmed by the subsequent enrichment analysis). This may indicate that it might depict the underlying regulatory network more accurately, rendering the unique genes found more reliable. The common as well as the unique genes found by EpiCID might have a significant role in the actual mechanism of pathogenesis. In that sense, the results we observed with EpiCID suggest it may represent a more relevant depiction of the "biological truth."

The third component of EpiDetect is the enrichment analysis using Enrichr to explore how accurately our approach depicted the core gene regulatory network underlying each trait. Central genes obtained with EpiCID performed similarly to NID, whereas MDR and BOOST lacked accuracy in the association with the respective traits. Interestingly, the degree of common genes between EpiDetect and NID was also informative of their performance after enrichment analysis. Both approaches had the best performances. However, EpiDetect was slightly more accurate in predicting the associated traits. These results also indicate the advantage of neural-network approaches compared to classical statistical methods in predicting epistatic interactions that lead to complex traits.

Further pathway analysis of EpiDetect-derived central genes offered insight into the possible underlying mechanisms of complex traits. GO analysis revealed both known and novel pathways associated with SBP, DBP, and PP. Regarding SBP, we found associations with ontologies related to the nervous system, such as synaptic transmission, brain development, and neural tube formation and closure. This is in line with previous studies with evidence of increased activation of the central nervous system as a contributor to hypertension [44] and effect of brain and nervous system development on blood pressure regulation [43]. Further enriched terms are confirmed by research linking hypertension with delayed wound healing [45] and plasma membrane defects [46]. MAPK signaling was found to be associated with DBP central genes. This result replicates previous studies [16] and provides further evidence of the significance of stress-associated p38-MAPK signaling in hypertension in experimental models of hypertensive mice and rats [60, 61]. A novel finding was that cellular sodium ion homeostasis was found to be associated with DBP. This may explain secondary hypertension in individuals suffering from other pathologies, such as aldosterone-producing adenomas [62] and kidney disease [63]. Regarding PP, regulation from RNA polymerase II [47] and cardiac muscle fiber development [16] are GOs with an established role in PP that were found to be associated with PP EpiDetect-derived central genes. Moreover, we found sarcomere organization to be enriched. Indeed, sarcomere gene mutations are the main genetic cause of hypertrophic cardiomyopathy [64], which can lead to higher PP in patients with an abnormal blood pressure response [48]. Notch signaling pathway was a significant finding that independently replicated findings of altered Notch activity during TNFα-induced hypertension [65], as well as hypoxic pulmonary vasoconstriction [66]. Even though the SNPs analyzed were already known to be associated with the three blood pressure traits, the pathway enrichment analysis aimed at identifying biological processes and molecular mechanisms that may be relevant to these traits, rather than merely reaffirming known associations. All the aforementioned data may provide evidence of the biological utility of EpiDetect-derived epistatic networks in both replicating previous findings and potentially discovering novel significant pathways in complex phenotypes. The effectiveness of EpiDetect might also warrant further investigation of the aforementioned novel pathways.

To maximize the utility of our methodology, we chose to focus on human data, particularly those from the UK Biobank, due to the large scale and the rich diversity of phenotype and genomic data it offers. This focus enables the application of EpiDetect to investigate medically relevant traits, emphasizing its potential to uncover novel epistatic interactions in human health and disease. As a possible direction for future work, applying the method to model organism datasets such as *Caenorhabditis elegans* could offer a complementary benchmark, especially given the availability of well-characterized epistatic interactions [67]. Although such datasets are valuable for method validation, their smaller scale and narrower trait coverage mean that they address a different scope of questions, whereas large human biobanks such as UK Biobank provide the breadth and complexity of data that motivated the design of EpiDetect.

The superior performance of our approach comes with some limitations. First, regarding the scaling with the input data, the approach scales very well with the number of patients, as EpiCID is not affected by the number of input samples (which affects the neural-network training only). However, the computational complexity depends on the number of input SNPs. As the number of SNPs increases, the number of parameters that the network requires to learn increases, and this has an effect on the requirements of computer mem-

ory and running time (a characteristic shared among all deep-learning approaches), as well as the number of pairwise interaction strength computations to be performed. Therefore, this shows the importance of a prior filter on robust genome-wide significant signals to reduce the number of analyzed genetic variants. However, the computation time required for our studies, with the provided datasets described in the "Materials and methods" section, was more than reasonable. As an indicative measure for reference, on a machine with an Intel Core i7-12700H with 4.70 GHz of maximum clock speed, an NVIDIA RTX 3060 GPU with 6 GB of dedicated memory, and 16 GB of RAM, the training of the neural network took on average 30 min. The EpiCID explanation phase required <9 s (as average on SBP, DBP, and PP traits). BOOST and NID also ran in a short time, taking on average around 28 and 41 s, respectively. In contrast, MDR took, on average, 28 min. However, because the application is not performed in a time-sensitive scenario, even longer execution times can be acceptable in the presence of high accuracy.

