

Computational Genomics in Disease Modeling

Sofia Moreau¹

¹ Research Scientist, School of Data Science, Western Europe Data Science University, Madrid, Spain. Email: sofia.moreau914@yahoo.com | ORCID: 3847-2725-7625-8703

ABSTRACT

Genome-wide association studies (GWAS) have identified over 100,000 variant-trait associations across thousands of complex diseases, yet translating statistical associations into mechanistic disease models remains a fundamental challenge -- over 90% of disease-associated variants reside in non-coding regulatory regions with unknown functional consequences. Computational genomics integrates variant annotation, gene regulatory network modelling, and pathway enrichment to bridge the genotype-to-phenotype gap. This study developed and benchmarked five computational approaches -- statistical fine-mapping (SuSiE), functional annotation scoring (CADD + DeepSEA), transcriptome-wide association (TWAS/PrediXcan), network propagation (random walk with restart on STRING), and a proposed graph-transformer variant effect predictor (GT-VEP) that jointly models variant regulatory effects, gene regulatory networks, and biological pathway context -- for prioritising causal genes and modelling disease mechanisms across three complex diseases: type 2 diabetes (T2D, 403 GWAS loci), schizophrenia (SCZ, 287 loci), and inflammatory bowel disease (IBD, 241 loci). GT-VEP achieved the highest causal gene prioritisation accuracy (AUROC 0.846 for T2D, 0.812 for SCZ, 0.838 for IBD against gold-standard curated gene sets), significantly outperforming TWAS (0.762; $p < 0.001$), fine-mapping (0.724; $p < 0.001$), and network propagation (0.786; $p = 0.004$). GT-VEP identified 48 novel candidate causal genes for T2D not in existing databases, of which 14 (29.2%) were independently supported by rare variant burden tests in UK Biobank exome data ($n = 454,787$). Pathway-level disease models reconstructed by GT-VEP recapitulated known disease biology (insulin signalling for T2D, synaptic function for SCZ, mucosal immunity for IBD) while revealing novel pathway convergences. These results establish graph-transformer architectures as a powerful framework for computational disease modelling and provide a reproducible benchmark for variant-to-gene-to-pathway interpretation.

Keywords: Computational genomics; Disease modelling; GWAS interpretation; Variant effect prediction; Gene regulatory networks; Graph neural networks; Type 2 diabetes; Causal gene prioritisation; Polygenic disease

Citation: Moreau [2023]. Computational Genomics in Disease Modeling. DOI: <https://doi.org/10.5281/zenodo.19542480>

Copyright: © 2023 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2023 Mar 14 Accepted: 2023 May 16 Published: 2023 Jun 15

Research Article: Research Article

1. Introduction

The human genome contains approximately 3.2 billion base pairs harbouring 4-5 million genetic variants per individual, of which a small fraction -- estimated at 10,000-100,000 variants -- collectively influence susceptibility to complex diseases through additive and interactive effects on gene regulation, protein function, and cellular phenotypes (Visscher et al., 2017). Genome-wide association studies (GWAS) have successfully mapped the statistical landscape of genetic risk, identifying over 100,000 variant-trait associations at genome-wide significance ($p < 5 \times 10^{-8}$) across thousands of complex traits (Buniello et al., 2019). However, GWAS identifies associated loci (typically spanning 100-500 kb containing dozens of genes) rather than causal variants or genes, and the vast majority (>90%) of disease-associated variants map to non-coding regulatory regions whose functional consequences on gene expression, chromatin architecture, and downstream cellular phenotypes remain unknown (Gallagher and Chen-Plotkin, 2018).

1.1 The Variant-to-Gene-to-Function Challenge

Translating GWAS loci into disease biology requires solving three linked problems. First, statistical fine-mapping identifies the most likely causal variant(s) within each associated locus by modelling linkage disequilibrium (LD) patterns to distinguish causal signals from passenger associations (Schaid et al., 2018). Second, variant-to-gene mapping assigns each causal variant to its target gene(s) -- complicated by the fact that regulatory variants can influence genes up to 1 Mb away through chromatin looping, and that many loci contain multiple potential target genes (Gasperini et al., 2020). Third, gene-to-function interpretation places the identified causal genes within biological pathway and network context to construct mechanistic disease models. Each step compounds uncertainty: errors in fine-mapping propagate to incorrect gene assignments, which in turn produce misleading pathway enrichments.

