Nucleic Acids Research, 2024, **52**, 10144–10160
https://doi.org/10.1093/nar/gkae697
Advance access publication date: 23 August 2024
Computational Biology

Network medicine-based epistasis detection in complex diseases: ready for quantum computing

Markus Hoffmann[1,2,3,*,†], Julian M. Poschenrieder[1,4,†], Massimiliano Incudini[5,†], Sylvie Baier[1,†], Amelie Fritz[6,7,†], Andreas Maier[4,†], Michael Hartung[4,†], Christian Hoffmann[4,†], Nico Trummer[1], Klaudia Adamowicz[4], Mario Picciani[8], Evelyn Scheibling[1], Maximilian V. Harl[1,9], Ingmar Lesch[1], Hunor Frey[1], Simon Kayser[1], Paul Wissenberg[1], Leon Schwartz[1], Leon Hafner[1,2], Aakriti Acharya[10,11], Lena Hackl[4], Gordon Grabert[10,11], Sung-Gwon Lee[3,12], Gyuhyeok Cho[13], Matthew E. Cloward[14], Jakub Jankowski[3], Hye Kyung Lee[3], Olga Tsoy[4], Nina Wenke[4], Anders Gorm Pedersen[6], Klaus Bønnelykke[7], Antonio Mandarino[15], Federico Melograna[16,17], Laura Schulz[18], Héctor Climente-González[19], Mathias Wilhelm[8,20], Luigi Iapichino[18], Lars Wienbrandt[21], David Ellinghaus[21], Kristel Van Steen[16,17], Michele Grossi[22], Priscilla A. Furth[23], Lothar Hennighausen[2,3], Alessandra Di Pierro[5,‡], Jan Baumbach[4,24,‡], Tim Kacprowski[9,10,‡], Markus List[1,25,*,‡] and David B. Blumenthal[,*,‡]

[1]Data Science in Systems Biology, School of Life Sciences, Technical University of Munich, Freising, Germany
[2]Institute for Advanced Study (Lichtenbergstrasse 2 a) Technical University of Munich, D-85748 Garching, Germany
[3]National Institute of Diabetes and Digestive and Kidney Diseases, Bethesda, MD 20892, USA
[4]Institute for Computational Systems Biology, University of Hamburg, Germany
[5]Dipartimento di Informatica, Universit'a di Verona, Strada le Grazie 15 - 34137 Verona, Italy
[6]Department of Health Technology, Section for Bioinformatics, Technical University of Denmark, DTU, 2800 Kgs. Lyngby, Denmark
[7]Copenhagen Prospective Studies on Asthma in Childhood (COPSAC), Herlev and Gentofte Hospital, University of Copenhagen, Copenhagen, Denmark
[8]Computational Mass Spectrometry, Technical University of Munich, Freising, Germany
[9]Department of Health Sciences and Technology, Neuroscience Center Zürich (ZNZ), Swiss Federal Institute of Technology (ETH Zürich), Zürich 8092, Switzerland
[10]Division Data Science in Biomedicine, Peter L. Reichertz Institute for Medical Informatics, Technische Universität Braunschweig and Hannover Medical School, Rebenring 56, 38106 Braunschweig, Germany
[11]Braunschweig Integrated Centre of Systems Biology (BRICS), Technische Universität Braunschweig, Rebenring 56, 38106 Braunschweig, Germany
[12]School of Biological Sciences and Technology, Chonnam National University, Gwangju, Korea
[13]Department of Chemistry, Gwangju Institute of Science and Technology, Gwangju, Korea
[14]Department of Biology, Brigham Young University, Provo, UT, USA
[15]International Centre for Theory of Quantum Technologies, University of Gdańsk, 80-309 Gdańsk, Poland
[16]BIO3 - Systems Genetics; GIGA-R Medical Genomics, University of Liège, Liège, Belgium
[17]BIO3 - Systems Medicine; Department of Human Genetics, KU Leuven, Leuven, Belgium
[18]Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities (LRZ), Garching b. München, Germany
[19]RIKEN Center for Advanced Intelligence Project, Tokyo, Japan
[20]Munich Data Science Institute (MDSI), Technical University of Munich, Garching, Germany
[21]Institute of Clinical Molecular Biology, Christian Albrechts University of Kiel, Kiel, Germany
[22]European Organization for Nuclear Research (CERN), Geneva1211, Switzerland
[23]Departments of Oncology & Medicine, Georgetown University, Washington, DC, USA
[24]Computational BioMedicine Lab, University of Southern Denmark, Denmark
[25]Biomedical Network Science Lab, Department Artificial Intelligence in Biomedical Engineering, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

[*]To whom correspondence should be addressed. Email: markus.hoffmann@nih.gov
Correspondence may also be addressed to Markus List. Email: markus.list@tum.de

Received: November 22, 2023. **Revised:** July 12, 2024. **Editorial Decision:** July 15, 2024. **Accepted:** August 1, 2024

Correspondence may also be addressed to David B. Blumenthal. Email: david.b.blumenthal@fau.de
†The first eight authors should be regarded as Joint First Authors.
‡The last five authors should be regarded as Joint Last Authors.

Abstract

Most heritable diseases are polygenic. To comprehend the underlying genetic architecture, it is crucial to discover the clinically relevant epistatic interactions (EIs) between genomic single nucleotide polymorphisms (SNPs) (1–3). Existing statistical computational methods for EI detection are mostly limited to pairs of SNPs due to the combinatorial explosion of higher-order EIs. With NeEDL (**ne**twork-based **e**pistasis **d**etection via **l**ocal search), we leverage network medicine to inform the selection of EIs that are an order of magnitude more statistically significant compared to existing tools and consist, on average, of five SNPs. We further show that this computationally demanding task can be substantially accelerated once quantum computing hardware becomes available. We apply NeEDL to eight different diseases and discover genes (affected by EIs of SNPs) that are partly known to affect the disease, additionally, these results are reproducible across independent cohorts. EIs for these eight diseases can be interactively explored in the Epistasis Disease Atlas (https://epistasis-disease-atlas.com). In summary, NeEDL demonstrates the potential of seamlessly integrated quantum computing techniques to accelerate biomedical research. Our network medicine approach detects higher-order EIs with unprecedented statistical and biological evidence, yielding unique insights into polygenic diseases and providing a basis for the development of improved risk scores and combination therapies.

Graphical abstract

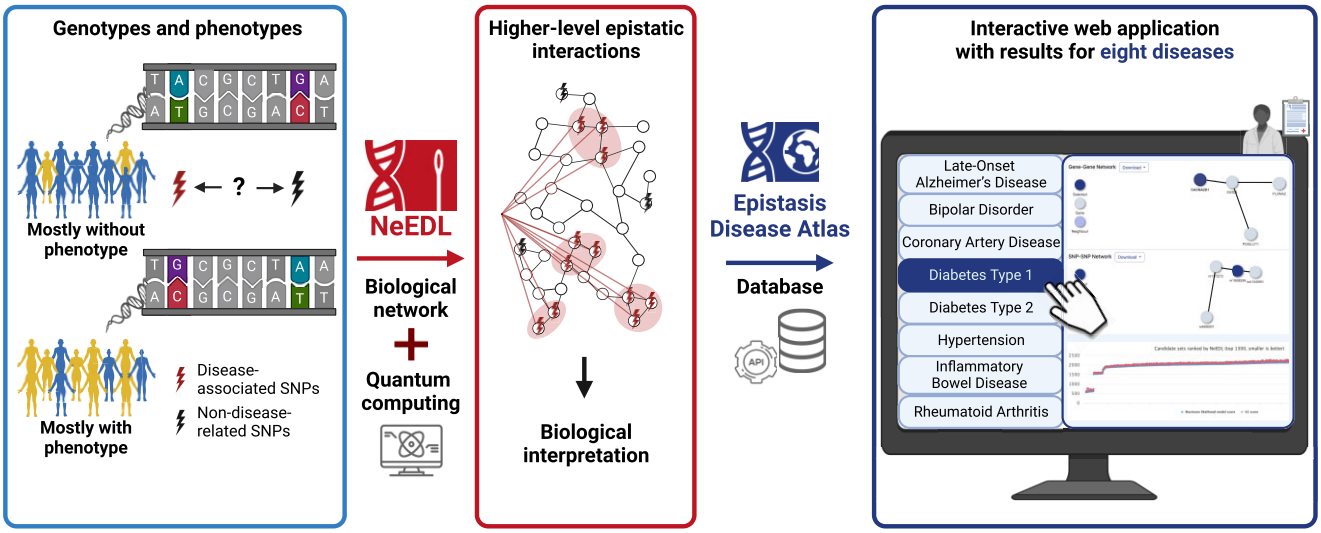

Introduction

Genome-wide association studies (GWAS) aim to identify genetic single nucleotide polymorphisms (SNPs) that are individually associated with a phenotype such as a disease (1–3). Thousands of individual SNPs have been associated with diseases since the early 1990s, yet they account only for a fraction of the investigated traits' heritability (4,5). The hypothesis is that a significant proportion of the missing heritability can be explained by epistatic SNP interactions that are jointly but not individually associated with the phenotype (6). However, no undisputed cases of epistasis in humans are known. Therefore, scalable detection tools for epistatic interactions that yield interpretable, high-quality results are needed to advance the understanding of possible genetic causes of diseases.

Detecting biologically plausible epistatic candidate SNP sets is difficult: Firstly, there is no consensus on the choice of the formal model in order to make this problem algorithmically accessible. Secondly, it is often unclear if predicted cases of epistasis are statistical artifacts since they are hardly interpretable. Thirdly, comprehensively testing higher-order interactions is computationally intractable due to the combinatorial explosion of the search space. Existing epistasis detection tools, including Potpourri (9), LinDen (10), PoCos (11), MACOED (12) and BiologicalEpistasis (13) hence do not scale to large datasets and are mostly restricted to pairwise interactions (Supplementary Table S1). Finally, easily accessible resources to browse and explore pre-computed epistasis candidates interactively are lacking. To overcome these challenges, we present NeEDL (**ne**twork-based **e**pistasis **d**etection via **l**ocal search), an epistasis detection tool leveraging quantum computing and network medicine to identify biologically interpretable candidate sets of epistatic interaction between SNPs (see Figure 1). NeEDL unifies various GWAS input formats and initially filters datasets for relevant SNPs (Supplementary Figure S1A). NeEDL then offers different statistical epistasis models (8) for associating higher-order SNP sets with phenotypes, prioritizing SNP sets using a biologically-informed SNP-SNP interaction (SSI) network. Our hypothesis is that SNPs affecting the same or functionally related proteins are more likely to be involved in relevant epistatic interaction and less likely to be statistical artifacts. Thus, SNPs are mapped to genes using dbSNP (7) and connected if they affect the same protein or a functionally associated protein (e.g., in a protein-protein interaction (PPI) network, Figure 1A). Within the SSI network, NeEDL uses local search with multi-start and simulated annealing to find connected subgraphs of a user-specified size that are

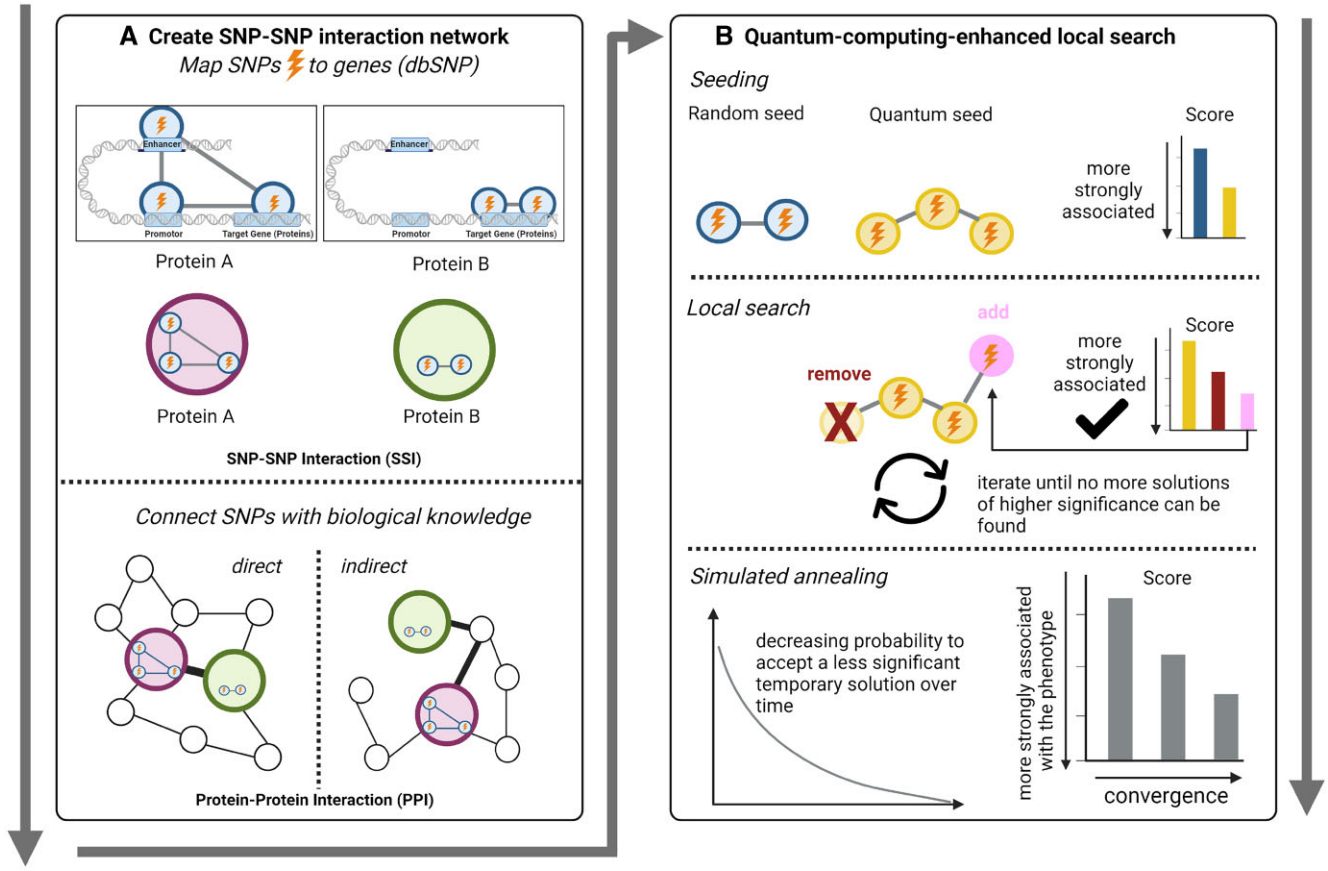

Figure 1. Overview of the methodology. (**A**) SNPs are mapped to genes via SNP-gene links obtained from dbSNP (7) and are then connected in an SNP-SNP interaction (SSI) network if they are mapped to the same protein or to adjacent proteins in a protein-protein interaction (PPI) network (i.e., directly when the proteins interact with each other or indirectly connected when the SNPs are linked through other proteins). (**B**) Next, NeEDL picks either random seeds or seeds selected by a quantum-computing optimization algorithm. It then uses local search (i.e., either adding new SNPs from the direct neighborhood in the SSI network or removing SNPs that worsen the score) to find SNP sets of a user-specified maximum size that induce connected subgraphs in the SSI network and are locally optimal w. r. t. statistical association of the higher-order genotype with the investigated phenotypic trait. To quantify association strength, NeEDL implements various statistical epistasis models suggested in the literature (8). Since local search can get stuck due to the requirement of constant improvement in each step, we use simulated annealing, which allows us to accept a less significant intermediate solution with decreasing probability over time.

locally optimal w. r. t. the selected statistical epistasis model (Figure 1B). Focusing on SNP sets inducing connected subgraphs in the SSI network not only increases the likelihood of uncovering biologically meaningful cases of epistasis but also dramatically reduces the size of the search space. Finally, NeEDL fully integrates quantum computing (QC) algorithms to yield high-quality initial solutions for the local search to reduce NeEDL's runtime further once QC hardware becomes more readily available. Our work hence adds epistasis detection to the growing array of bioinformatics problems that can be tackled using QC, where prior studies have focused on applications such as molecular docking (14), genome assembly (15), DNA assembly (16), DNA reconstruction (17) and drug design (18).