The second limitation also comes from the fact that our approach is based on neural networks. Despite the fact that in many artificial-intelligence applications neural networks have superior performance (often significant) compared to more classic methods, this superiority comes without guarantees and strong theoretical support. It is current research to understand the theory behind the superiority of neural networks, but the current results are preliminary and often limited to shallow models [68, 69]. Thus, as we mentioned in the introduction, those architectures are treated as black boxes. This also means that the results that our method obtains do not have some theoretical or statistical justification, but they have to be evaluated experimentally; first with approaches such as those used in this paper, and eventually *in vitro*.

To conclude, complex traits are typically influenced by multiple genes and environmental factors, making their analysis challenging. By developing the EpiDetect pipeline and its core element, the EpiCID algorithm, which detects central genes and pathways in an epistatic network, we aim to offer researchers a tool to identify the most influential genetic factors and pathways that contribute to the trait of interest. This information can provide valuable insights into the underlying biology of the trait, inform the development of new treatments, and potentially lead to the discovery of novel drug targets. This approach has the potential to be applied to a wide range of complex traits, including cardiovascular diseases, cancer, and neurological disorders, among others. We provided evidence supporting that our proposed framework may act as a guide to disease-associated experimental research or an independent approach to validate experimental observations.

Acknowledgements

This work was conducted using UK Biobank data under application number 22102.

We dedicate this paper to the memory of our co-author and friend Vangelis Evangelou, who passed away prematurely. We will deeply miss his insight, kindness, and energy.

Prof. Evangelou passed away before the paper was submitted and therefore was not able to approve the final accepted version and cannot be held accountable for the work post-publication.

Author contributions: Conceptualization was done by A.A. and E.E., methodology was developed by A.M., data were gathered and processed by A.M., code was written by A.M., investigation was carried out by A.M., G.M., E.E., and A.A., analysis was done by A.M. and G.M., supervision was done by A.A. and E.E., the original draft was written by A.M., G.M, E.E., and A.A, review and editing were done by A.M., G.M., and A.A.

Supplementary data

Supplementary data is available at NAR Genomics & Bioinformatics online.

Conflict of interest

None declared.

Funding

Supported by the PNRR MUR (Ministero dell'Università e della Ricerca) project "PE0000013-FAIR," the PNRR MUR project IR0000013-SoBigData.it, and the MUR PRIN project 2022EKNE5K "Learning in Markets and Society."

Data availability

The data underlying this article were provided by UK Biobank under application number 22102. Supplementary data are available at NARGAB online. The code produced for this work is available on Zenodo [70] at https://doi.org/10.5281/zenodo.13152209 and on GitHub at https://github.com/AndMastro/EpiDetect.

References

1. Uffelmann E, Huang QQ, Munung NS *et al.* Genome-wide association studies. *Nat Rev Methods Prim* 2021;1:59. https://doi.org/10.1038/s43586-021-00056-9
2. Wang MH, Cordell HJ, Van Steen K. Statistical methods for genome-wide association studies. *Semin Cancer Biol* 2019;55:53–60. https://doi.org/10.1016/j.semcancer.2018.04.008
3. VanderWeele TJ. Epistatic interactions. *Stat Appl Genet Mol Biol* 2010;9. https://doi.org/10.2202/1544-6115.1517
4. Nachtegael C, Gravel B, Dillen A *et al.* Scaling up oligogenic diseases research with OLIDA: the Oligogenic Diseases Database. *Database* 2022;2022:baac023. https://doi.org/10.1093/database/baac023
5. Niel C, Sinoquet C, Dina C *et al.* A survey about methods dedicated to epistasis detection. *Front Genet* 2015;6:285. https://doi.org/10.3389/fgene.2015.00285
6. Hahn LW, Ritchie MD, Moore JH. Multifactor dimensionality reduction software for detecting gene–gene and gene–environment interactions. *Bioinformatics* 2003;19:376–82. https://doi.org/10.1093/bioinformatics/btf869
7. Wan X, Yang C, Yang Q *et al.* BOOST: a fast approach to detecting gene-gene interactions in genome-wide case-control studies. *Am J Hum Genet* 2010;87:325–40. https://doi.org/10.1016/j.ajhg.2010.07.021
8. Yoshida M, Koike A. SNPInterForest: a new method for detecting epistatic interactions. *BMC Bioinformatics* 2011;12:1–10. https://doi.org/10.1186/1471-2105-12-469
9. Jiang R, Tang W, Wu X *et al.* A random forest approach to the detection of epistatic interactions in case-control studies. *BMC Bioinformatics* 2009;10:1–12. https://doi.org/10.1186/1471-2105-10-S1-S65