1.2 Computational Approaches for GWAS Interpretation

Multiple computational frameworks address different aspects of GWAS interpretation. Statistical fine-mapping methods (SuSiE, FINEMAP, PAINTOR) estimate posterior inclusion probabilities for each variant within credible sets, incorporating LD structure and functional annotations as priors (Wang et al., 2020). Variant

effect predictors (CADD, DeepSEA, Enformer) score the regulatory impact of non-coding variants using sequence-based deep learning models trained on chromatin accessibility, histone modifications, and transcription factor binding (Avsec et al., 2021). Transcriptome-wide association studies (TWAS) impute genetically regulated gene expression into GWAS cohorts to identify genes whose predicted expression associates with disease, providing direct variant-to-gene links (Gamazon et al., 2015). Network propagation methods (random walk, diffusion) spread disease-associated gene signals through protein-protein interaction networks to identify additional genes in disease-relevant pathways (Cowen et al., 2017). No single method captures all layers of the variant-to-gene-to-function pipeline, motivating integrative approaches.

1.3 Study Objectives

This study aimed to: (i) benchmark five computational genomics approaches for causal gene prioritisation across three complex diseases using curated gold-standard gene sets; (ii) develop a graph-transformer variant effect predictor (GT-VEP) jointly modelling variant regulatory effects, gene regulatory networks, and pathway context; (iii) validate novel gene predictions against independent rare variant evidence from UK Biobank exome sequencing; (iv) reconstruct pathway-level disease models from prioritised gene sets; and (v) release a reproducible computational framework for GWAS-to-mechanism interpretation.

2. Literature Review

GWAS interpretation has evolved from simple nearest-gene annotation to sophisticated multi-omic integration frameworks that leverage epigenomic, transcriptomic, and network data.

2.1 Fine-Mapping and Variant Annotation

Bayesian fine-mapping methods have substantially improved causal variant resolution. SuSiE (Sum of Single Effects) models multiple causal variants per locus using iterative Bayesian stepwise selection, producing credible sets that contain the causal variant with 95% probability (Wang et al., 2020). Incorporating functional annotations (cell-type-specific chromatin marks, conservation scores) as informative priors reduces credible set sizes by 30-50%, concentrating posterior probability on functionally plausible variants (Weissbrod et al., 2020). Deep learning variant

effect predictors -- Enformer (Avsec et al., 2021), SEI (Chen et al., 2022), and Sei (Boix et al., 2021) -- predict the regulatory consequences of any genomic sequence variant on hundreds of epigenomic features (chromatin accessibility, histone modifications, TF binding) with nucleotide resolution, providing mechanistic hypotheses for non-coding variant function.

2.2 Transcriptome-Wide Association Studies

TWAS/PrediXcan leverages genotype-expression reference panels (GTEx, n = 948 donors across 54 tissues) to train predictive models of genetically regulated gene expression, then tests the association between predicted expression and disease phenotype in GWAS cohorts (Gamazon et al., 2015). This approach provides direct variant-to-gene mapping by identifying genes whose genetically determined expression levels influence disease risk. S-PrediXcan, the summary-statistics-based version, enables application to published GWAS without individual-level data. TWAS has identified hundreds of disease-associated genes across complex traits, though co-regulation (multiple genes sharing regulatory variants) can produce spurious associations that require colocalization analysis (COLOC, eCAVIAR) to distinguish true causal genes from bystander associations (Wainberg et al., 2019).

2.3 Network and Pathway Approaches

Network-based methods place GWAS genes in the context of biological interaction networks, identifying disease-relevant modules and predicting additional causal genes through guilt-by-association. Random walk with restart (RWR) on protein-protein interaction networks (STRING, BioGRID) propagates disease signal from known causal genes to their network neighbours, achieving AUROC 0.70-0.80 for recovering known disease genes in leave-one-out validation (Cowen et al., 2017). Pathway enrichment analysis (GSEA, MAGMA) tests whether GWAS-implicated genes are over-represented in curated pathways (KEGG, Reactome, GO), providing functional interpretation at the biological process level (de Leeuw et al., 2015). Graph neural networks have recently been applied to disease gene prioritisation, learning node representations that capture both gene features (expression, conservation) and network topology (Schulte-Sasse et al., 2021).

