Multi-omic deep learning identifies exercise-responsive ageing pathways in humans

Ciara G. Juan[1], Lazaros Ntasis[2]

1. Faculty of Life and Health Sciences, Ulster University, Belfast, BT15 1ED
2. Department of Sport Organisation and Management, University of Peloponnese, Sparta, Greece 231 00

Abstract

Genome-wide association studies of physical activity traits have mapped numerous loci, yet the molecular mechanisms through which exercise influences human biology remain poorly defined. Mechanistic progress has been limited by heritability-dominated signals, siloed single-omic analyses, and the lack of integrative models that connect genetic associations to causal, system-level pathways. We introduce the first deep learning, multi-omic framework for exercise genomics, unifying causal inference, molecular topology, protein structure, and functional context within a supervised Graph Neural Network (GNN) with experimental validation. Using linkage disequilibrium-aware Mendelian randomisation with accelerometer-derived vigorous physical activity as the exposure, we integrated four omic layers—plasma proteomics, blood CpG methylation, blood single-cell transcriptomics, and plasma glycomics. The GNN prioritised a coherent, multi-omic network of exercise-responsive genes spanning glycosylation and immunity, apoptosis/stress signalling, growth and transcription factors, proteostasis/autophagy, metabolism, oxidative stress/redox, mitochondrial/oxidative phosphorylation, chromatin/epigenetic, translation/ribosome, RNA splicing/processing, DNA damage and repair, cell cycle, and cytoskeleton/ECM functions. Pathway annotation and alignment with ageing biology using differentially methylated regions and principal component features from five biological ageing clocks (Horvath, Hannum, DunedinPACE, PhenoAge, Proteomic Clock), used as contextual markers of ageing-related regulation, mapped to: 1) autophagy (Fas signalling); 2) stress and DNA damage response pathways (Hypoxia-Inducible Factor, apoptosis, p53, p38, PI3K, Ras, oxidative stress, DNA replication, purine/pyrimidine metabolism and biosynthesis, pentose phosphate, Rho GTPase, G protein, ubiquitin proteasome, Epidermal Growth Factor Receptor); 3) immune and inflammatory pathways (cytokine and interleukin signalling, T and B cell activation, Toll-like receptor signalling); 4) physiological adaptation (VEGF, Wnt, GnRH, Fibroblast Growth Factor, Platelet-Derived Growth Factor, Thyrotropin-Releasing Hormone, glutamate receptor, plasminogen, endothelin, heme biosynthesis, cholecystokinin, Corticotropin-Releasing Hormone signalling); 5) neuroendocrine and neurotransmitter pathways (acetylcholine, dopamine, oxytocin, opioid, β-adrenergic signalling); 6) neurodegeneration pathways (Alzheimer's, Parkinson's, Huntington's disease); 7) protein metabolism (leucine, isoleucine, valine biosynthesis); and exercise-responsive epigenetic regulatory pathways (S-adenosylmethionine, thiamine, vitamin D biosynthesis, bZIP transcription, and circadian rhythm). Finally, we partially validated GNN predictions in humans, demonstrating acute exercise-induced shifts in plasma glycomic markers consistent with GNN-predicted glycomic remodelling. This work establishes a deep learning, multi-omic map of exercise-responsive pathways in humans and identifies actionable regulators that couple habitual vigorous physical activity to stress-resilient immunometabolic regulation and healthy ageing trajectories.

Keywords: deep learning, graph neural network, mendelian randomisation, multi-omics, physical activity, accelerometry, ageing, ageing clocks

Introduction

Regular physical activity is one of the most effective behavioural interventions for preserving functional capacity and delaying the onset of age-related multimorbidity and mortality. Yet, despite robust epidemiological evidence, the molecular pathways through which habitual exercise reconfigures proteomic, transcriptomic, epigenomic, and glycomic networks in humans remain incompletely understood. Dissecting these pathways is essential for moving beyond observational associations to mechanistic insight, and for identifying molecular regulators that could be targeted to mimic or potentiate the benefits of physical activity.

Large-scale genome-wide association studies (GWAS) in the UK Biobank (UKBB) have begun to map the genetic architecture of physical activity. Early efforts using self-reported and accelerometer-derived measures identified loci associated with habitual activity and sedentary behaviour (Klimentidis et al., 2018). Subsequent analyses of seven-day wrist accelerometry extended these findings to objectively measured activity, sleep, and circadian traits, identifying loci associated with overall movement and rhythmicity (Doherty et al., 2018). More recent multi-trait accelerometry GWAS further refined genetic signals across activity intensity, timing, and bout structure, using supervised algorithms to extract higher-resolution behavioural phenotypes from raw device data (Qi et al., 2022). Collectively, these studies demonstrate that physical activity is highly polygenic, with modest per-variant effects and limited variance explained, but also illustrate how machine learning-derived phenotypes can improve downstream molecular inference.

Importantly, device-based epidemiology now suggests that activity intensity is a key determinant of health. Using wrist-worn accelerometry in UKBB, Biswas et al. (2025) showed that as little as one minute per day of vigorous physical activity (VPA) was associated with health benefits comparable to tens to over 100 minutes of light-intensity activity, with outcome-specific equivalence estimates of approximately ~90 minutes for diabetes risk, ~70–85 minutes for cardiovascular disease risk, and >150 minutes for cancer outcomes. These findings reinforce a growing consensus that VPA carries disproportionate biological benefits, and motivate a mechanistic focus on pathways that respond specifically to higher-intensity exercise exposures.