10. Beam AL, Motsinger-Reif A, Doyle J. Bayesian neural networks for detecting epistasis in genetic association studies. *BMC Bioinformatics* 2014;15:1–12. https://doi.org/10.1186/s12859-014-0368-0
11. Novakovsky G, Dexter N, Libbrecht MW *et al.* Obtaining genetics insights from deep learning via explainable artificial intelligence. *Nat Rev Genet* 2023;24:125–37. https://doi.org/10.1038/s41576-022-00532-2
12. Ross AS, Hughes MC, Doshi-Velez F. Right for the right reasons: Training differentiable models by constraining their explanations. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 2662–70. Melbourne, Australia. 2017. https://doi.org/10.24963/ijcai.2017/371
13. Ribeiro MT, Singh S, Guestrin C. 'Why should i trust you?' Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–44. San Francisco, CA, U.S.A. 2016.
14. Tsang M, Cheng D, Liu Y. Detecting statistical interactions from neural network weights. In: *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. 2018.
15. Bodria F, Giannotti F, Guidotti R *et al.* Benchmarking and survey of explanation methods for black box models. *Data Min Knowl Disc* 2023;37:1719–78.
16. Evangelou E, Warren HR, Mosen-Ansorena D *et al.* Genetic analysis of over 1 million people identifies 535 new loci associated with blood pressure traits. *Nat Genet* 2018;50:1412–25. https://doi.org/10.1038/s41588-018-0205-x
17. Sudlow C, Gallacher J, Allen N *et al.* UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med* 2015;12:e1001779. https://doi.org/10.1371/journal.pmed.1001779
18. Bycroft C, Freeman C, Petkova D *et al.* The UK Biobank resource with deep phenotyping and genomic data. *Nature* 2018;562:203–9. https://doi.org/10.1038/s41586-018-0579-z
19. Haykin S. *Neural networks: a comprehensive foundation*. Englewood Cliffs, New Jersey, U.S.A.: Prentice Hall PTR, 1994.
20. Zou J, Huss M, Abid A *et al.* A primer on deep learning in genomics. *Nat Genet* 2019;51:12–8. https://doi.org/10.1038/s41588-018-0295-5
21. Kingma DP, Ba J. Adam: a method for stochastic optimization. In: Bengio Y, LeCun Y, (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. 2015.
22. Simonyan K, Vedaldi A, Zisserman A. Deep inside convolutional networks: visualising image classification models and saliency maps. In: Bengio Y, LeCun Y, (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*. 2014.
23. Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. In: Bengio Y, LeCun Y, (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. 2015.
24. Pruitt KD, Harrow J, Harte RA *et al.* The consensus coding sequence (CCDS) project: Identifying a common protein-coding gene set for the human and mouse genomes. *Genome Res* 2009;19:1316–23. https://doi.org/10.1101/gr.080531.108
25. Harrow J, Frankish A, Gonzalez JM *et al.* GENCODE: the reference human genome annotation for the ENCODE Project. *Genome Res* 2012;22:1760–74. https://doi.org/10.1101/gr.135350.111
26. Sherry ST, Ward MH, Kholodov M *et al.* dbSNP: the NCBI database of genetic variation. *Nucleic Acids Res* 2001;29:308–11. https://doi.org/10.1093/nar/29.1.308
27. Birney E, Andrews TD, Bevan P *et al.* An overview of Ensembl. *Genome Res* 2004;14:925–8. https://doi.org/10.1101/gr.1860604
28. Consortium EP *et al.* An integrated encyclopedia of DNA elements in the human genome. *Nature* 2012;489:57. https://doi.org/10.1038/nature11247