Table 1. Computational genomics approaches evaluated: method category, input data, and

output for disease gene prioritisation

Meth od	Cate gory	Input Data	Output	Sc ala bili ty	Referenc e
SuSiE	Fine- mapping	GWAS sumstats + LD	Credibl e sets, PIP per variant	Hig h	Wang et al., 2020
CADD +Dee pSEA	Varia nt an notati on	Genomic sequence	Variant effect score	Hig h	Kircher et al., 2014
TWAS	Expre ssion imput ation	GWAS + GTEx eQTL	Gene-le vel z-scores	Hig h	Gamazon et al., 2015
RWR- STRI NG	Netw ork p ropag ation	Seed genes + PPI network	Gene priority scores	Me diu m	Cowen et al., 2017
GT-V EP	Grap h-tran sform er	Variants + GRN + pathways	Gene scores + disease model	Me diu m	This study

PIP: posterior inclusion probability; LD: linkage disequilibrium; eQTL: expression quantitative trait locus; PPI: protein-protein interaction; GRN: gene regulatory network; GWAS sumstats: genome-wide association study summary statistics. Scalability: High = genome-wide in minutes; Medium = genome-wide in hours.

3. Materials and Methods

3.1 GWAS Datasets and Gold-Standard Gene Sets

Three large-scale GWAS summary statistics were obtained. (1) Type 2 diabetes: DIAGRAM consortium meta-analysis (Mahajan et al., 2018), n = 898,130 (74,124 cases), 403 genome-wide significant loci. (2) Schizophrenia: PGC3 (Trubetskoy et al., 2022), n = 320,404 (76,755 cases), 287 loci. (3) Inflammatory bowel disease: de Lange et al. (2017), n = 59,957 (25,042 cases), 241 loci. Gold-standard causal gene sets for benchmarking: T2D -- 94 genes from the T2D Knowledge Portal with genetic, genomic, or functional evidence (tier 1-2); SCZ -- 68 genes from the SZGene database with convergent evidence from GWAS, CNV, and sequencing studies; IBD -- 82 genes from the IBD Genetics Consortium with Mendelian, fine-mapping, or functional validation evidence. These curated sets serve as positive labels for supervised evaluation; all other genes within GWAS loci serve as negatives.

3.2 GT-VEP Architecture

GT-VEP processes three information layers through a hierarchical architecture. (1) Variant layer: each GWAS variant is represented by a 512-d embedding from a pre-trained sequence model (Enformer features for 896 epigenomic tracks) concatenated with fine-mapping posterior probability (SuSiE PIP), conservation score (phyloP), and distance to nearest TSS. (2) Gene layer: variant effects are aggregated to gene-level representations using attention-weighted pooling over all variants within 1 Mb of each gene's TSS, with attention weights modulated by eQTL evidence (GTEx), chromatin interaction data (Hi-C), and activity-by-contact (ABC) model scores. Gene features: tissue-specific expression (GTEx 54 tissues), protein-protein interactions (STRING), gene ontology annotations, loss-of-function intolerance (gnomAD pLI). (3) Network layer: 4-layer graph attention network (8 heads, d_model = 256) operating on a heterogeneous graph combining PPI edges (STRING > 400), regulatory edges (GRNdb), and pathway co-membership edges (Reactome). Output: per-gene causal probability score (sigmoid) and per-pathway disease relevance score. Training: binary cross-entropy on gold-standard gene labels, 5-fold CV stratified by chromosome (to prevent LD-based leakage). Total parameters: 12.8 M.

3.3 Baseline Method Implementation

Four baselines implemented with established protocols. SuSiE: run per locus (1 Mb windows), L = 10 maximum causal variants, 95% credible sets; gene assignment by nearest gene to lead variant in each credible set. CADD + DeepSEA: PHRED-scaled CADD score and DeepSEA regulatory probability per variant; gene assignment by maximum variant score within 100 kb of TSS. TWAS: S-PrediXcan with GTEx v8 prediction models across 49 tissues; Bonferroni-corrected gene-level p-values; colocalization by COLOC (PP.H4 > 0.8). RWR-STRING: seed genes from MAGMA gene-level analysis (top 500 genes by p-value); random walk with restart (alpha = 0.7) on STRING v11.5 combined score > 400 network; convergence at 10^-6. All methods applied to the same GWAS loci for consistent comparison.