We benchmark NeEDL against frequently used epistasis detection tools on GWAS data for eight human diseases — Late-Onset Alzheimer's Disease (LOAD), Bipolar Disorder (BD), Coronary Artery Disease (CAD), Diabetes Type 1 (T1D), Diabetes Type 2 (T2D), Hypertension (HT), Inflammatory Bowel Disease (IBD) and Rheumatoid Arthritis (RA) (see Supplementary Tables S2–S10 for an overview over sample numbers and SNP numbers and Datasets for information on datasets). Our validation shows that NeEDL markedly outperforms the currently used epistasis detection tools in terms of the discovered candidate SNP sets' associations with the diseases (i.e. the statistical score). The top-scoring candidate SNPs sets are further supported by (i) significant biological associations with the respective diseases, and (ii) in the case of LOAD, by replication in independent cohorts (i.e., discovery with NeEDL in the TGen cohort and replication in the UK Biobank cohort). We make the results of NeEDL for the eight diseases, which we obtained in over two million CPU hours, widely accessible through the Epistasis Disease Atlas (https://epistasis-disease-atlas.com, Supplementary Figure S1B) via an application programming interface (API), an R and Python package, an FTP server, and an interactive, feature-rich web application.

Materials and methods

Implementation details

NeEDL is implemented in C++. It employs the Boost Graph Library version 1.71.0 (19) and iGraph version 0.9.8 (20) for the construction and handling of graphs. We use OpenMP (21) for the parallelization of the initialization of NeEDL (i.e. read

the input data, map the SNPs to genes, construct the SSI network, and get random start seeds) and the local search with simulated annealing. We further included the following external dependencies: CMake v. 2.6 or higher, Doxygen, Catch version 2.11.0, CLI11 version 1.9.0 and Eigen version 3.3.7 (http://eigen.tuxfamily.org). We provide an installer script that installs most external dependencies (excluding CMake, Doxygen, and OpenMP). The dockerized version of NeEDL can be directly pulled and executed from: https://hub.docker.com/r/bigdatainbiomedicine/needl.

We conducted computation on the high-performance computing systems provided by the Leibniz Supercomputing Center of the Bavarian Academy of Sciences and Humanities (LRZ). Each MACOED task was run using a single core (2.6 GHz nominal frequency) on LRZ's Large Memory Teramem cluster (HP DL580 shared memory system) since MACOED has extremely high memory requirements (approx. 1.3 terabytes of RAM). For each NeEDL task, we used 28 threads on 14 physical cores (2.6 GHz nominal frequency) on LRZ's CoolMUC-2 cluster (28-way Haswell-EP nodes).

Since MACOED uses a randomized optimization technique (ant colony optimization), we ran it 100 times on each dataset to minimize the impact of random bias. All obtained SNP sets were used for downstream evaluation. Also, NeEDL includes a randomized subroutine, namely, the seeding of the initial solutions for the local search. To account for this, we ran NeEDL's local search with multi-start and introduced a global time limit of 12 h.

The behavior of NeEDL can be influenced through various parameters, including (i) seeding procedure (default: random seeding), (ii) maximal size of SNPs in a candidate SNP set (default: 10), (iii) maximum iteration of lookups inside one local neighborhood (default: 300), (iv) statistical model used for the local search (default: maximum likelihood model) and (v) SNP filters inside the network (e.g. minor allele frequency, maximum marginal association, linkage disequilibrium cutoff, default: None). For further parameters, please check the GitHub repository and the manual of NeEDL.

Data format converter, data cleaning and data filtering

We faced challenges in using different genotype and phenotype data formats with NeEDL and other epistasis detection tools due to the lack of available data processing tools for converting these formats to the JSON format required by NeEDL. Furthermore, most epistasis detection tools require formatted, pre-cleaned, and pre-filtered datasets as input. To automate preprocessing, cleaning, filtering, and conversion into a joint machine-readable format, we developed epiJSON which supports VCF, PED/MAP, TPED/TFAM, BED/BIM/FAM and other formats and follows recommendations laid out by Marees *et al.* (22): (i) excluding samples with missing phenotypes; (ii) excluding SNPs with more than 20% of missing data across individuals and individuals with more than 20% of missing data across SNPs; (iii) checking an individual's assigned sex using X chromosome data, and removing those with a discrepancy; (iv) considering the number of samples within the given dataset to determine a suitable minor allele frequency threshold; (v) removing variants that fail the Hardy–Weinberg test at a threshold recommended for binary traits of $1e^{-10}$ in controls and $1e^{-6}$ in cases; (vi) excluding individuals with heterozygosity that is too high or too low, indicating sample contamination or inbreeding and (vii) excluding SNPs with any missing data across individuals to ensure that the dataset has no missing values. Cleaning processes and filtering steps can be adjusted with parameters following the guidelines in the GitHub repository.

Datasets

Eight datasets were included in our benchmark: LOAD, BD, CAD, T1D, T2D, HT, IBD and RA (see Section Data Availability for links to the databases). The LOAD dataset had controls included. All other datasets had no controls included, and we used the British Cohort 1958 provided by the Wellcome Trust Case Control Consortium (WTCCC) as controls. In Supplementary Tables S2–S10, we show the detailed sample and SNP numbers before and after the epiJSON tool for the datasets discussed in this manuscript. The LOAD dataset was provided by the TGen consortium (23,24), the other datasets were provided by the WTCCC consortium (25) (see Data Availability).

A replication study for SNP sets discovered in the disease LOAD was conducted using data from the UK Biobank database (www.ukbiobank.ac.uk) (26). Individuals were selected using ICD-10 coding (27) (field 41 270), see Supplementary Table S11. We constructed two replication data sets. One containing individuals with mixed ancestry and a sub-data set containing only individuals with British ancestry. Individuals in controls matched the individuals in cases in age (field 34) and gender (field 31) proportionally. Replication was performed on the results of the discovery study using the tools NeEDL, higher-order baseline, second-order baseline, MACOED, and LinDen. The Pearson correlation coefficient was calculated (28) to determine the correlation of statistical scores between the discovery and replication dataset for the candidate SNP sets for LOAD of the NeEDL.

Epistasis models

For statistical modeling, we use the framework introduced by Blumenthal *et al.* (8). Let $\mathbf{G} = (g_{i,s}) \in \{0, 1, 2\}^{\mathcal{I}\times\mathcal{S}}$ be a genotype matrix, where $\mathcal{I}$ and $\mathcal{S}$ denote the sets of all individuals (patients) and SNPs, respectively, and the entry $g_{i,s}$ encodes the number of minor alleles of individual i at SNP s. Moreover, let $\mathbf{y} \in \mathbb{P}^{\mathcal{I}}$ be a phenotype vector, where $\mathbb{P}$ is the set of all possible phenotypes (e.g., $\mathbb{P} = \{0, 1\}$ for case-versus-control data). We call a tuple $M = (f_{\mathrm{G,y}}, \sigma)$ an *epistasis model* if $\sigma \in \{MIN, MAX\}$ is the model sense and $f_{\mathrm{G,y}} : \mathcal{P}(\mathcal{S}) \to \mathbb{R}$ is an objective function that assigns objective values $f_{\mathrm{G,y}}(S)$ to SNP sets $S \subseteq \mathcal{S}$ which quantify the statistical evidence for the SNPs $s \in S$ being involved in epistatic interaction.

A wide variety of epistasis models have been proposed in the literature, including Bayesian network scores (such as the K2-score (29–34)), variants of multifactor dimensionality reduction (35–39), variants of polynomial regression (40–42), the P-value of the χ^2-test, which is arguably the most widely used epistasis model (29–34), or the maximum likelihood model (MLM) introduced by Blumenthal *et al.* (8).

Equipped with these definitions (and assuming $\sigma = MIN$; if $\sigma = MAX$ we instead have to maximize in (Equation (1)), the problem of detecting SNPs involved in epistatic interaction can, in a first attempt, be formalized as the problem of finding an SNP set $S^\star$ that solves the optimization problem

$$S^\star \in \arg\min\{f_{\mathrm{G,y}}(S) \mid S \subseteq \mathcal{S} \wedge LB \leq |S| \leq UB\}, \qquad (1)$$

where $LB \geq 2$ and $UB \geq LB$ are user-specified lower and upper bounds on the size of the solution SNP set. The problem with this naïve formulation is that the search space is huge, and solving the optimization problem is hence computationally very expensive. In NeEDL, we mitigate this shortcoming by restricting the search space to SNP sets S, which induce a connected subgraph $\mathcal{N}[S]$ in an SSI network $\mathcal{N} = (\mathcal{S}, \mathcal{E})$ where two SNPs $s_1, s_2 \in \mathcal{S}$ are connected by an edge $s_1 s_2 \in \mathcal{E}$ if there exists prior evidence that they might be involved in a biologically meaningful interaction (see the following subsection for details on the construction of $\mathcal{N}$). That is, NeEDL uses the following computational model:

$$S^\star \in \arg\min\{ f_{\mathrm{G,y}}(S) \mid S \subseteq \mathcal{S} \wedge LB \leq |S| \leq UB \wedge \mathcal{N}[S] \text{ is connected}\} \quad (2)$$

For our benchmark, we use four epistasis models: the P-value of the χ^2-test, the K2-score, the MLM score, and the NLL gain of a quadratic regression model in comparison to a linear regression model. The former three models are all based on the penetrance table $T_{\mathrm{G,y},S}(\mathbf{g}) = \{\{y_i \mid i \in \mathcal{I} : g_{i,s} = g_s \forall s \in S\}\}$ of the scored SNP set S. For each possible genotype $\mathbf{g} \in \{0, 1, 2\}^S$ at the scored SNP set S, the cell $T_{\mathrm{G,y},S}(\mathbf{g})$ of the penetrance table contains the multiset of phenotypes of individuals whose genotype at S matches $\mathbf{g}$. For all three models, the score $f_{\mathrm{G,y}}(S)$ is small if the phenotypes are unevenly distributed across the cells of $T_{\mathrm{G,y},S}$ (see Blumenthal *et al.* (8) for details). For the P-value of the χ^2-test, the K2-score, and the MLM score, we hence have $\sigma = MIN$ and small scores $f_{\mathrm{G,y}}(S)$ indicate that the genotype at S is predictive of phenotypic variation.

The NLL gain of quadratic versus linear regression captures another dimension of the concept of epistasis. Here, we first fit two linear or logistic (depending on whether phenotypes are quantitative are categorical) regression models as follows ($\mathbf{G}_{\bullet,s}$ is the column for the SNP s in $\mathbf{G}$ and $\odot$ denotes element-wise multiplication):

$$\mathbf{y} \sim \theta_0 + \sum_{s \in S} \theta_s \mathbf{G}_{\bullet,s} + \sum_{\{s,s'\} \in \binom{S}{2}} \theta_{s,s'} \mathbf{G}_{\bullet,s} \odot \mathbf{G}_{\bullet,s'} \quad (3)$$

$$\mathbf{y} \sim \theta'_0 + \sum_{s \in S} \theta'_s \mathbf{G}_{\bullet,s} \quad (4)$$

Subsequently, we compute NLLs for the fit models $\boldsymbol{\theta}$ and $\boldsymbol{\theta'}$ and define our score as $f_{\mathrm{G,y}}(S) = \mathrm{NLL}(\boldsymbol{\theta'}) - \mathrm{NLL}(\boldsymbol{\theta})$. For the NLL gain, large scores hence indicate that we can better predict the phenotypes when considering multiplicative interactions between all pairs of SNPs contained in the scored SNP set S than when considering only marginal effects (i.e. we have $\sigma = MAX$). Since the NLL gain can result in negative values, we excluded such values from the analysis.