29. Cordell HJ. Detecting gene–gene interactions that underlie human diseases. *Nat Rev Genet* 2009;10:392–404. https://doi.org/10.1038/nrg2579
30. Guo X, Meng Y, Yu N *et al.* Cloud computing for detecting high-order genome-wide epistatic interaction via dynamic clustering. *BMC Bioinformatics* 2014;15:1–16. https://doi.org/10.1186/1471-2105-15-102
31. Ma'ayan A. Introduction to network analysis in systems biology. *Science signaling* 2011;4:tr5. https://doi.org/10.1126/scisignal.2001965
32. Chen EY, Tan CM, Kou Y *et al.* Enrichr: interactive and collaborative HTML5 gene list enrichment analysis tool. *BMC Bioinformatics* 2013;14:1–14. https://doi.org/10.1186/1471-2105-14-S18-S1
33. Kuleshov MV, Jones MR, Rouillard AD *et al.* Enrichr: a comprehensive gene set enrichment analysis web server 2016 update. *Nucleic Acids Res* 2016;44:W90–7. https://doi.org/10.1093/nar/gkw377
34. Xie Z, Bailey A, Kuleshov MV *et al.* Gene set knowledge discovery with enrichr. *Curr Protoc* 2021;1:e90. https://doi.org/10.1002/cpz1.90
35. Carbon S, Ireland A, Mungall CJ *et al.* AmiGO: online access to ontology and annotation data. *Bioinformatics* 2009;25:288–9. https://doi.org/10.1093/bioinformatics/btn615
36. Piñero J, Queralt-Rosinach N, Bravo A *et al.* DisGeNET: a discovery platform for the dynamical exploration of human diseases and their genes. *Database* 2015;2015. https://doi.org/10.1093/database/bav028
37. Piñero J, Bravo À, Queralt-Rosinach N *et al.* DisGeNET: a comprehensive platform integrating information on human disease-associated genes and variants. *Nucleic Acids Res* 2016; D833–39.
38. Piñero J, Ramírez-Anguita JM, Saüch-Pitarch J *et al.* The DisGeNET knowledge platform for disease genomics: 2019 update. *Nucleic Acids Res* 2020;48:D845–55.
39. Sollis E, Mosaku A, Abid A *et al.* The NHGRI-EBI GWAS Catalog: knowledgebase and deposition resource. *Nucleic Acids Res* 2023;51:D977–85. https://doi.org/10.1093/nar/gkac1010
40. Urbanowicz RJ, Kiralis J, Sinnott-Armstrong NA *et al.* GAMETES: a fast, direct algorithm for generating pure, strict, epistatic models with random architectures. *BioData Min* 2012;5:1–14. https://doi.org/10.1186/1756-0381-5-16
41. McKinney B, Pajewski NM. Six degrees of epistasis: statistical network models for GWAS. *Front Genet* 2012;2:109. https://doi.org/10.3389/fgene.2011.00109
42. Taylor MB, Ehrenreich IM. Higher-order genetic interactions and their contribution to complex traits. *Trends Genet* 2015;31:34–40. https://doi.org/10.1016/j.tig.2014.09.001
43. Dono R, Texido G, Dussel R *et al.* Impaired cerebral cortex development and blood pressure regulation in FGF-2-deficient mice. *EMBO J* 1998;17:4213–25. https://doi.org/10.1093/emboj/17.15.4213
44. Carmichael CY, Wainford RD. Hypothalamic signaling mechanisms in hypertension. *Curr Hypertension Rep* 2015;17:1–9.
45. Ahmed AA, Mooar PA, Kleiner M *et al.* Hypertensive patients show delayed wound healing following total hip arthroplasty. *PLoS One* 2011;6:e23224. https://doi.org/10.1371/journal.pone.0023224
46. Rinaldi G, Bohr D. Plasma membrane and its abnormalities in hypertension. *Am J Med Sci* 1988;295:389–95. https://doi.org/10.1097/00000441-198804000-00028
47. de Leeuw CA, Stringer S, Dekkers IA *et al.* Conditional and interaction gene-set analysis reveals novel functional pathways for blood pressure. *Nat Commun* 2018;9:3768. https://doi.org/10.1038/s41467-018-06022-6