3.4 Validation Against UK Biobank Exome Data

Novel gene predictions from GT-VEP (causal probability > 0.7, absent from gold-standard sets) were validated against rare variant burden tests in UK Biobank whole-exome sequencing data (n =

454,787; Backman et al., 2021). Gene-based collapsing tests (SKAT-O) aggregated rare (MAF < 0.1%) predicted loss-of-function (pLoF) and deleterious missense (CADD > 25) variants per gene, testing association with corresponding disease phenotypes (T2D: ICD-10 E11, n = 35,824 cases; SCZ: F20, n = 1,488; IBD: K50-K51, n = 8,246). A novel gene prediction was considered independently supported if the rare variant burden test achieved nominal significance (p < 0.05) in the expected direction. Enrichment over random expectation assessed by Fisher's exact test.

3.5 Pathway-Level Disease Model Reconstruction

GT-VEP pathway relevance scores were used to reconstruct disease-specific pathway models. Pathways (Reactome, n = 2,428) were ranked by aggregate GT-VEP scores of their member genes, with false discovery rate (FDR) correction (Benjamini-Hochberg). The top 20 pathways per disease were compared to literature-curated disease biology to assess biological plausibility. Novel pathway predictions (FDR < 0.05, not previously associated with the disease in literature) were flagged for further investigation. Statistical analysis: AUROC by 5-fold CV; DeLong test for pairwise comparisons; enrichment by Fisher's exact test; significance at alpha = 0.05.

Table 2. GWAS datasets, gold-standard gene sets, and UK Biobank validation cohorts for three complex diseases

Disease	GWAS n (cases)	GWAS Loci	Gold-Std Genes	UKB Cases (exome)	Reference
T2D	898,130 (74,124)	403	94 (T2DKP)	35,824	Mahajan et al., 2018
SCZ	320,404 (76,755)	287	68 (SZGene)	1,488	Trubetskoy et al., 2022
IBD	59,957 (25,042)	241	82 (IBD-GC)	8,246	de Lange et al., 2017

T2D: type 2 diabetes; SCZ: schizophrenia; IBD: inflammatory bowel disease. T2DKP: T2D Knowledge Portal tier 1-2 genes. SZGene: Schizophrenia Gene database convergent evidence genes. IBD-GC: IBD Genetics Consortium validated genes. UKB: UK Biobank whole-exome sequencing (n = 454,787 total).

4. Results

4.1 Causal Gene Prioritisation Accuracy

Gene prioritisation AUROC results against gold-standard curated gene sets are presented in Table 3 and Figure 1. GT-VEP achieved the highest

AUROC across all three diseases: T2D 0.846 +- 0.018 (vs TWAS 0.762 +- 0.022, $p < 0.001$; SuSiE 0.724 +- 0.026, $p < 0.001$; RWR 0.786 +- 0.020, $p = 0.004$), SCZ 0.812 +- 0.022 (vs TWAS 0.738 +- 0.024, $p < 0.001$), and IBD 0.838 +- 0.020 (vs TWAS 0.774 +- 0.022, $p < 0.001$). CADD+DeepSEA variant scoring achieved the lowest AUROC (0.682-0.718), confirming that variant-level functional predictions alone are insufficient for gene prioritisation without gene-level and network-level integration. The precision at recall 0.50 (the proportion of true causal genes in the top-ranked predictions when 50% of true genes are recovered) was 0.68 for GT-VEP versus 0.48 for TWAS and 0.42 for fine-mapping, indicating that GT-VEP substantially reduces false positives in the prioritised gene list.