Construction of the SNP–SNP interaction network

To construct $\mathcal{N} = (\mathcal{S}, \mathcal{E})$, we require a many-to-many SNP-to-gene mapping $\pi \subseteq (\mathcal{S} \times \mathcal{G})$ ($\mathcal{G}$ is the set of all genes), as well as a gene-gene network $\mathcal{N}' = (\mathcal{G}, \mathcal{E}')$ (Figure 1A). In NeEDL, we obtain π from dbSNP (7) and use the BioGRID (43) PPI network to construct $\mathcal{N}'$. That is, two genes $g_1, g_2 \in \mathcal{G}$ are connected by an edge $g_1 g_2 \in \mathcal{E}'$ if their encoded proteins are connected by an edge in BioGRID. Note that in BioGRID, two proteins are connected by an edge if and only if a physical interaction between them is supported by at least two publications or two experimental systems (e.g. affinity purification-mass spectrometry and yeast-2-hybrid). However, NeEDL can be easily extended to use other SNP-to-gene mappings and/or gene-gene networks. For each SNP $s \in \mathcal{S}$, let $\pi[s] = \{g \in \mathcal{G} \mid (s, g) \in \pi\}$ be the image of s under the mapping π. Equipped with π and $\mathcal{N}'$, we construct $\mathcal{N}$ by connecting to SNPs $s_1, s_2 \in \mathcal{S}$ by an edge $s_1 s_2 \in \mathcal{E}$ if and only if at least one of the following conditions holds:

- The SNPs s_1 and s_2 are mapped to at least one common gene, i.e. $\pi[s_1] \cap \pi[s_2] \neq \emptyset$.
- There is a SNP s_3 and genes $g_1, g_2 \in \pi[s_3]$ such that $g_1 \in \pi[s_i]$ and $g_2 \in \pi[s_2]$.
- The SNPs s_1 and s_2 are mapped to genes that are connected in the gene-gene network $\mathcal{N}'$, i.e. there is an edge $g_1 g_2 \in \mathcal{E}'$ such that $g_1 \in \pi[s_1]$ and $g_2 \in \pi[s_2]$.

Note that, with this construction, SNPs that are left unmapped by π correspond to isolated nodes in $\mathcal{N}$ and are hence not contained in any feasible solution of the optimization problem specified in Equation (2). In NeEDL, we hence remove all unmapped SNPs before (heuristically) solving the optimization problem, which further reduces the size of the search space.

Besides the dbSNP mapping procedure described above, we also assessed a mapping based on expression quantitative trait loci (eQTL) information. We downloaded all available studies with gene expression data and condition label 'naive' from the eQTL Catalogue (44). The studies in the catalog contain data from various different tissues. We converted the Ensembl gene identifiers to gene symbols with the help of biomaRt (45) to be compatible with the BioGRID dataset. The retrieved data was condensed into a single tabular file containing the dbSNP rsID, the gene name, the tissue, and the reported P-value. At runtime, NeEDL allows the user to select all tissues or a subset of those of interest. All mapping entries are retrieved from the mentioned condensed eQTL file and filtered for the given criteria. Afterward, NeEDL filters the remaining interactions based on their Benjamini-Hochberg-corrected P-value (46). For comparison with the dbSNP-based mapping, we selected all mapping information from all available tissues and applied a significance level of 0.05 after multiple testing correction. We then conducted a NeEDL run with the same parameters as our dbSNP-based runs but exchanged the dbSNP mapping with the eQTL mapping.

Local search with multi-start and simulated annealing

We use the local search with multi-start and simulated annealing to heuristically solve the optimization problem specified in Equation (2). NeEDL supports all epistasis models benchmarked by Blumenthal *et al.* (8). Simulated annealing is a general meta-heuristic to escape local optima in local search by accepting deteriorations from local optima with probabilities that decrease as the algorithm runs (47). For NeEDL, we adapted a simulating annealing algorithm for graph edit distance computation proposed by Riesen *et al.* (48), using its implementation available in GEDLIB (49,50).

We decided to use local search to solve our optimization problem instead of other meta-heuristics such as ant colony optimization (51) or genetic programming (52) because it can be easily parallelized and yields anytime algorithms that can be interrupted before convergence. Both of these properties

are important in the context of epistasis detection where runtime efficiency is crucial. Similarly, there are various other strategies besides simulated annealing which can be used to escape local optima in local search, e.g. variable neighborhood descent (53) or randomized postprocessing (54). Again, runtime considerations were the main reason for our choice of simulated annealing as a strategy to escape local optima.

Algorithm 1 provides a high-level description of our algorithm. It computes a set $\mathcal{S}^\star \subseteq \mathcal{F}_{LB,UB}$ of up to n locally optimal SNP sets, where n is a hyper-parameter that can be set by the user and $\mathcal{F}_{LB,UB} = \{S \subseteq \mathcal{S} \mid LB \leq |S| \leq UB \wedge \mathcal{N}[S] \text{ is connected}\}$ is the set of all feasible solutions. Computation of the n locally optimal SNP sets is parallelized, which is the main reason for NeEDL's excellent runtime performance (see 'Implementation Details' for details).

Algorithm 1: Overview of local search and simulated annealing.

Input: Genotype matrix $\mathbf{G}$, phenotype vector $\mathbf{y}$, SSI network $\mathcal{N} = (\mathcal{S}, \mathcal{E})$, number of initial SNP sets n, lower bound on SNP set size LB, upper bound on SNP set size UB, time limit L, maximal number of iterations I, start and end probabilities $1 > p_1 > p_I > 0$.
Output: Set $\mathcal{S}^\star \subseteq \mathcal{F}_{LB,UB}$ of up to n locally optimal SNP sets.

$\mathcal{S}^\star \leftarrow \emptyset$; $\alpha \leftarrow (\log(p_1)/\log(p_I))^{1/(I-1)}$;
parallel for $j \in \{1, \ldots, n\}$ **do**
  $S^\star \leftarrow$ `connected-subgraph`$(size = LB)$;
  $f^\star \leftarrow f_{\mathbf{G},\mathbf{y}}(S^\star)$; $S' \leftarrow S^\star$; $f' \leftarrow f_{\mathbf{G},\mathbf{y}}(S^\star)$;
  `converged` $\leftarrow$ `False`; $\Delta \leftarrow 0$; $i \leftarrow 1$;
  while $i \leq I$ **and not** *time limit reached* **and not** `converged` **do**
    $S'' \leftarrow \arg\min\{f_{\mathbf{G},\mathbf{y}}(S) \mid S \in \mathcal{F}_{LB,UB}(S')\}$; $f'' \leftarrow f_{\mathbf{G},\mathbf{y}}(S'')$;
    if $f'' < f^\star$ **then**
      $S' \leftarrow S''$; $f' \leftarrow f''$;
      if $f' < f^\star$ **then**
        $S^\star \leftarrow S'$; $f^\star \leftarrow f'$;
    else
      $\delta \leftarrow (f'' - f')$; $\Delta \leftarrow \Delta + \delta$; $\overline{\delta} \leftarrow \Delta/i$; $p_i \leftarrow p_1^{\alpha^{-(i-1)}}$;
      $r \leftarrow$ `random-number`$(min = 0, max = 1)$;
      if $r < p_i^{\delta/\overline{\delta}}$ **then**
        $S' \leftarrow S''$; $f' \leftarrow f''$;
      else
        `converged` $\leftarrow$ `True`;
    $\mathcal{S}^\star \leftarrow \mathcal{S}^\star \cup \{S^\star\}$;
return $\mathcal{S}^\star$;

To construct the j^{th} locally optimal SNP set, we initialize it as a randomly drawn connected subgraph of size LB (again, our objective to keep NeEDL's runtime low was the reason why we decided against more elaborate seeding strategies). We maintain the best encountered SNP set $S^\star$, as well as the currently processed SNP set S', together with their objective values $f^\star$ and f'. As long as a given time limit or a maximal number of iterations has not been reached and our descent in the jth region of the search space has not converged, we compute a candidate SNP set S'' as the best SNP set among the set $\mathcal{F}_{LB,UB}(S')$ of all locally reachable SNP sets from S'. We define $\mathcal{F}_{LB,UB}(S') = \{S \in \mathcal{F} \mid |S \triangle S'| = 1\}$ as the set of all feasible SNP sets that differ from S' in exactly one SNP (details on this step are provided below). We then check whether S'' improves the currently processed SNP set S'. If so, we update S' and also $S^\star$ if S' is now the best encountered SNP set.

On the other hand, if the objective value S'' exceeds the objective value of S' by $\delta > 0$, we nonetheless accept it as our new currently processed SNP set with probability $p_i^{\delta/\overline{\delta}}$, where $\overline{\delta}$ is the mean deterioration from the objective value of the currently processed SNP set up to the current iteration i. That is, the larger δ, the lower the probability of accepting S'' as our new currently processed SNP set. The baseline acceptance probability p_1 for the first iteration is a hyper-parameter that can be provided by the user (defaulted to $p_1 = 0.8$ in NeEDL). As the algorithm runs, it is updated using a suitably defined cooling factor α such that, when the maximum number of iterations I has been reached, it equals a user-specified end probability $p_I < p_1$ (defaulted to $p_I = 0.01$ in NeEDL).

To compute S'', we exhaustively enumerate the set $\mathcal{F}_{LB,UB}(S')$ of all locally reachable SNP sets from S'. This can be done by constructing modified SNP sets using the following three types of edit operations:

- SNP insertions (only if $|S'| < UB$): Insert an SNP $s \in \mathcal{S} \setminus S'$ that is connected to one of the SNPs already contained in S'.
- SNP removals (only if $|S'| > LB$): Remove a SNP $s \in S'$ such that the induced subgraph $\mathcal{N}[S' \setminus \{s\}]$ remains connected (i.e., s must not be an articulation point).
- SNP substitutions: Substitute a SNP $s_1 \in S'$ by a SNP $s_2 \in \mathcal{S} \setminus S'$ such that the induced subgraph $\mathcal{N}[(S' \setminus \{s_1\}) \cup \{s_2\}]$ remains connected (i.e., if s_1 is an articulation point, s_2 must bridge the connected components resulting from removal of s_1).

Statistical methods

We generate random SNP sets for empirical estimation of the false-positive rate. For epistasis detection tools other than NeEDL, we sample 1000 random SNP sets of size two as a baseline, whereas for NeEDL, we estimate the distribution of the number of SNPs in candidate sets based on the findings of our results and then randomly sample 1000 candidate SNP sets following this size distribution.

We generate two random SSI networks to evaluate the information that can be gained by transforming a PPI network into an SSI network. We either shuffle the labels to obtain a topology-preserving SSI network or shuffle the edges by a probability function to obtain an expected degree-preserving network as discussed in Lazareva *et al.* (55). Either way, this leads to a network where each SNP has other neighbors in the SSI network that we expect to result in the loss of biological information (Supplementary Figure S2A).

All P-values obtained by NeEDL and all competitors are corrected with the Benjamini-Hochberg procedure (46). In addition, we provide a baseline of randomly chosen candidate SNP sets, whose SNPs do not need to be connected via the SSI network. The baseline for the competitors contains 1000 randomly sampled sets with the size $|s| = 2$, as all competitors (except PoCos, which return hundreds of SNPs, and since the statistical scores cannot be calculated with such a high number of SNPs, PoCos was excluded) only return candidate SNP sets with the size of two. NeEDL finds candidate SNP sets with variable sizes $|s| \in [2, \infty[$.

Phenotype shuffling

We compared our original results with runs on datasets where we shuffled the phenotype vector while keeping the genotype matrix unchanged. These newly created datasets thus do not have any connection between the individuals' genotypes and phenotypes. We repeated this procedure 100 times for each dataset to reduce the influence of random effects on our analysis. This procedure led to 100 shuffled datasets for each disease. Due to the computational complexity of this analysis, we restricted the experiments to the three diseases BD, T2D and RA. We conducted one NeEDL run, and one Linden run for each shuffled dataset. We excluded MACOED from this analysis as the tool's extremely high runtime and memory requirements made it impossible to run it 300 times with our computational resources. Further, we replicated our analysis of the influence of the SSI network on these shuffled datasets as well. For that, we ran both of our network randomization methods ten times on each of the 100 phenotype-shuffled datasets,

resulting in 2,000 NeEDL runs per disease. In total we conducted 6300 NeEDL runs and 300 Linden runs on the shuffled datasets (Supplementary File 1).

Gene set enrichment analysis

Affected genes present in the highest ranks of penetrance SNP combinations for LOAD, IBD, and T1D were subjected to GSEA (http://www.gsea-msigdb.org/gsea/index.jsp, date of last access 8 March 2023) to determine if statistically significant enrichment in specific Gene Ontology (GO) gene sets could be identified (56,57). The query gene sets for GSEA are contained in Supplementary File 2.

Quantum computing

While large, fault-tolerant quantum computers are expected to outperform classical devices, complexity theory suggests that quantum computing may not be able to solve combinatorial optimization problems in polynomial time, similar to classical computing. Nevertheless, Grover's algorithm (58) can provide a quadratic speedup in terms of query complexity for quantum circuit-based quantum hardware (Supplementary Materials 1 and 2), and quantum annealers (Supplementary Materials 1 and 2) can provide the same speedup (59), under ideal, noiseless conditions. For particular instances of NP-hard optimization problems, super-polynomial speedups are possible (60), and certain problems can achieve even larger speedups by exploiting their specific structural features (61). Currently, only prototypes of quantum hardware are available, and heuristics must be run on them, with satisfactory results depending on the problem (62). The relation between Grover's algorithm and the heuristics used in our work is detailed in Supplementary Material 4.

We propose a novel quantum-enhanced method to address the initial seeding problem in NeEDL. Specifically, instead of selecting the initial seed randomly, our method employs an optimization technique to identify the most promising candidates. The underlying idea is that a better initial seed can significantly improve the quality of local search and reduce the time required to obtain high-quality solutions, assuming the optimization technique adds negligible overhead. To achieve this, we utilize a variant of the maximum clique problem to analyze the matrix of pairwise SNP correlations and identify the candidate set of SNPs that maximizes the sum of pairwise correlations. This problem is formulated as a quadratic unconstrained binary optimization (QUBO) whose objective function to minimize is defined as a quadratic polynomial over binary variables and consists of a linear combination of the solution cost and the associated constraints. By adjusting the weights of the linear combination, we can enforce or relax different constraints, such as the size of the candidate set. We formulate the QUBO problem as follows:

$$Q = \sum_{\ell=1}^{N} \left[\lambda_0 \left(\sum_i x_{i\ell} - K \right)^2 - \lambda_1 \sum_{(i,j)\notin E} w_{i\ell,j\ell} x_{i\ell} x_{j\ell} \right] + \lambda_2 \sum_{\ell=1}^{N} \sum_{m=\ell+1}^{N} \sum_{i=1}^{|V|} x_{i\ell} x_{im} \quad (5)$$

The formulation returns, saved in the matrix of binary variables x, N candidate sets of SNPs, each consisting of K SNPs. The parameter λ_0 penalizes candidate sets that are not of size K, λ_1 rewards sets with a higher sum of pairwise interactions w, and λ_2 penalizes solutions with similar sets (63). For additional details, see Supplementary Materials 3. We solve the optimization problem by applying quantum annealing on a D-Wave quantum annealer, QAOA (quantum approximate optimization algorithm) on IBM Perth and IBM Lagos quantum processing units, thermal annealing (64), parallel tempering (65), and Gurobi (66) on classical hardware as a classical baseline. These algorithms are heuristics, and their potential speedup must be determined empirically on a case-by-case basis. The functioning of these quantum algorithms is detailed in Supplementary Materials 4, Supplementary Figures S3 and S4.