48. Heffernan KS, Maron MS, Patvardhan EA *et al.* Relation of pulse pressure to blood pressure response to exercise in patients with hypertrophic cardiomyopathy. *Am J Cardiol* 2011;107:600–3. https://doi.org/10.1016/j.amjcard.2010.10.023
49. Greenside P, Shimko T, Fordyce P *et al.* Discovering epistatic feature interactions from neural network models of regulatory DNA sequences. *Bioinformatics* 2018;34:i629–37. https://doi.org/10.1093/bioinformatics/bty575
50. Moore JH, Williams SM. Traversing the conceptual divide between biological and statistical epistasis: systems biology and a more modern synthesis. *Bioessays* 2005;27:637–46. https://doi.org/10.1002/bies.20236
51. Gjuvsland AB, Hayes BJ, Omholt SW *et al.* Statistical epistasis is a generic feature of gene regulatory networks. *Genetics* 2007;175:411–20. https://doi.org/10.1534/genetics.106.058859
52. Aylor DL, Zeng ZB. From classical genetics to quantitative genetics to systems biology: modeling epistasis. *PLoS Genet* 2008;4:e1000029. https://doi.org/10.1371/journal.pgen.1000029
53. Phillips PC. Epistasis—the essential role of gene interactions in the structure and evolution of genetic systems. *Nat Rev Genet* 2008;9:855–67. https://doi.org/10.1038/nrg2452
54. Koschützki D, Schreiber F. Centrality analysis methods for biological networks and their application to gene regulatory networks. *Gene Regul Syst Biol* 2008;2:GRSB–S702.
55. Lin CY, Chin CH, Wu HH *et al.* Hubba: hub objects analyzer—a framework of interactome hubs identification for network biology. *Nucleic Acids Res* 2008;36:W438–43. https://doi.org/10.1093/nar/gkn257
56. Wang N, Xu Hl, Zhao X *et al.* Network-based identification of novel connections among apoptotic signaling pathways in cancer. *Appl Biochem Biotechnol* 2012;167:621–31. https://doi.org/10.1007/s12010-012-9704-x
57. Deng SP, Zhu L, Huang DS. Predicting hub genes associated with cervical cancer through gene co-expression networks. *IEEE/ACM T Comput Biol Bioinform* 2015;13:27–35. https://doi.org/10.1109/TCBB.2015.2476790
58. Tang J, Kong D, Cui Q *et al.* Prognostic genes of breast cancer identified by gene co-expression network analysis. *Front Oncol* 2018;8:374. https://doi.org/10.3389/fonc.2018.00374
59. Shannon P, Markiel A, Ozier O *et al.* Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Res* 2003;13:2498–504. https://doi.org/10.1101/gr.1239303
60. Callera GE, Montezano AC, Yogi A *et al.* c-Src–dependent nongenomic signaling responses to aldosterone are increased in vascular myocytes from spontaneously hypertensive rats. *Hypertension* 2005;46:1032–8. https://doi.org/10.1161/01.HYP.0000176588.51027.35
61. Bao W, Behm DJ, Nerurkar SS *et al.* Effects of p38 MAPK Inhibitor on angiotensin II-dependent hypertension, organ damage, and superoxide anion production. *J Cardiovasc Pharm* 2007;49:362–8. https://doi.org/10.1097/FJC.0b013e318046f34a
62. Beuschlein F, Boulkroun S, Osswald A *et al.* Somatic mutations in ATP1A1 and ATP2B3 lead to aldosterone-producing adenomas and secondary hypertension. *Nat Genet* 2013;45:440–4. https://doi.org/10.1038/ng.2550
63. Schiffl H, Küchle C, Lang S. Dietary salt, intracellular ion homeostasis and hypertension secondary to early-stage kidney disease. *Miner Electrol Metab* 1996;22:178–81.
64. Sabater-Molina M, Pérez-Sánchez I, Hernández Del Rincón J *et al.* Genetics of hypertrophic cardiomyopathy: A review of current state. *Clin Genet* 2018;93:3–14. https://doi.org/10.1111/cge.13027
65. Hurst LA, Dunmore BJ, Long L *et al.* TNFα drives pulmonary arterial hypertension by suppressing the BMP type-II receptor and altering NOTCH signalling. *Nat Commun* 2017;8:14079. https://doi.org/10.1038/ncomms14079
66. Smith KA, Voiriot G, Tang H *et al.* Notch activation of Ca2+ signaling in the development of hypoxic pulmonary vasoconstriction and pulmonary hypertension. *Am J Resp Cell Mol Biol* 2015;53:355–67. https://doi.org/10.1165/rcmb.2014-0235OC
67. Toch K, Buczek M, Labocha MK. Genetic interactions in various environmental conditions in *Caenorhabditis elegans*. *Genes* 2023;14:2080. https://doi.org/10.3390/genes14112080
68. Goldt S, Loureiro B, Reeves G *et al.* The gaussian equivalence of generative models for learning with shallow neural networks. In: *Proceedings of Mathematical and Scientific Machine Learning (PMLR)*. Vol. 2022, 145, 426–71.
69. Veiga R, Stephan L, Loureiro B *et al.* Phase diagram of Stochastic Gradient Descent in high-dimensional two-layer neural networks. *Adv Neur Inf Proc Syst* 2022;35:23244–55.
70. Mastropietro A. EpiDetect: v1.0.0. Zenodo, 2024. https://doi.org/10.5281/zenodo.13152209

Received: August 23, 2024. **Revised:** December 15, 2025. **Accepted:** December 22, 2025