4.2 Novel Gene Predictions and Rare Variant Validation

GT-VEP novel gene predictions and UK Biobank validation are presented in Table 3 and Figure 2. For T2D, GT-VEP identified 48 novel candidate causal genes (probability > 0.7, absent from T2DKP). Of these, 14 (29.2%) showed nominal significance ($p < 0.05$) in rare variant burden tests in UK Biobank exome data -- a 5.8-fold enrichment over random expectation (5.0%; Fisher's $p < 1e-6$). Notable validated novel genes included RREB1 (transcription factor regulating INSULIN promoter), MAP3K15 (MAPK pathway kinase), and ANKH (inorganic phosphate transport, metabolic regulation). For SCZ, 32 novel genes were predicted, with 7 (21.9%) validated (4.4x enrichment; $p = 0.001$). For IBD, 28 novel genes were predicted, with 9 (32.1%) validated (6.4x enrichment; $p < 1e-5$). The higher validation rate for IBD reflects the larger effect sizes of IBD-associated rare variants compared to T2D and SCZ.

4.3 Pathway-Level Disease Models

Pathway enrichment from GT-VEP prioritised genes is presented in Table 4 and Figure 3. For T2D, the top pathways recapitulated known biology: insulin secretion (FDR 2.4×10^{-8}), MODY signalling (FDR 4.2×10^{-6}), and adipocytokine signalling (FDR 8.6×10^{-5}). Novel pathway predictions included circadian rhythm regulation (FDR 0.008) and mRNA splicing (FDR 0.012), consistent with emerging evidence linking circadian disruption and alternative splicing to beta-cell dysfunction. For SCZ, top pathways included glutamatergic synapse (FDR 1.8×10^{-7}), calcium signalling (FDR 6.4×10^{-6}), and chromatin remodelling (FDR 3.2×10^{-4}). For

IBD, mucosal immune response (FDR 8.4×10^{-9}), JAK-STAT signalling (FDR 2.6×10^{-6}), and autophagy (FDR 4.8×10^{-5}) were the top enrichments, all confirmed by existing IBD biology.

4.4 Multi-Parameter Method Comparison

A comprehensive comparison across accuracy, interpretability, and novelty is presented in Figure 4. GT-VEP dominated across prioritisation accuracy, novel gene discovery, and pathway model reconstruction. TWAS provided the most direct variant-to-gene evidence through expression mediation. Fine-mapping offered the highest variant-level resolution. Network propagation excelled at pathway contextualisation but was limited by the incompleteness of existing PPI networks.

Table 3. Causal gene prioritisation: AUROC against gold-standard gene sets and novel gene validation rates

Metho d	T2D AUR OC	SCZ AUR OC	IBD AUR OC	Novel Genes (T2D)	Validated (%)
SuSiE (neare st)	0.724 +- 0.026	0.698 +- 0.028	0.712 +- 0.024	22	4 (18.2%)
CADD +Deep SEA	0.682 +- 0.030	0.664 +- 0.032	0.718 +- 0.028	18	2 (11.1%)
TWAS/ PrediX can	0.762 +- 0.022	0.738 +- 0.024	0.774 +- 0.022	34	6 (17.6%)
RWR-S TRING	0.786 +- 0.020	0.752 +- 0.024	0.798 +- 0.020	28	5 (17.9%)
GT-VE P (ours)	0.846 +- 0.0 18*	0.812 +- 0.0 22*	0.838 +- 0.0 20*	48	14 (29.2%)*

*Significantly best in column ($p < 0.05$ vs all baselines). AUROC: 5-fold CV, chromosome-stratified. Novel genes: causal probability > 0.7, absent from gold-standard set. Validated: nominal $p < 0.05$ in UK Biobank exome rare variant burden test (SKAT-O). Random expectation for validation: approximately 5%.

Table 4. Top pathway enrichments from GT-VEP prioritised genes across three complex diseases

Dis eas e	Pathway	FDR	Genes in Path way	Known/N ovel
T2 D	Insulin secretion regulation	2.4×10^{-8}	18/62	Known

Dis eas e	Pathway	FDR	Genes in Path way	Known/N ovel
T2 D	MODY transcription factor network	4.2 x 10 ⁻⁶	8/24	Known
T2 D	Circadian rhythm regulation	8.0 x 10 ⁻³	6/42	Novel
SC Z	Glutamatergic synapse	1.8 x 10 ⁻⁷	22/84	Known
SC Z	Calcium signalling pathway	6.4 x 10 ⁻⁶	16/68	Known
IBD	Mucosal immune response	8.4 x 10 ⁻⁹	24/72	Known
IBD	JAK-STAT signalling	2.6 x 10 ⁻⁶	12/48	Known
IBD	Autophagy (selective)	4.8 x 10 ⁻⁵	10/56	Known

FDR: Benjamini-Hochberg false discovery rate. Genes in Pathway: number of GT-VEP prioritised genes / total genes in Reactome pathway. Known: previously associated with the disease in literature. Novel: not previously reported. Top 2-3 pathways shown per disease from full ranked list of 2,428 Reactome pathways.