To address the limitations in quantum hardware resources, we utilize the community detection method Leiden (67) to identify densely connected sections of arbitrary size in the SNPs network. This approach allows us to divide the problem into smaller sub-instances, enabling efficient computation on any quantum hardware. Specifically, quantum annealing algorithms can handle problem sizes on the order of 100 SNPs, while IBM quantum processing units can manage up to 10 SNPs. Our quantum software module called *quepistasis*, implements the seeding procedure using the optimization technique. It is responsible for creating the QUBO formulation based on the pairwise correlation of SNPs and some configuration, interfacing with the platforms of D-Wave and IBM (and also supporting quantum hardware available on Microsoft Azure), and returning the set of SNPs. Its structure is shown in Supplementary Materials 4 and Supplementary Figure S5).

We test NeEDL with the quantum computing seeding procedure over three different datasets: BD, RA, and T2D. The purpose of our experiments is to verify how much quantum computing helps in finding better seeds and thus speeds up the local search process. To accommodate the limited amount of resources (namely the number of qubits and the number of gates) of the quantum computer, we subsample each of the three datasets into three different subsets, having 100, 500 and 1000 SNPs. Minimizing the quantum resources is essential for keeping the *fidelity*, which is the measure of the accuracy of the actual quantum operation compared to the ideal one, above a certain threshold, guaranteeing reliable results. The timing information in Supplementary Figure S5 is obtained from the logs generated by NeEDL. For the quantum procedure run on IBM devices, such timing is spoiled by the queue we must wait to execute each task, and such time is normalized. The scalability discussion in Figure 5 is obtained by linear interpolation of the seeding time using each technique for different sample sizes of each dataset. We show the experimental setup and provide some comments about the results obtained in Supplementary Materials 6, Supplementary Figure S6, Supplementary Table S12.

Results

NeEDL outperforms existing epistasis detection tools

We compared NeEDL to the existing epistasis detection tools LinDen (10) and MACOED (12) on the same input data as NeEDL. All tools were run with default hyper-parameters (see Material and Methods, for NeEDL, see Supplementary Tables S13– S15). PoCos (11), Potpourri (9) and BiologicalEpistasis (13) were excluded due to unresolvable implementation problems and/or excessive runtime or memory requirements

(Supplementary Table S16). To ensure fair competition, we compared the SNP sets using four widely used scores in this research area (8), including the *P*-value of the χ^2-test (optimized for by LinDen), the Bayesian network score K2 (optimized for by MACOED), the negative log-likelihood score of the maximum likelihood model (MLM) (optimized for by NeEDL per default) and the differences between the negative log-likelihoods of fitted linear and quadratic regression models (NLL gain) (8). For the *P*-value of the χ^2-test, the Bayesian network score K2, and the MLM score, low scores indicate promising SNP sets. Conversely, the higher the NLL gain, the more promising the scored SNP set (see Material and Methods). To show the robustness of the results and to control if higher-order candidate SNP sets achieve more significant associations just because of the number of SNPs, we generated baselines by randomly sampling 1000 size-matched SNP sets (size 2 for LinDen and MACOED, higher-order for NeEDL) to estimate the false positive rate (see Material and Methods). All *P*-values (both for the background distributions and the results) were adjusted using the Benjamini-Hochberg procedure (46).

Since the local search algorithm returns a result for each start seed, we obtain thousands of results. To find a suitable cut-off, we show the MLM and K2 scores for the 100 top-ranked SNP sets computed by NeEDL (Figure 2A, Supplementary Figures S7–S9) Interestingly, we observe phase transitions (i.e., an abrupt change to a lower significance of the score over linear increasing ranks) for both scores, which typically occur after 50 SNP sets (except for the LOAD and IBD datasets, with transitions between the ranks 20 and 30, and the RA dataset, with transitions at 60 SNP sets). We hypothesize that this score gap represents the barrier between promising and not promising epistasis candidate sets and thus focused on the top 50 candidates (i.e. where most transitions occur, Supplementary Figure S8) selected based on the respective metric each tool is optimizing for (NeEDL: MLM score, LinDen: *P*-value of χ^2-test, MACOED: K2 score).

MACOED was executed 100 times to account for randomized subroutines, leading to the union of the top 50 candidate SNP sets for each run for MACOED. The data shown in Figure 2B, Supplementary Figure S7B, and Supplementary File 1 demonstrates that NeEDL outperforms LinDen, MACOED, and the baselines across all metrics and datasets. Even in instances where it performs comparably on some metrics on certain datasets, NeEDL still surpasses the competitor on at least one other metric across all datasets in terms of statistical association. LinDen, as well as the second-order- and higher-order baseline, are outperformed on all metrics over all datasets. On all datasets, the χ^2-test *P*-values obtained for the SNP sets computed by NeEDL approach the maximum number of decimal places available in standard computation. With the more complex MLM and K2 scores, differences become more visible, and we see that NeEDL markedly outperforms MACOED in all datasets (Figure 2B, Supplementary Figure S7B). It is noteworthy that when the cut-off of the rank is set at 25 for LOAD, and IBD and 60 for RA, the advantage of NeEDL over MACOED becomes more distinct (Supplementary Figure S9). NeEDL also yields significantly lower K2 scores than MACOED on all datasets, even though NeEDL did not optimize for this score. Since NeEDL optimizes for the MLM score, we can see that NeEDL outperforms all competitors and baselines on this score. In terms of NLL gain, NeEDL again outperforms its competitors on all datasets except IBD, where the NLL gain is similar to MACOED. Also, for LOAD, very low values of the NLL gain in MACOED indicate that, unlike NeEDL, MACOED mainly finds SNP sets with strong additive main effects. In terms of required computational resources, NeEDL and LinDen (both executable on desktop computers) are much more efficient than MACOED, which requires a super-computer (Supplementary Table S17).

Figure 2C and Supplementary Figures S7C and 9 show that, in most datasets, the best scores are achieved with three to seven interacting SNPs. Since we ran NeEDL with a maximum SNP set size of 10, NeEDL likely does not exploit combinatorial statistical artifacts but prioritizes SNPs that may be involved in protein-protein interactions and thus be more likely to be biologically meaningful. Additionally, we investigated if eQTL-based SNP-to-gene mappings could provide more insights into SNP-SNP relationships than the mostly positional mappings provided by dbSNP. For this, we constructed an SSI network based on eQTLs from the eQTL Catalogue (44) (see Material and Methods for details) and then ran NeEDL on this network. Supplementary Figure S10C shows the obtained MLM scores in comparison to the scores obtained with the dbSNP-based network. The eQTL-based approach leads to worse scores on all of our datasets. A possible explanation for this is that the eQTL-based network is substantially smaller than the dbSNP-based network (Supplementary Figure S10A and B), suggesting that information on eQTLs is still too sparse to be useful for epistasis detection.

To further examined whether the use of the biological priors encoded in the dbSNP-based SSI network indeed guides NeEDL toward more promising SNP sets, we generated 100 randomized networks with preserved topology (by shuffling the node labels) and 100 randomized networks where the SNPs' node degrees are preserved in expectation (Supplementary Figure S2A). For this analysis, we randomly selected three of the eight datasets (BD, T2D and RA) to avoid excessive runtime. MLM scores on the original SSI networks are, on average, better than those on the randomized networks (Figure 2D).

NeEDL identifies epistatic associations beyond random effects

To investigate whether the scores obtained by NeEDL may be explainable by random effects alone, we ran NeEDL and its competitors (except for MACOED due to its extremely high resource requirements) on the BD, T2D and RA datasets with randomly shuffled phenotypes (Results in Supplementary File 1). We observe that the score gaps we obtained with the original phenotypes after around 50 SNP sets disappear (Figure 3A). Moreover, w. r. t. all four scoring metrics, NeEDL yields substantially better results for the original dataset than for the phenotype-shuffled datasets, and the gap between NeEDL and LinDen disappears for the shuffled phenotypes (Figure 3B). Further, the SNP sets returned by NeEDL for the shuffled datasets contain more SNPs (Figure 3C) than those returned for the original datasets (Figure 2C), potentially due to the fact that, in the absence of a true signal in the data, a higher number of variables leads to an improvement of the statistical scores because of combinatorial statistical artifacts. Additionally, we observe that, for shuffled phenotypes, running NeEDL with a PPI-based SSI network does not lead to improved scores compared to random networks (Figure 3D)—in contrast to the results obtained for the orig-

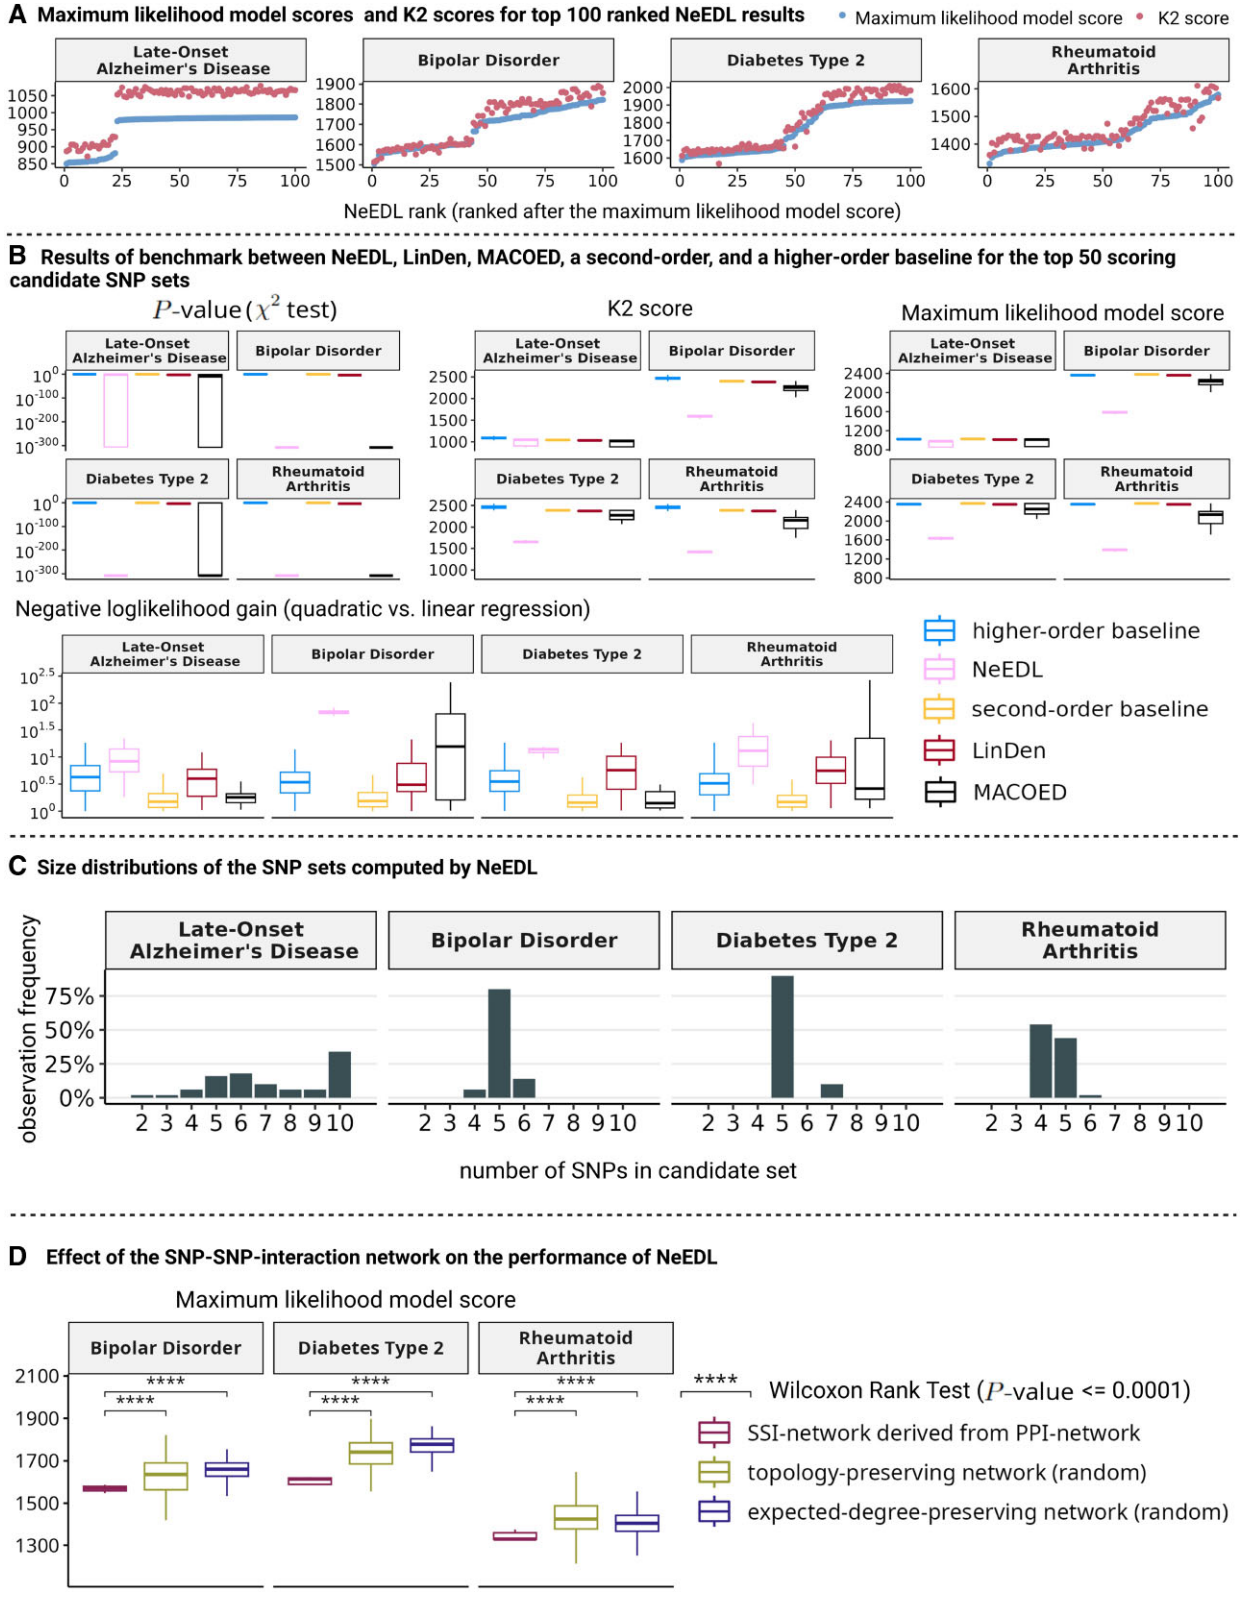

Figure 2. Quantitative evaluation of the SNP sets computed by NeEDL for Late-Onset Alzheimer's Disease, Bipolar Disorder, Diabetes Type 2, and Rheumatoid Arthritis. Results for Coronary Artery Disease, Diabetes Type 1, Hypertension, and Inflammatory Bowel Disease can be found in Supplementary Figure S7. (**A**) Visualization of the maximum likelihood model score and the K2 score of the top 100 candidate SNP sets ranked by NeEDL. (**B**) A benchmark study between NeEDL, LinDen, MACOED, a second-order baseline (i.e. 1000 random sampled pairs of SNPs), and a higher-order baseline (i.e. 1000 randomly sampled sets consisting of multiple SNPs) shows that NeEDL outperforms in statistical significance existing epistasis detection tools w. r. t. four different evaluation metrics. (**C**) Analyzing the number of SNPs included in NeEDL's output SNP sets reveals that the most promising SNP sets typically contain between three and seven SNPs. (**D**) Comparing maximum likelihood model scores of NeEDL results against those obtained using randomized networks demonstrates that the use of the SSI network indeed leads to the discovery of more promising SNP sets.

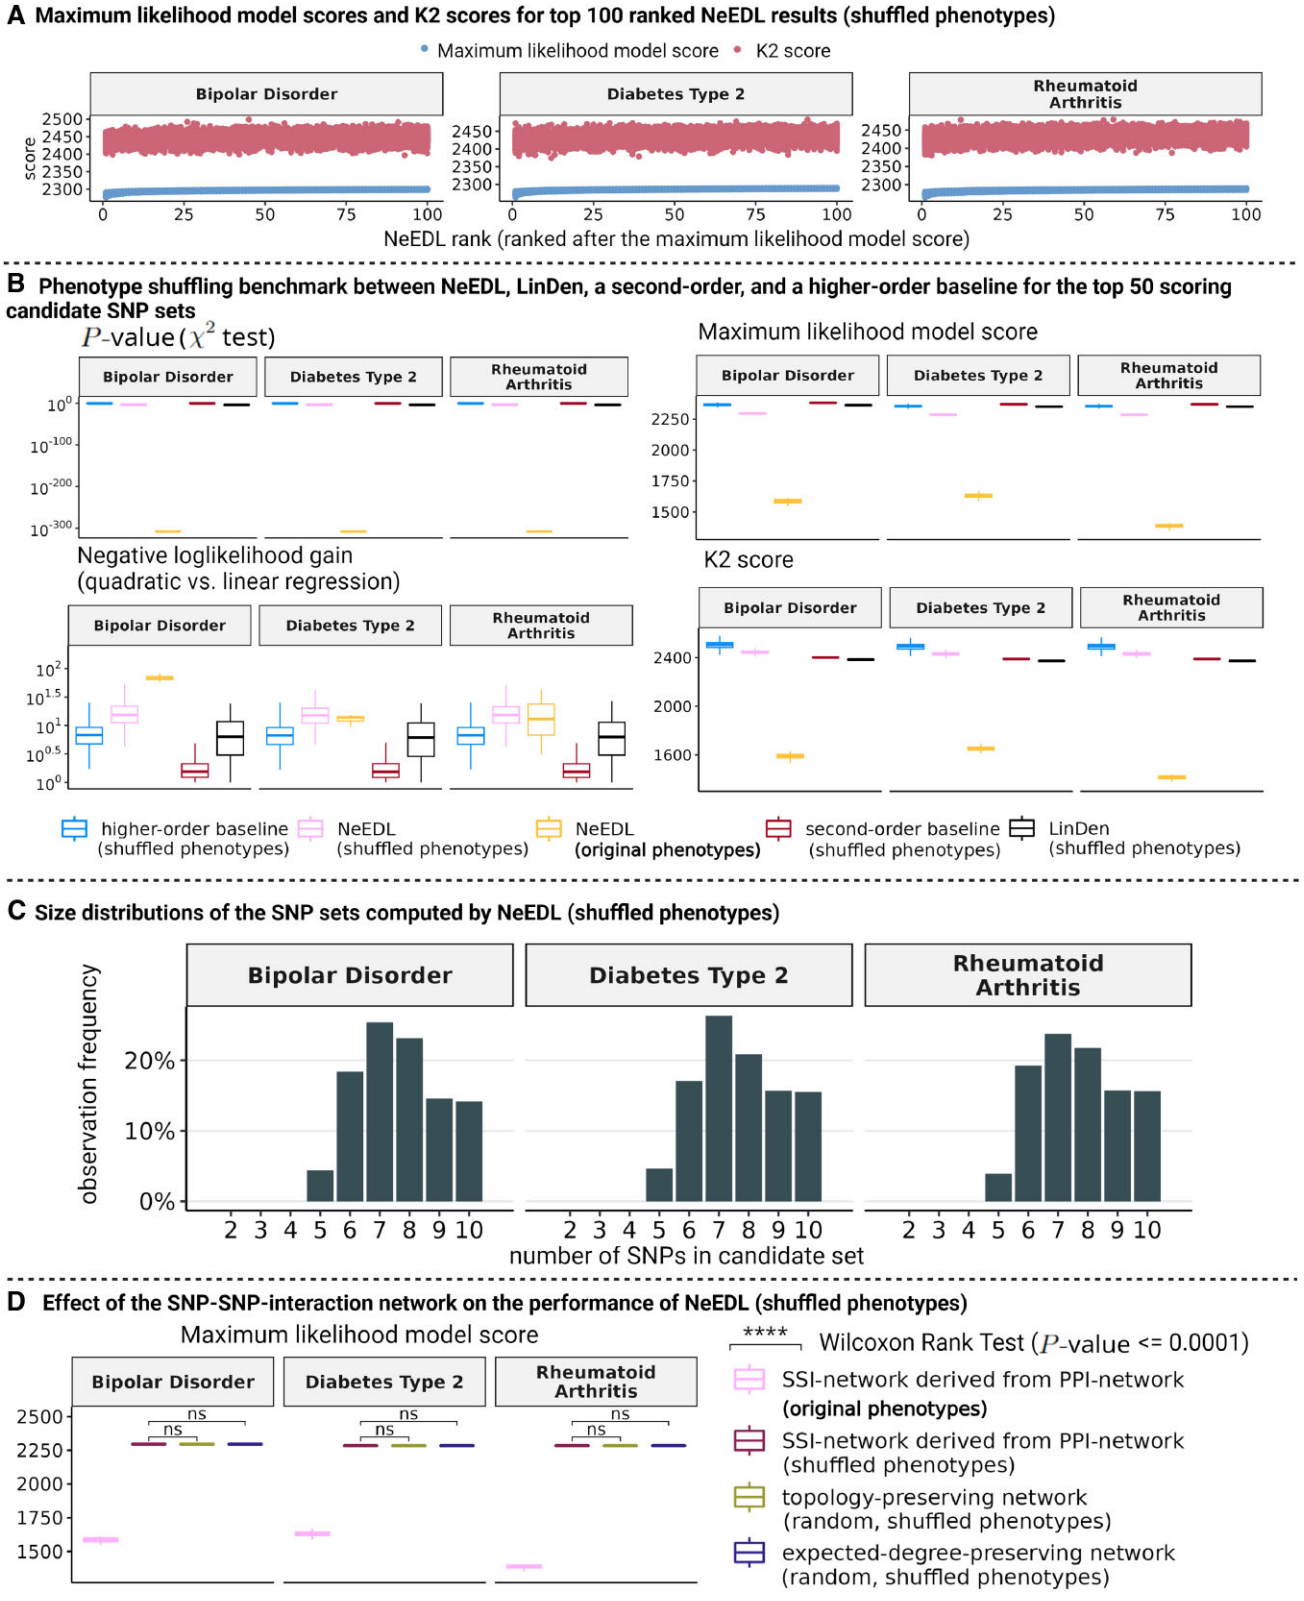

Figure 3. Evaluation of the SNP sets computed by NeEDL and LinDen with shuffled phenotypes for Bipolar Disorder, Diabetes Type 2, and Rheumatoid Arthritis. (**A**) Maximum likelihood model and K2 scores for the top 100 SNP sets reported by NeEDL (ranked by the maximum likelihood model score). (**B**) Comparison of the scores of the top 50 results of NeEDL and LinDen and the second- and higher-order baseline. As a reference, the scores of the original NeEDL runs without phenotype shuffling are also shown (yellow box plots). (**C**) SNP set size distribution of NeEDL's results. (**D**) Maximum likelihood model scores obtained by NeEDL with and without network rewiring and phenotype shuffling.

inal phenotypes where we did observe statistically significant improvements (Figure 2D). Overall, these results are strong evidence that NeEDL indeed picks up on true biological signals and that its usage of a PPI-based SSI network is crucial for this.

Consistent results across independent cohorts

We conducted a replication study to validate the findings of NeEDL and to demonstrate the consistency of the results across independent datasets for LOAD. For this, we retrieved two independent replication datasets (independent from the discovery dataset used in the benchmark) from the UK Biobank (see Material and methods), one with individuals of British ancestry (similar to the LOAD dataset) and one with individuals of mixed ancestry (i.e. 10% of the samples are of mixed ancestries like Latino, African, or Asian) to estimate population-specific effects. From the SNP sets computed by NeEDL on the LOAD dataset (i.e. the dataset used for discovery), we selected the 50 SNP sets with the most significant MLM scores (to be consistent with the benchmark) and recomputed their χ^2-test *P*-values, MLM scores, K2 scores and NLL gains.

Figure 4 a shows that we could replicate the phase transition observed in the LOAD dataset used for the benchmark (Figure 2A): The top 22 SNP sets and the top 23 SNP sets of the replication data set of British ancestry and mixed ancestry, respectively, replicate with a significant χ^2-test *P*-values <0.05 (Supplementary Figure S11). The MLM scores of replication correlate (Pearson correlation of 0.995) highly with the MLM scores of the benchmark study (Figure 4B). We can further observe a drop in the Pearson correlation to 0.993 in the mixed ancestry dataset, suggesting that the population-specific genetic architecture might play a role in epistasis detection. Similar to the LOAD dataset, a phase transition occurs in *P*-values and MLM scores after the top 22 SNP sets (Supplementary Figure S11). We also re-computed χ^2-test *P*-values, MLM scores and K2 scores for the 50 best SNP sets computed by MACOED and LinDen on the LOAD benchmark dataset. Consistent with the discovery study, NeEDL with its top 50 candidate sets outperforms LinDen in both replication datasets (Figure 4C). MACOED is outperformed in the *P*-values of the χ^2-test, the K2 score, and the MLM score. However, MACOED yielded similar results or slightly less significant scores with the NLL-gain. Overall, the replication results are consistent with the results of the benchmark in the discovery dataset.

Epistasis candidates uncovered by NeEDL are biologically plausible

By utilizing NeEDL, we identified potential epistatic candidate SNP sets that have been selected based on their maximum likelihood model scores. We selected LOAD, T1D and IBD (the latter two in Supplementary Materials 7) as they are well-represented in the literature. In each of these diseases, various SNPs of a single gene in different combinations with SNPs of several other genes are present in the majority of the statistically most significant candidate SNP sets (Figure 4B, Supplementary Figure S11). We evaluated the top genes before the gap that were affected by the predicted SNP interactions by performing gene set enrichment analysis (GSEA) (56,57) to identify statistically significant enrichment in Gene Ontology (56) and KEGG (68) (see Materials and methods for details).

APOE was present in the 25 highest penetrance SNP combinations for the disease LOAD. APOE and the lipoprotein metabolism pathway have been previously recognized as the most critical gene and pathway for Alzheimer's disease risk (24,69). Since APOE mutation status alone is not a reliable predictor of Alzheimer's disease, a better understanding of epistatic interactions may be instrumental in developing better diagnostics. NeEDL identified 47 additional genes exhibiting potential epistatic SNP combinations with APOE (Supplementary File 2). GSEA determined that 15 genes (APOE, SORL1, SMAD3, LRP2, YBX3, TP63, PRKCB, DAB1, LRP1, IFI16, UBQLN1, A2M, DYRK1A, HMGXB4 and TRAF3IP1), including APOE, were enriched in a gene set that includes the process to reduce the response to the stimulus (GO:0048585; $k/K = 0.0088$; *P*-value $= 6.25 \times 10^{-11}$; FDR *q*-value $= 9.82 \times 10^{-7}$). Of note, one of the epistatic SNP combinations identified by NeEDL was rs429358 and rs7412 in the APOE gene, which has been previously identified as a potential epistatic SNP combination (70) (note that the good results of NeEDL's competitor MACOED on the LOAD dataset can be explained by the fact that this combination of SNPs was found also by MACOED).

In the Human Protein Atlas (71), we can observe that the identified genes are significantly expressed in tissues that are reported to be affected by LOAD, T1D or IBD (Supplementary Figure S12) (72–84).

Quantum computing improves seeding of local search

The GWAS datasets employed in this study cover up to 140 000 SNPs due to necessary cleaning and filtering steps (Supplementary Tables S2–S10), a small fraction of the 84.7 million SNPs (85) in the human genome. It is hence likely that, in the future, we will see a sharp increase in the number of covered and mappable SNPs. In the NeEDL workflow, this would result in an SSI network containing millions of SNPs. To explore such a huge network via randomly seeded local search with reasonable coverage, we would have to dramatically increase the number of initial solutions—certainly, beyond the capacities of classical high-performance computing clusters. We thus explored the use of quantum computing in NeEDL (see Materials and methods for details) to find a promising set of initial solutions for the local search instead of starting with random seeds as in the classical implementation (Supplementary Figure S2B). To be able to use a quantum computer, we modeled the problem of finding promising initial solutions as a variant of the maximum clique problem, transformed it into an instance of the QUBO problem (86), and then solved it heuristically using several quantum computing machines and simulators, including the quantum annealer D-Wave Advantage 6.1, the superconducting-based IBM Perth and IBM Lagos devices.