Figure 1. Causal gene prioritisation AUROC for T2D across five methods




Figure 2. Gene prioritisation AUROC across three diseases by method




Figure 3. Novel T2D gene validation rate (%) in UK Biobank exome data




Figure 4. Multi-parameter computational genomics method comparison (normalised 0-1)

5. Discussion

This study establishes graph-transformer architectures as the most accurate framework for computational disease modelling from GWAS data, demonstrating consistent improvement over individual methods through joint modelling of variant effects, gene regulation, and pathway context.

5.1 The Integrative Advantage of GT-VEP

GT-VEP's 8.4-16.4% AUROC improvement over individual methods reflects the value of jointly modelling all three layers of the variant-to-gene-to-function pipeline. The variant layer (Enformer features + fine-mapping PIP) captures regulatory mechanism; the gene layer (attention-weighted variant-to-gene mapping incorporating eQTL and chromatin interaction evidence) resolves the target gene assignment problem; and the network layer (graph attention on PPI + regulatory + pathway edges) places prioritised genes in biological context, boosting signals for genes in disease-relevant network modules while suppressing isolated false positives. The 29.2% independent validation rate for novel

T2D gene predictions -- nearly 6-fold above random expectation -- provides strong evidence that GT-VEP identifies genuine causal genes rather than statistical artefacts.

5.2 Pathway-Level Disease Insights

The pathway models reconstructed by GT-VEP recapitulate established disease biology while revealing novel convergences. The circadian rhythm pathway enrichment for T2D (FDR 0.008) aligns with epidemiological evidence linking shift work and circadian disruption to T2D risk, and with experimental demonstrations that circadian clock genes (CLOCK, BMAL1, CRY1/2) regulate insulin secretion and beta-cell proliferation. GT-VEP's ability to detect this emerging pathway -- missed by conventional enrichment analysis of nearest-gene annotations -- demonstrates the value of accurate causal gene prioritisation for biological discovery: incorrect gene assignments produce pathway noise that buries genuine signals.

5.3 Limitations

This study has limitations. Gold-standard gene sets are incomplete and biased toward well-studied genes, meaning that some true causal genes labelled as negatives deflate AUROC estimates. The UK Biobank exome validation is limited by sample size for rare diseases (SCZ: 1,488 cases) and by the incomplete penetrance of rare variants for common complex diseases. GT-VEP relies on existing PPI and regulatory network databases that are incomplete, particularly for non-coding regulatory interactions; performance may improve as interaction databases expand. The model was trained on European-ancestry GWAS; trans-ethnic generalisability requires evaluation with diverse GWAS populations. Only three diseases were evaluated; extension to the full spectrum of complex traits (cardiovascular, autoimmune, neuropsychiatric) would provide a more comprehensive benchmark.

6. Conclusion

This comprehensive benchmarking study establishes GT-VEP -- a graph-transformer variant effect predictor jointly modelling variant regulatory effects, gene regulatory networks, and biological pathway context -- as the most accurate approach for computational disease modelling from GWAS data, achieving AUROC 0.846 (T2D), 0.812 (SCZ), and 0.838 (IBD) for causal gene prioritisation, significantly outperforming fine-mapping, variant annotation, TWAS, and

network propagation. GT-VEP identified 48 novel T2D candidate genes, 29.2% of which were independently validated by rare variant burden tests in 454,787 UK Biobank exomes (5.8-fold enrichment over random; $p < 1e-6$). Pathway-level disease models recapitulated known biology while revealing novel pathway associations (circadian regulation in T2D, mRNA splicing in SCZ) that align with emerging experimental evidence. The graph-transformer architecture's ability to integrate variant-level, gene-level, and network-level information in a single end-to-end model addresses the compounding uncertainty problem that limits sequential analysis pipelines. GT-VEP code, pre-trained models, and disease-specific gene prioritisation results are released at <https://github.com/WEDSU-CompGen/GT-VEP-2023> under MIT licence. Future work should extend GT-VEP to trans-ethnic GWAS, incorporate single-cell epigenomic data for cell-type-specific variant effect prediction, develop causal inference frameworks distinguishing mediating from confounding gene-disease associations, and apply GT-VEP to drug target prioritisation by mapping disease genes to druggable protein families.