To test this approach, we subsampled our BD, RA, and T2D datasets to 100, 500 and 1000 SNPs and ran NeEDL with random and quantum-computing-based seeding. Running quantum-computing-based seeding on the entire datasets is still infeasible due to limitations of current quantum computing hardware. Supplementary Figure S5 shows that when using quantum computing, fewer but higher-quality candidate SNP sets are returned, and that the speed up these high-

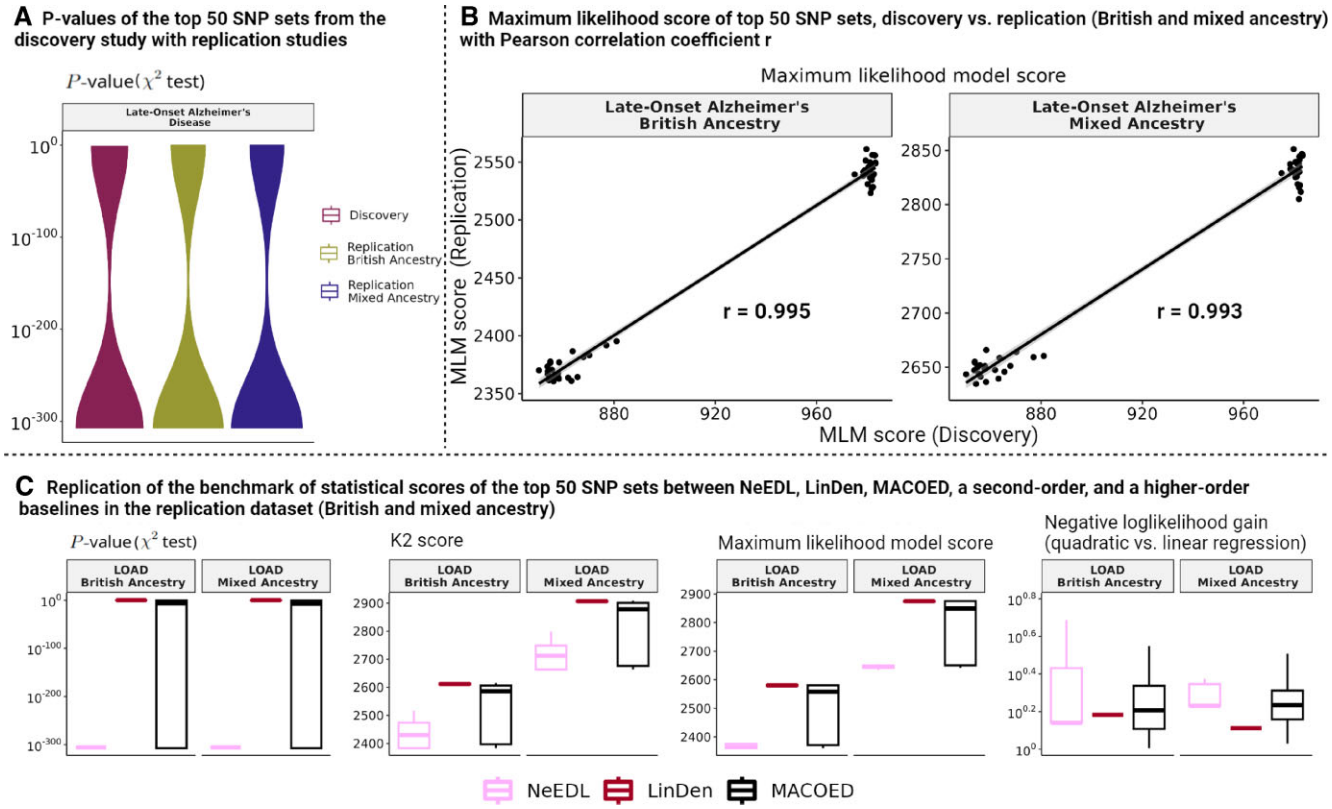

Figure 4. Replication studies in two independent data sets from UK Biobank with British and mixed ancestry. (**A**) *P*-values across discovery and the two replication studies. (**B**) Correlation of the MLM score between the discovery study and the two replication studies. (**C**) Replication of the benchmarking between NeEDL, LinDen, MACOED, a second-order baseline and a higher-order baseline in two replication data sets.

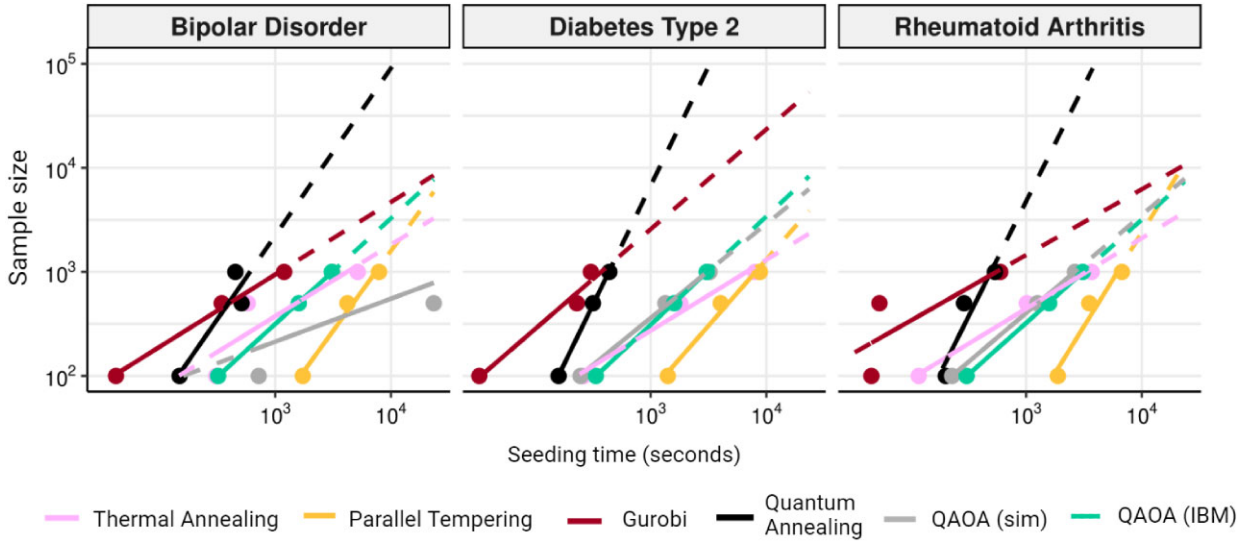

Figure 5. Experiments on the quantum computers: Expected scaling of the quantum hardware. The different quantum and classical devices used to perform the optimization process show different scaling, with the quantum devices (Quantum Annealing, QAOA) having a higher slope than classical devices (thermal annealing, parallel tempering, and Gurobi), suggesting a possible scaling advantage.

quality initial solutions induce in the downstream local search outweighs the additional runtime needed for the quantum-computing-based seeding. Since quantum computing hardware cannot yet process big data, we estimate the speed up on realistically sized GWAS data across different devices (Figure 5). Notably, any quantum device, whether simulated or executed on actual hardware, scales better than the baseline classical algorithm. We may hence anticipate quantum computation to surpass the classical baseline in the future, showcasing a potential quantum advantage once more powerful quantum computing hardware will become available, even though currently only small subproblems can be addressed. To conduct our quantum computing experiments, we used more than 4.1 million seconds of computing time on quantum resources.

Exploring epistatic interactions in the Epistasis Disease Atlas

Since more than two million CPU hours were necessary to calculate the results reported here, we make them widely and easily accessible through the web-based Epistasis Disease Atlas (https://epistasis-disease-atlas.com). The Epistasis Disease Atlas provides results of eight heritable diseases (see Section Datasets). It enables users to browse, visualize, interpret, and link candidate SNP sets and their interactions within the SSI network to published literature and a variety of external databases and tools (see Materials and methods for details, (87)). The Epistasis Disease Atlas can further be queried through an application programming interface (API) where we offer both R and Python packages (Supplementary Materials 8, including Supplementary Figures S12–S15).

Discussion

The initially high expectations towards genome-wide association studies have thus far not completely been fulfilled due to the missing heritability problem (4,5), which suggests that individual genetic variants cannot explain the genetic architecture of complex and polygenic heritable diseases. However, the missing heritability is likely overestimated since epistatic interactions are implicitly ignored in these estimates (88). While statistical models and computational tools for identifying epistatic interactions have been proposed, the space of possible interactions to be tested is vast and could thus far mostly tackle pairwise interactions. With NeEDL, we overcame this problem and suggest promising higher-order SNP sets. NeEDL outperforms the existing tools LinDen and MACOED, which are limited to pairwise interactions, across all commonly used statistical models and offers new insights into the genetic architecture of complex diseases. We show evidence that candidate SNP sets that are strongly associated with disease phenotypes are in the order of three to seven, are biologically plausible, and can be verified in independent cohorts.

The candidate SNP sets identified by NeEDL represent locally optimal solutions within specific areas of an SSI network and are thus limited by the quality of the priors in this network. While the current version of NeEDL generally outperforms random priors, we see room for further improvement in designing the SSI networks. For instance, PPI networks such as BioGRID may contain a substantial amount of missing interactions such that a globally optimal SNP set cannot be found by NeEDL, i.e., due to yet unknown or tissue-specific PPIs. Similarly, existing PPI networks neglect the consequences of alternative splicing on functional interactions (89,90). Furthermore, SNPs that alter codon efficiency during the translation process can reduce or inhibit protein activity (91), which ultimately can alter a PPI network. More disease-associated variants are found in non-coding compared to coding regions of the DNA, suggesting a major role of gene regulation in human diseases (92). This motivates further research in SSI networks that consider enhancer-gene interaction and chromatin remodeling. Including gene-regulatory interactions will also pose a significant challenge for interpretability, as the current strategy of GSEA and KEGG enrichment is limited to established interactions between genes and proteins. Moreover, for the dbSNP- and eQTL based SSI networks used in this study, we consider only SNPs that have previously been associated with genes in, respectively, dbSNP and the eQTL Catalogue. Other associations are neglected and possible misannotation in the external databases used to generate our SSI network are inherited. In summary, the higher the quality of the prior knowledge NeEDL can use to reduce the search space, the more meaningful the results will be.

We should further consider the limitations of the statistical models employed in NeEDL. Recent work by Blumenthal *et al.* (8) have demonstrated significant differences in the detection power of various statistical models, concluding that the MLM score offers the best tradeoff between quality and performance, outperforming the quadratic regression model, Bayesian model, and variance model. The *P*-values obtained through the χ^2-test are extremely small (approaching the minimum C++ double value) and preclude further differentiation between candidate SNP sets. Hence, we recommend using the MLM score. Covariates, such as age, comorbidities, environmental factors, lifestyle factors, and population structure, could negatively impact the discovery of epistatic interactions. While regression models can, in principle, control for such effects (which is also supported by NeEDL), this further increases the computational demand beyond currently available resources. Further algorithmic improvements in mining the SSI network, maybe through graph neural networks, are conceivable and will be a promising research direction (90). As a potential solution to further speed up the local search algorithm of NeEDL, we considered various quantum algorithms and devices, showing that epistasis detection poses a real-world scenario in which quantum advantage could be achieved and exploited. In the near future, quantum algorithms (potentially enhanced via more sophisticated approaches such as warm-started QAOA (93)) may hence allow NeEDL to handle increasingly complex datasets with millions of SNPs.

Replication of epistasis candidate SNP sets in independent cohorts, as shown here, can provide valuable confirmation of statistical associations and offer a testbed for assessing the clinical potential of epistatic interaction scores similar to polygenic risk scores. To date, multi-SNP epistasis hypotheses are difficult to verify experimentally. Despite these challenges, future gene editing studies will be needed for their verification and to unravel the functional consequences of epistatic interactions in disease pathophysiology. To optimally support this process, we developed the Epistasis Disease Atlas, a user-friendly and interactive web resource that will allow the biomedical community to extract testable hypotheses directly and work towards an improved understanding of disease mechanisms as the basis for developing targeted therapeutic interventions. We note that the application of NeEDL is wider than just human disease research and see the potential for applying NeEDL in model organisms such as Arabidopsis thaliana or Mus musculus, where gene editing and functional studies are more readily achievable.

In conclusion, NeEDL is an epistasis detection tool to leverage prior biological knowledge for the systematic discovery of higher-order candidate SNP sets and quantum computing. NeEDL thus lays the foundation for charting the complex genetic architecture underlying heritable diseases. NeEDL lifts epistasis detection to a systems-oriented network biology level and shows that seamless integration of quantum computing has the potential to transform computational genomics, highlighting its role in biomedical research.

Data availability

The LOAD data set is restricted and can be found at https://www.tgen.org. The users need to apply for the dataset. The BD (EGAD00000000003), CAD (EGAD00000000004), T1D (EGAD00000000008), T2D (EGAD00000000009), HT (EGAD00000000006), IBD (EGAD00000000005), RA (EGAD00000000007) and the British 1958 British Birth Cohort (EGAD00000000001) datasets are restricted and can be found at https://www.sanger.ac.uk/legal/DAA/MasterController and https://edam.sanger.ac.uk/#/. The user needs to apply for the datasets. The results of NeEDL stored in the Epistasis Disease Atlas are freely available under the CC BY-NC 4.0 license.

In the future, we plan to include the following datasets in the Epistasis Disease Atlas: AS (EGAD00000000010), ATD (EGAD00000000011), MS (EGAD00000000012), BRCA (EGAD00000000013), TB (EGAD00000000016), UC (EGAD00000000025), PD (EGAD00000000057), AK (EGAD00010000150), SP (EGAD00010000262), IS (EGAD00010000264), CD (EGAD00010000246); these data sets are restricted and can be found at https://www.sanger.ac.uk/legal/DAA/MasterController. The user needs to apply for the data sets. The datasets for replication are restricted and can be found at the UK Biobank database (www.ukbiobank.ac.uk), under project ID 54273.

A graphical overview of the whole project is provided in Supplementary Figure S16. Supplementary File 1 (all results and scripts to plot them) is available on figshare: https://doi.org/10.6084/m9.figshare.25780350.v1.

Code availability

The source code of NeEDL, the quantum computing module, and the R Shiny App is freely available under the GPLv3 license at GitHub: https://github.com/biomedbigdata/NeEDL. All features of NeEDL are dockerized and available at Dockerhub: https://hub.docker.com/r/bigdatainbiomedicine/needl.

Supplementary data

Supplementary Data are available at NAR Online.

Acknowledgements

We want to thank Christina Trummer for the design of the NeEDL logo and the Epistasis Disease Atlas logo, as well as the design of the webpage. We further want to thank Simona Dedurkeviciute for her advice on webpage design and architecture. Images were created with https://www.biorender.com. Parts of the images were designed using resources from Flaticon.com under a paid license. ChatGPT version 4, under a paid license, was used to rephrase parts of this manuscript. The database diagrams were created with the support of the LucidApp, which gratefully provided us with a free premium account. Grammarly, under a paid license, was used for spell and grammar checks. Overleaf, under a paid license, was used to draft the manuscript in LaTeX. Paperpile, under a paid license, was used to collect references in the right format. We would like to thank Matt Huentelman from The Translational Genomics Research Institute in Phoenix, Arizona, for sharing the TGen II cohort data (LOAD). This study makes use of data generated by the Wellcome TrustCase-ControlConsortium. A full list of the investigators who contributed to the generation of the data is available from www.wtccc.org.uk. Funding for the project was provided by the Wellcome Trust under awards 076113, 085475 and 090355. The authors gratefully acknowledge the Leibniz Supercomputing Centre for funding this project by providing computing time and support on its Linux Cluster. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer SuperMUC-NG at Leibniz Supercomputing Centre (www.lrz.de). We acknowledge the CINECA award under the ISCRA initiative, for the availability of high-performance computing resources and support (QPU D-Wave). Access to the IBM Quantum Services was obtained through the IBM Quantum Hub at CERN. The views expressed are those of the authors and do not reflect the official policy or position of IBM, the IBM Q team. We acknowledge support from Microsoft's Azure Quantum for providing credits and access to the quantum hardware used in this paper.

Author contributions: M.H.O. designed and developed most parts of the software architectural concepts, database design, and features of epiJSON, NeEDL, calculate_scores, the quantum-computing module, the Epistasis Disease Atlas, and the R Shiny App of the Epistasis Disease Atlas. MHO implemented the initial idea, conceptualized and conducted most analyses, wrote the manuscript, and held the role of the project manager and project leader. JP implemented large parts of NeEDL, benchmarked NeEDL against the competitors, and implemented the concept of the real-time calculation of the scores for the Epistasis Disease Atlas. M.I. implemented the quantum computing module of NeEDL and conducted several analyses on various quantum-computing hardware. S.B. conducted literature research, applied for most of the data sets, and implemented the R Shiny App of the Epistasis Disease Atlas. MHO and SB produced the overview video on the EDA. A.F., A.G.P. and K.B. conducted the replication studies of the potential candidate SNP sets in independent cohorts. AM and CH helped with the technical conceptualization and implementation of the backend (database, API and file server) of the Epistasis Disease Atlas. M.H.A. implemented and designed the front end of the Epistasis Disease Atlas. C.H. implemented the automatic pipeline to integrate new data sets into the database of the Epistasis Disease Atlas. N.T. implemented the features of the Integrated Genome Browser. K.A. implemented the R and Python packages of the Epistasis Disease Atlas. M.P. and M.W. conceptualized and executed the big memory runs and long-time runs of NeEDL and MACOED on the LRZ supercluster. E.S. implemented the epiJSON tool and the linkage-disequilibrium options of NeEDL. ES further reformatted the datasets of the eight discussed diseases into NeEDL readable format. P.W., H.F. and S.K. helped with analyses. M.V.H. implemented the co-variate and ancestry options. I.L. tested the software of NeEDL and the Epistasis Disease Atlas for bugs. LSCHW and LHAF conducted further research on epistatic SNPs in regulatory elements. G.G. and A.A. processed the raw Illumina data for the additional ten datasets to NeEDL readable format that are about to be integrated into the Epistasis Disease Atlas but not further discussed in this manuscript. L.H.A.C. and O.T. conducted further research on epistatic SNPs in alternative splicing site regions. N.W. designed the architecture of the Figures and revised the manuscript. SGL and GC conducted further research on epistatic SNPs in structural biology and binding sites. MC conducted further research on

epistatic SNPs that change translational speed. J.J., H.K.L., P.F. and L.H.E.N. conducted biological plausibility studies and drafted parts of the manuscript. A.Man. performed the theoretical analysis for the quantum computing module of NeEDL. F.M., H.C.G. and K.V.S. provided expertise in epistasis detection throughout the project, suggested algorithmic changes, and revised the manuscript. L.S.C.H.U., L.I,. and M.G. provided quantum-computing expertise and quantum-computing hardware and helped calibrate the systems. L.W. and D.E. provided expertise on the math for the co-variant and mixed-ancestry options of NeEDL. A.D.P., A.Man. and M.G. supervised the quantum-computing part of NeEDL. P.F. and L.H.E.N. supervised the biomedical part of the project. T.K., J.B., M.L. and D.B. jointly supervised the project. J.B. and T.K. had the initial idea for network-enhanced higher-level epistasis detection. DB had crucial ideas regarding algorithmic optimization solutions. J.B. and D.B. first suggested the use of quantum computing. D.B. implemented the statistical models and, together with MHO, wrote most parts of the manuscript. This project took more than five and a half years to implement, execute, and evaluate—which is the reason for eight first authors and five senior authors. All authors revised the manuscript and gave their approval.

Funding

Technical University Munich—Institute for Advanced Study, funded by the German Excellence Initiative (in part); Intramural Research Program (IRP) of the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK, in part); M.G. is supported by CERN through the CERN Quantum Technology Initiative; A.Man. is supported by the Foundation for Polish Science (FNP), IRAP project ICTQT [2018/MAB/5], co-financed by the EU Smart Growth Operational Programme; German Federal Ministry of Education and Research (BMBF) within the framework of the *e:Med* research and funding concept [*ZX1908A/01ZX2208A*, *01ZX1910D/01ZX2210D*]; European Union's Horizon 2020 research and innovation programme [777111]; this publication reflects only the authors' view and the European Commission is not responsible for any use that may be made of the information it contains; Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) [422216132, 516188180]; the research of LI and LSCHU is partially funded by the Bavarian State Ministry of Science and the Arts as part of the Munich Quantum Valley; European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie [813533 to M.L.F.P.M., 860895 to TranSYS]; FNRS convention PDR T.0294.24 'Expanded PRS embracing pathways and interactions for increased clinical utility'. Funding for open access charge: TUM (one of the corresponding authors markus.list@tum.de, has the right to publish at OUP journals). M.P. was supported by the German Federal Ministry of Education and Research (BMBF; Grant No, 031L0305A).

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author(s) used ChatGPT 4 under a paid premium license to rephrase the text for more clarity, conciseness, and grammatical correctness. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

Conflict of interest statement

During the course of the project, H.C.G. became a full-time employee of Novo Nordisk Ltd. M.W. is a founder and shareholder of MSAID GmbH and OmicScouts GmbH. He has no operational role in either company.

References

1. Heap,G.A., Trynka,G., Jansen,R.C., Bruinenberg,M., Swertz,M.A., Dinesen,L.C., Hunt,K.A., Wijmenga,C., Vanheel,D.A. and Franke,L. (2009) Complex nature of SNP genotype effects on gene expression in primary human leucocytes. *BMC Med. Genom.*, **2**, 1.
2. Bush,W.S. and Moore,J.H. (2012) Chapter 11: Genome-wide association studies. *PLoS Comput. Biol.*, **8**, e1002822.
3. MacArthur,J., Bowler,E., Cerezo,M., Gil,L., Hall,P., Hastings,E., Junkins,H., McMahon,A., Milano,A., Morales,J., *et al.* (2017) The new NHGRI-EBI Catalog of published genome-wide association studies (GWAS Catalog). *Nucleic Acids Res.*, **45**, D896–D901.
4. Gibson,G. (2012) Rare and common variants: twenty arguments. *Nat. Rev. Genet.*, **13**, 135–145.
5. Lippert,C., Listgarten,J., Davidson,R.I., Baxter,S., Poon,H., Kadie,C.M. and Heckerman,D. (2013) An exhaustive epistatic SNP association analysis on expanded Wellcome Trust data. *Sci. Rep.*, **3**, 1099.
6. Manolio,T.A., Collins,F.S., Cox,N.J., Goldstein,D.B., Hindorff,L.A., Hunter,D.J., McCarthy,M.I., Ramos,E.M., Cardon,L.R., Chakravarti,A., *et al.* (2009) Finding the missing heritability of complex diseases. *Nature*, **461**, 747–753.
7. Sherry,S.T., Ward,M.H., Kholodov,M., Baker,J., Phan,L., Smigielski,E.M. and Sirotkin,K. (2001) dbSNP: the NCBI database of genetic variation. *Nucleic Acids Res.*, **29**, 308–311.
8. Blumenthal,D.B., Baumbach,J., Hoffmann,M., Kacprowski,T. and List,M. (2020) A framework for modeling epistatic interaction. *Bioinformatics*, **37**, 1708–1716.
9. Caylak,G., Tastan,O. and Cicek,A.E. (2020) Potpourri: an epistasis test prioritization algorithm via diverse SNP selection. *J. Comput. Biol.*, **28**, 365–377.
10. Cowman,T. and Koyutürk,M. (2017) Prioritizing tests of epistasis through hierarchical representation of genomic redundancies. *Nucleic Acids Res.*, **45**, e131.
11. Ayati,M. and Koyutürk,M. (2014) Prioritization of genomic locus pairs for testing epistasis. In: *Proceedings of the 5th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics, BCB '14*. Association for Computing Machinery, NY, pp. 240–248.
12. Jing,P.-J. and Shen,H.-B. (2015) MACOED: a multi-objective ant colony optimization algorithm for SNP epistasis detection in genome-wide association studies. *Bioinformatics*, **31**, 634–641.
13. Duroux,D., Climente-González,H., Azencott,C.-A. and Van Steen,K. (2022) Interpretable network-guided epistasis detection. *Gigascience*, **11**, giab093.
14. Banchi,L., Fingerhuth,M., Babej,T., Ing,C. and Arrazola,J.M. (2020) Molecular docking with Gaussian boson sampling. *Sci. Adv.*, **6**, eaax1950.
15. Boev,A., Rakitko,A., Usmanov,S., Kobzeva,A., Popov,I., Ilinsky,V., Kiktenko,E. and Fedorov,A. (2021) Genome assembly using quantum and quantum-inspired annealing. *Sci. Rep.*, **11**, 13183.
16. Nałęcz-Charkiewicz,K. and Nowak,R.M. (2022) Algorithm for DNA sequence assembly by quantum annealing. *BMC Bioinformatics*, **23**, 122.
17. Sarkar,A., Al-Ars,Z. and Bertels,K. (2021) QuASeR: Quantum Accelerated de novo DNA sequence reconstruction. *PLoS One*, **16**, e0249850.