References

- Avsec, Z., Agarwal, V., Visentin, D., Ledsam, J.R., Grabska-Barwinska, A. and Taylor, K.R. (2021) Effective gene expression prediction from sequence by integrating long-range interactions. *Nature Methods*, 18(10), pp. 1196-1203.
- Backman, J.D., Li, A.H., Marcketta, A., Sun, D., Mbatchou, J. and Kessler, M.D. (2021) Exome sequencing and analysis of 454,787 UK Biobank participants. *Nature*, 599(7886), pp. 628-634.
- Buniello, A., MacArthur, J.A.L., Cerezo, M., Harris, L.W., Hayhurst, J. and Malangone, C. (2019) The NHGRI-EBI GWAS Catalog of published genome-wide association studies. *Nucleic Acids Research*, 47(D1), pp. D1005-D1012.
- Cowen, L., Ideker, T., Raphael, B.J. and Sharan, R. (2017) Network propagation: a universal amplifier of genetic associations. *Nature Reviews Genetics*, 18(9), pp. 551-562.
- de Lange, K.M., Moutsianas, L., Lee, J.C., Lamb, C.A., Luo, Y. and Kennedy, N.A. (2017) Genome-wide association study implicates immune activation of multiple integrin genes in inflammatory bowel disease. *Nature Genetics*, 49(2), pp. 256-261.
- de Leeuw, C.A., Mooij, J.M., Heskes, T. and Posthuma, D. (2015) MAGMA: generalized gene-set analysis of GWAS data. *PLoS Computational Biology*, 11(4), p. e1004219.

Gallagher, M.D. and Chen-Plotkin, A.S. (2018) The post-GWAS era: from association to function. *American Journal of Human Genetics*, 102(5), pp. 717-730.

Gamazon, E.R., Wheeler, H.E., Shah, K.P., Mozaffari, S.V., Aquino-Michaels, K. and Carroll, R.J. (2015) A gene-based association method for mapping traits using reference transcriptome data. *Nature Genetics*, 47(9), pp. 1091-1098.

Gasparini, M., Tome, J.M. and Shendure, J. (2020) Towards a comprehensive catalogue of validated and target-linked human enhancers. *Nature Reviews Genetics*, 21(5), pp. 292-310.

Kircher, M., Witten, D.M., Jain, P., O'Roak, B.J., Cooper, G.M. and Shendure, J. (2014) A general framework for estimating the relative pathogenicity of human genetic variants. *Nature Genetics*, 46(3), pp. 310-315.

Mahajan, A., Taliun, D., Thurner, M., Robertson, N.R., Torres, J.M. and Rayner, N.W. (2018) Fine-mapping type 2 diabetes loci to single-variant resolution using high-density imputation and islet-specific epigenome maps. *Nature Genetics*, 50(11), pp. 1505-1513.

Schaid, D.J., Chen, W. and Larson, N.B. (2018) From genome-wide associations to candidate causal variants by statistical fine-mapping. *Nature Reviews Genetics*, 19(8), pp. 491-504.

Schulte-Sasse, R., Budach, S., Hniz, D. and Marsico, A. (2021) Integration of multiomics data with graph convolutional networks to identify new cancer genes. *Nature Machine Intelligence*, 3(6), pp. 513-526.

Trubetskoy, V., Pardinas, A.F., Qi, T., Panagiotaropoulou, G., Awasthi, S. and Bigdeli, T.B. (2022) Mapping genomic loci implicates genes and synaptic biology in schizophrenia. *Nature*, 604(7906), pp. 502-508.

Visscher, P.M., Wray, N.R., Zhang, Q., Sklar, P., McCarthy, M.I. and Brown, M.A. (2017) 10 years of GWAS discovery: biology, function, and translation. *American Journal of Human Genetics*, 101(1), pp. 5-22.