18. Vakili,M.G., Gorgulla,C., Nigam,A., Bezrukov,D., Varoli,D., Aliper,A., Polykovsky,D., Padmanabha Das,K.M., Snider,J., Lyakisheva,A., *et al.* (2024) Quantum computing-enhanced algorithm unveils novel inhibitors for KRAS. arXiv doi: https://arxiv.org/abs/2402.08210, 13 february 2024, preprint: not peer reviewed.
19. Siek,J., Lumsdaine,A. and Lee,L.-Q. (2002) *The Boost Graph Library: User Guide and Reference Manual*. Addison-Wesley.
20. Csardi,G., Nepusz,T. and Others (2006) The igraph software package for complex network research. *InterJournal, Complex Syst.*, **1695**, 1–9.
21. Liu,F. and Chaudhary,V. (2003) A practical OpenMP compiler for system on chips. In: *OpenMP Shared Memory Parallel Programming*. Springer, Berlin, Heidelberg, pp. 54–68.
22. Marees,A.T., de Kluiver,H., Stringer,S., Vorspan,F., Curis,E., Marie-Claire,C. and Derks,E.M. (2018) A tutorial on conducting genome-wide association studies: quality control and statistical analysis. *Int. J. Methods Psychiatr. Res.*, **27**, e1608.
23. Kunkle,B.W., Grenier-Boley,B., Sims,R., Bis,J.C., Damotte,V., Naj,A.C., Boland,A., Vronskaya,M., van der Lee,S.J., Amlie-Wolf,A., *et al.* (2019) Genetic meta-analysis of diagnosed Alzheimer's disease identifies new risk loci and implicates Aβ, tau, immunity and lipid processing. *Nat. Genet.*, **51**, 414–430.
24. Reiman,E.M., Webster,J.A., Myers,A.J., Hardy,J., Dunckley,T., Zismann,V.L., Joshipura,K.D., Pearson,J.V., Hu-Lince,D., Huentelman,M.J., *et al.* (2007) GAB2 alleles modify Alzheimer's risk in APOE epsilon4 carriers. *Neuron*, **54**, 713–720.
25. Wellcome Trust Case Control Consortium, Craddock,N., Hurles,M.E., Cardin,N., Pearson,R.D., Plagnol,V., Robson,S., Vukcevic,D., Barnes,C., Conrad,D.F., *et al.* (2010) Genome-wide association study of CNVs in 16,000 cases of eight common diseases and 3,000 shared controls. *Nature*, **464**, 713–720.
26. Bycroft,C., Freeman,C., Petkova,D., Band,G., Elliott,L.T., Sharp,K., Motyer,A., Vukcevic,D., Delaneau,O., O'Connell,J., *et al.* (2018) The UK Biobank resource with deep phenotyping and genomic data. *Nature*, **562**, 203–209.
27. World Health Organization (1993) In: *The ICD-10 classification of mental and behavioural disorders: diagnostic criteria for research*, World Health Organization.
28. Best,D. and Roberts,D. (1975) Algorithm AS 89: the upper tail probabilities of Spearman's rho. *J. Roy. Stat. Soc. Ser. C (Appl. Stat.)*, **24**, 377–379.
29. Caylak,G., Tastan,O. and Cicek,A.E. (2021) Potpourri: an epistasis test prioritization algorithm via diverse SNP Selection. *J. Comput. Biol.*, **28**, 365–377.
30. Cowman,T. and Koyutürk,M. (2017) Prioritizing tests of epistasis through hierarchical representation of genomic redundancies. *Nucleic Acids Res.*, **45**, e131.
31. Guan,B., Zhao,Y. and Sun,W. (2018) Ant colony optimization with an automatic adjustment mechanism for detecting epistatic interactions. *Comput. Biol. Chem.*, **77**, 354–362.
32. Guan,B. and Zhao,Y. (2019) Self-adjusting ant colony optimization based on information entropy for detecting epistatic interactions. *Genes*, **10**, 114.
33. Schüpbach,T., Xenarios,I., Bergmann,S. and Kapur,K. (2010) FastEpistasis: a high performance computing solution for quantitative trait epistasis. *Bioinformatics*, **26**, 1468–1469.
34. Wang,Y., Liu,X., Robbins,K. and Rekaya,R. (2010) AntEpiSeeker: detecting epistatic interactions for case-control studies using a two-stage ant colony optimization algorithm. *BMC Research Notes*, **3**, 117.
35. Cao,X., Yu,G., Ren,W., Guo,M. and Wang,J. (2019) DualWMDR: Detecting epistatic interaction with dual screening and multifactor dimensionality reduction. *Hum. Mutat.*, **41**, 719–734.
36. Gola,D., John,J. M.M., van Steen,K. and König,I.R. (2015) A roadmap to multifactor dimensionality reduction methods. *Brief. Bioinform.*, **17**, 293–308.
37. Yu,W., Lee,S. and Park,T. (2016) A unified model based multifactor dimensionality reduction framework for detecting gene–gene interactions. *Bioinformatics*, **32**, i605–i610.
38. Ritchie,M.D., Hahn,L.W., Roodi,N., Bailey,L.R., Dupont,W.D., Parl,F.F. and Moore,J.H. (2001) Multifactor-dimensionality reduction reveals high-order interactions among estrogen-metabolism genes in sporadic breast cancer. *Am. J. Hum. Genet.*, **69**, 138–147.
39. Sinnott-Armstrong,N.A., Greene,C.S. and Moore,J.H. (2010) Fast genome-wide epistasis analysis using ant colony optimization for multifactor dimensionality reduction analysis on graphics processing units. In: *Proceedings of the 12th annual conference on Genetic and evolutionary computation - GECCO '10*. ACM Press.
40. Ansarifar,J. and Wang,L. (2019) New algorithms for detecting multi-effect and multi-way epistatic interactions. *Bioinformatics*, **35**, 5078–5085.
41. Wu,T.T., Chen,Y.F., Hastie,T., Sobel,E. and Lange,K. (2009) Genome-wide association analysis by lasso penalized logistic regression. *Bioinformatics*, **25**, 714–721.
42. North,B.V., Curtis,D. and Sham,P.C. (2005) Application of logistic regression to case-control association studies involving two causative loci. *Hum. Hered.*, **59**, 79–87.
43. Oughtred,R., Stark,C., Breitkreutz,B.-J., Rust,J., Boucher,L., Chang,C., Kolas,N., O'Donnell,L., Leung,G., McAdam,R., *et al.* (2019) The BioGRID interaction database: 2019 update. *Nucleic Acids Res.*, **47**, D529–D541.
44. Kerimov,N., Hayhurst,J.D., Peikova,K., Manning,J.R., Walter,P., Kolberg,L., Samoviča,M., Sakthivel,M.P., Kuzmin,I., Trevanion,S.J., *et al.* (2021) A compendium of uniformly processed human gene expression and splicing quantitative trait loci. *Nat. Genet.*, **53**, 1290–1299.
45. Durinck,S., Spellman,P.T., Birney,E. and Huber,W. (2009) Mapping identifiers for the integration of genomic datasets with the R/Bioconductor package biomaRt. *Nat. Protoc.*, **4**, 1184–1191.
46. Benjamini,Y. and Hochberg,Y. (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc.*.**125**, 289–300.
47. Kirkpatrick,S., Gelatt,C.D. and Vecchi,M.P. (1983) Optimization by simulated annealing. *Science*, **220**, 671–680.
48. Riesen,K., Fischer,A. and Bunke,H. (2017) Improved graph edit distance approximation with simulated annealing. In: Foggia,P., Liu,C.-L. and Vento,M. (eds.) *GbRPR 2017, Vol. 10310 of LNCS*. Springer, Cham, pp. 222–231.
49. Blumenthal,D.B., Bougleux,S., Gamper,J. and Brun,L. (2019) GEDLIB: A C++ library for graph edit distance computation. In: Conte,D., Ramel,J.-Y. and Foggia,P. (eds.) *GbRPR 2019, Vol. 11510 of LNCS*. Springer, Cham, pp. 14–24.
50. Blumenthal,D.B., Boria,N., Gamper,J., Bougleux,S. and Brun,L. (2020) Comparing heuristics for graph edit distance computation. *VLDB J.*, **29**, 419–458.
51. Dorigo,M., Birattari,M. and Stutzle,T. (2006) Ant colony optimization. *IEEE Comput. Intell. Mag.*, **1**, 28–39.
52. Koza,J.R. and Poli,R. (2005) *Genetic Programming*. Springer, US, Boston, MA, pp. 127–164.
53. Duarte,A., Sánchez-Oro,J., Mladenović,N. and Todosijević,R. (2018) *Variable Neighborhood Descent*. *Springer International Publishing*, **Cham**, pp. 341–367.
54. Boria,N., Blumenthal,D.B., Bougleux,S. and Brun,L. (2019) Improved local search for graph edit distance. *Pattern Recognit. Lett.*, **129**, 19–25.
55. Lazareva,O., Baumbach,J., List,M. and Blumenthal,D.B. (2021) On the limits of active module identification. *Brief. Bioinform.*, **22**, bbab066.
56. Subramanian,A., Tamayo,P., Mootha,V.K., Mukherjee,S., Ebert,B.L., Gillette,M.A., Paulovich,A., Pomeroy,S.L., Golub,T.R., Lander,E.S., *et al.* (2005) Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proc. Natl. Acad. Sci. U.S.A.*, **102**, 15545–15550.

57. Mootha,V.K., Lindgren,C.M., Eriksson,K.-F., Subramanian,A., Sihag,S., Lehar,J., Puigserver,P., Carlsson,E., Ridderstråle,M., Laurila,E., *et al.* (2003) PGC-1alpha-responsive genes involved in oxidative phosphorylation are coordinately downregulated in human diabetes. *Nat. Genet.*, **34**, 267–273.
58. Grover,L.K. (1997) Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.*, **79**, 325.
59. Roland,J. and Cerf,N.J. (2002) Quantum search by local adiabatic evolution. *Phys. Rev. A*, **65**, 042308.
60. Pirnay,N., Ulitzsch,V., Wilde,F., Eisert,J. and Seifert,J.-P. (2024) An in-principle super-polynomial quantum advantage for approximating combinatorial optimization problems via computational learning theory, *Sci. Adv.*, **10**, eadj5170.
61. Aaronson,S. (2022) How Much structure is needed for huge quantum speedups? arXiv doi: https://arxiv.org/abs/2209.06930, 14 September 2022, preprint: not peer reviewed.
62. King,A.D., Raymond,J., Lanting,T., Harris,R., Zucca,A., Altomare,F., Berkley,A.J., Boothby,K., Ejtemaee,S., Enderud,C., *et al.* (2023) Quantum critical dynamics in a 5,000-qubit programmable spin glass. *Nature*, **617**, 16–66.
63. Incudini,M., Tarocco,F., Mengoni,R., Di Pierro,A. and Mandarino,A. (2022) Computing graph edit distance on quantum devices. *Quant. Mach. Intell.*, **4**, 24.
64. Kirkpatrick,S., Gelatt,C.D. and Vecchi,M.P. (1983) Optimization by simulated annealing. *science*, **220**, 671–680.
65. Earl,D.J. and Deem,M.W. (2005) Parallel tempering: theory, applications, and new perspectives. *Phys. Chem. Chem. Phys.*, **7**, 3910–3916.
66. Gurobi Optimization, LLC. (2024) Gurobi optimizer reference manual. https://www.gurobi.com/documentation/current/refman/index.html. version 11.
67. Traag,V.A., Waltman,L. and van Eck,N.J. (2019) From Louvain to Leiden: guaranteeing well-connected communities. *Sci. Rep.*, **9**, 5233.
68. Kanehisa,M. (2002) The KEGG database. *Novartis Found. Symp.*, **247**, 91–101.
69. Guo,X., Qiu,W., Garcia-Milian,R., Lin,X., Zhang,Y., Cao,Y., Tan,Y., Wang,Z., Shi,J., Wang,J., *et al.* (2017) Genome-wide significant, replicated and functional risk variants for Alzheimer's disease. *J. Neural. Transm.*, **124**, 1455–1471.
70. Kulminski,A.M., Shu,L., Loika,Y., He,L., Nazarian,A., Arbeev,K., Ukraintseva,S., Yashin,A. and Culminskaya,I. (2020) Genetic and regulatory architecture of Alzheimer's disease in the APOE region. *Alzheimers. Dement.*, **12**, e12008.
71. Thul,P.J. and Lindskog,C. (2018) The human protein atlas: a spatial map of the human proteome. *Protein Sci.*, **27**, 233–244.
72. Braak,H. and Braak,E. (1991) Neuropathological stageing of Alzheimer-related changes. *Acta Neuropathol.*, **82**, 239–259.
73. Braak,H., Thal,D.R., Ghebremedhin,E. and Del Tredici,K. (2011) Stages of the pathologic process in Alzheimer disease: age categories from 1 to 100 years. *J. Neuropathol. Exp. Neurol.*, **70**, 960–969.
74. Thal,D.R., Rüb,U., Orantes,M. and Braak,H. (2002) Phases of A beta-deposition in the human brain and its relevance for the development of AD. *Neurology*, **58**, 1791–1800.
75. Filip,P., Canna,A., Moheet,A., Bednarik,P., Grohn,H., Li,X., Kumar,A.F., Olawsky,E., Eberly,L.E., Seaquist,E.R., *et al.* (2020) Structural alterations in deep brain structures in type 1 diabetes. *Diabetes*, **69**, 2458–2466.
76. Knapp,M., Tu,X. and Wu,R. (2019) Vascular endothelial dysfunction, a major mediator in diabetic cardiomyopathy. *Acta Pharmacol. Sin.*, **40**, 1–8.
77. Schuster,D.P. and Duvuuri,V. (2002) Diabetes mellitus. *Clin. Podiatr. Med. Surg.*, **19**, 79–107.
78. Eizirik,D.L., Pasquali,L. and Cnop,M. (2020) Pancreatic β-cells in type 1 and type 2 diabetes mellitus: different pathways to failure. *Nat. Rev. Endocrinol.*, **16**, 349–362.
79. Gillespie,K.M. (2006) Type 1 diabetes: pathogenesis and prevention. *CMAJ*, **175**, 165–170.
80. Granlund,L., Hedin,A., Korsgren,O., Skog,O. and Lundberg,M. (2022) Altered microvasculature in pancreatic islets from subjects with type 1 diabetes. *PLoS One*, **17**, e0276942.
81. Stefański,A., Wolf,J., Harazny,J.M., Miszkowska-Nagórna,E., Wolnik,B., Murawska,J., Narkiewicz,K. and Schmieder,R.E. (2023) Impact of type 1 diabetes and its duration on wall-to-lumen ratio and blood flow in retinal arterioles. *Microvasc. Res.*, **147**, 104499.
82. Kiseleva,E., Ryabkov,M., Baleev,M., Bederina,E., Shilyagin,P., Moiseev,A., Beschastnov,V., Romanov,I., Gelikonov,G. and Gladkova,N. (2021) Prospects of intraoperative multimodal OCT application in patients with acute mesenteric ischemia. *Diagnostics (Basel)*, **11**, 705.
83. Ahmed,M. (2021) Ischemic bowel disease in 2021. *World J. Gastroenterol.*, **27**, 4746–4762.
84. Green,B.T. and Tendler,D.A. (2005) Ischemic colitis: a clinical review. *South. Med. J.*, **98**, 217–222.
85. 1000 Genomes Project Consortium, Auton,A., Brooks,L.D., Durbin,R.M., Garrison,E.P., Kang,H.M., Korbel,J.O., Marchini,J.L., McCarthy,S., McVean,G.A., *et al.* (2015) A global reference for human genetic variation. *Nature*, **526**, 68–74.
86. Chapuis,G., Djidjev,H., Hahn,G. and Rizk,G. (2019) Finding maximum cliques on the D-wave quantum annealer. *J. Signal Process. Syst.*, **91**, 363–377.
87. Maier,A., Hartung,M., Abovsky,M., Adamowicz,K., Bader,G.D., Baier,S., Blumenthal,D.B., Chen,J., Elkjaer,M.L., Garcia-Hernandez,C., *et al.* (2024) Drugst.One — a plug-and-play solution for online systems medicine and network-based drug repurposing. *Nucleic Acids Res.*, **52**, W481–W488.
88. Zuk,O., Hechter,E., Sunyaev,S.R. and Lander,E.S. (2012) The mystery of missing heritability: genetic interactions create phantom heritability. *Proc. Natl. Acad. Sci.*, **109**, 1193–1198.
89. Louadi,Z., Yuan,K., Gress,A., Tsoy,O., Kalinina,O.V., Baumbach,J., Kacprowski,T. and List,M. (2021) DIGGER: exploring the functional role of alternative splicing in protein interactions. *Nucleic Acids Res.*, **49**, D309–D318.
90. Hernández-Lorenzo,L., Hoffmann,M., Scheibling,E., List,M., Matías-Guiu,J.A. and Ayala,J.L. (2022) On the limits of graph neural networks for the early diagnosis of Alzheimer's disease. *Sci. Rep.*, **12**, 17632.
91. Cortazzo,P., Cerveñansky,C., Marín,M., Reiss,C., Ehrlich,R. and Deana,A. (2002) Silent mutations affect in vivo protein folding in Escherichia coli. *Biochem. Biophys. Res. Commun.*, **293**, 537–541.
92. Watanabe,K., Stringer,S., Frei,O., Umićević Mirkov,M., de Leeuw,C., Polderman,T.J., van der Sluis,S., Andreassen,O.A., Neale,B.M. and Posthuma,D. (2019) A global overview of pleiotropy and genetic architecture in complex traits. *Nat. Genet.*, **51**, 1339–1348.
93. Tate,R., Moondra,J., Gard,B., Mohler,G. and Gupta,S. (2023) Warm-started QAOA with custom mixers provably converges and computationally beats Goemans-Williamson's Max-Cut at low circuit depths. *Quantum*, **7**, 1121.

Received: November 22, 2023. **Revised:** July 12, 2024. **Editorial Decision:** July 15, 2024. **Accepted:** August 1, 2024