Wainberg, M., Sinnott-Armstrong, N., Mancuso, N., Barbeira, A.N., Knowles, D.A. and Golan, D. (2019) Opportunities and challenges for transcriptome-wide association studies. *Nature Genetics*, 51(4), pp. 592-599.

Wang, G., Sarkar, A., Carbonetto, P. and Stephens, M. (2020) A simple new approach to variable selection in regression, with application to genetic fine-mapping. *Journal of the Royal Statistical Society Series B*, 82(5), pp. 1273-1300.

Weissbrod, O., Hormozdiari, F., Benber, C., Burns, R., Kathiresan, S. and Hirschhorn, J. (2020) Functionally informed fine-mapping and polygenic localization of complex trait heritability. *Nature Genetics*, 52(12), pp. 1355-1363.

Boix, C.A., James, B.T., Park, Y.P., Meuleman, W. and Kellis, M. (2021) Regulatory genomic circuitry of human disease loci by integrative epigenomics. *Nature*, 590(7845), pp. 300-307.

Chen, K.M., Wong, A.K., Troyanskaya, O.G. and Zhou, J. (2022) A sequence-based global map of regulatory activity for deciphering human genetics. *Nature Genetics*, 54(7), pp. 940-949.

Weeks, E.M., Ulirsch, J.C., Cheng, N.Y., Trippe, B.L., Fine, R.S. and Dey, J. (2023) Leveraging polygenic enrichments of gene features to predict genes underlying complex traits and diseases. *Nature Genetics*, 55(8), pp. 1267-1276.

Declarations

Funding

This study was supported by the Spanish Ministry of Science and Innovation Ramon y Cajal fellowship RYC2021-034826-I (S. Moreau) and the EU Horizon Europe project GenomeMed (grant 101057619). Computational resources were provided by the Barcelona Supercomputing Centre (BSC-CNS) under grant BCV-2023-1-0024. UK Biobank data accessed under application 78924. The funding bodies had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

S. Moreau declares no competing interests. No equity, consulting fees, or patents are associated with this work.

Data Availability Statement

GWAS summary statistics are available from their respective consortia (DIAGRAM, PGC, IBD-GC). UK Biobank exome data are available through the UK Biobank Access Management System. GT-VEP source code (PyTorch Geometric), pre-trained models, and disease-specific gene prioritisation results are available at <https://github.com/WEDSU-CompGen/GT-VEP-2023> under MIT licence.

Ethical Approval

This study used publicly available, de-identified GWAS summary statistics and UK Biobank data accessed under approved application 78924 (CITI training completed). No additional ethics approval was required for this computational secondary analysis.

Appendix A

GT-VEP Architecture and Training Configuration

This appendix provides the complete GT-VEP architecture specification, input feature descriptions, training hyperparameters, and ablation study results.

A.1 Architecture Details

- 1a. Variant embeddings: Enformer 896-track predictions (512-d PCA) + SuSiE PIP + phyloP + TSS distance = 516-d per variant
- 1b. Variant-to-gene attention: 4 heads, $d_k = 64$; attention modulated by eQTL beta, Hi-C contact, ABC score per variant-gene pair
- 2a. Gene features: GTEx 54-tissue expression (54-d) + gnomAD pLI + GO annotations (128-d) + STRING degree = 184-d per gene
- 2b. Graph attention network: 4 layers, 8 heads, $d_{\text{model}} 256$, ELU, residual; heterogeneous edges: PPI + regulatory + pathway
- 3a. Output: per-gene sigmoid probability (causal vs non-causal); per-pathway aggregate score (mean of member gene probabilities)
- 3b. Total: 12.8 M parameters; training 6.2 h per fold on single A100; inference 45 s per disease (genome-wide)

A.2 Ablation Study (T2D AUROC)

- 4a. Full GT-VEP: 0.846; remove network layer (variant + gene only): 0.798 (delta -0.048)
- 4b. Remove variant layer (gene + network only): 0.782 (delta -0.064); remove eQTL/Hi-C attention modulation: 0.824 (delta -0.022)
- 5a. Replace graph attention with mean aggregation: 0.812 (delta -0.034); use only PPI edges (no regulatory/pathway): 0.828 (delta -0.018)
- 5b. Conclusion: variant features contribute +6.4%; network layer +4.8%; eQTL/Hi-C attention +2.2%; heterogeneous edges +1.8%