Despite increasingly refined behavioural phenotypes, translating polygenic physical activity signals into coherent molecular mechanisms remains challenging. Exercise-responsive pathways are dispersed across molecular layers and rarely converge through standard association-based prioritisation. This reflects broader limitations of GWAS translation, where loci are typically ranked by statistical significance rather than by biological efficiency or functional relevance. A recent appraisal of GWAS highlighted several structural barriers to translation—including reliance on additive heritability, proliferation of small effect associations, under-utilisation of multi-omic data, and the absence of frameworks that prioritise genes by how efficiently they realise a trait (Huang et al., 2025). In response, Huang and

colleagues introduced the concept of trait efficiency loci (TEL), arguing that the most informative regulators of complex traits may be those that mediate efficient structural and functional realisation rather than those with the largest marginal effects.

In parallel, Spence et al. (2025) demonstrated that conventional GWAS ranking strategies are systematically biased by gene length, pleiotropy, and stochastic variation, and formalised trait importance as a quantitative property of gene-trait relationships. Together, these frameworks suggest that key regulators of complex phenotypes—such as exercise adaptation—may be genes that are both causally implicated and functionally efficient, even if they do not achieve top statistical significance in any single GWAS.

Here, we build on these conceptual advances by combining causal inference from Mendelian randomisation (MR) with a supervised Graph Attention Network (GAT) to prioritise exercise-responsive genes across molecular layers. The model integrates causal MR effects with structural, functional, and topological neighbourhoods to learn attention-weighted representations of exercise adaptation. To connect exercise biology with human ageing, we map these Graph Neural Network or GNN-prioritised genes onto two independent classes of ageing signatures: (i) differentially methylated regions underlying four established epigenetic ageing clocks (Horvath, Hannum, DunedinPACE, and PhenoAge), and (ii) genes contributing to the first principal component of plasma proteomic ageing clocks across multiple cohorts.

Finally, we experimentally validate key glycomic predictions in a human vigorous exercise intervention, linking genetically proxied habitual activity to acute changes in plasma glycosylation. Together, this framework—combining causal inference, supervised graph learning, and experimental validation—provides a mechanistic map of exercise-responsive pathways and illustrates how AI-guided multi-omics can bridge the long-standing gap between GWAS discovery and biological mechanism.

Results

Omic-specific MR analyses identify widespread but largely non-overlapping associations with vigorous activity

We first performed a coding region GWAS with linkage disequilibrium (LD) mapping to identify single-nucleotide polymorphisms (SNPs) associated with accelerometry-derived VPA in UKBB (n ≈ 75,000, predominantly European). We identified 87 VPA SNPs and performed LD-aware, overlap-aware MR to identify molecular processes causally influenced by genetically proxied habitual VPA across four orthogonal omic layers: proteomic (UKBB plasma proteomics, n ≈ 54,000, predominantly European), epigenomic (GoDMC whole blood CpG methylation, n = 27,750, European), glycomic (TPNG GWAMA plasma glycan

peaks, n = 10,172, European), and transcriptomic (OneK1K Peripheral Blood Mononuclear Cells, n = 982, predominantly European).

Using a predefined nominal threshold of $p < 0.05$, each omic layer exhibited distinct degrees of MR signal density. In the proteomic layer, 2,991 proteins had valid MR estimates, of which 90.0% were nominally significant ($p < 0.05$) and 89.7% remained significant after Benjamini–Hochberg false discovery rate (FDR) correction. Comparable proportions were observed for LD-aware fixed-effect MR estimates (88.9% nominal significance; 88.5% FDR-adjusted). These results indicate that multiple testing correction does not materially attenuate signal density in proteomic MR analyses, consistent with extensive shared genetic architecture and horizontal pleiotropy in large-scale proteomic datasets. Accordingly, proteomic MR signals were treated as noisy causal priors rather than independent causal discoveries, motivating the use of continuous effect size-based scoring and network-level prioritisation within the GNN framework rather than binary significance filtering.

Strong but variable MR support was also observed across other molecular layers. In the epigenomic layer, 151 of 184 VPA SNP-matched CpG sites (82.1%) showed significant causal associations, while in the glycomic layer, 101 of 117 VPA SNP-matched glycan traits (86.3%) reached MR significance. In contrast, the single-cell transcriptomic layer exhibited more selective evidence, with 119 of 364 gene-cell type tests (32.7%) meeting criteria for MR support. Within this layer, the strongest signals were observed in B cells (33 of 82 tests) and dendritic cells (15 of 38), followed by CD8 T cells (21 of 68), monocytes (17 of 60), NK cells (16 of 55), and CD4 T cells (17 of 61).

Because molecular layers represent features at distinct biological scales, MR results were harmonised to gene symbols prior to integration. This yielded 2,304 genes with nominally significant proteomic MR associations, 19 genes supported by epigenomic MR, 31 genes supported by glycomic MR, and 35 genes supported by single-cell transcriptomic MR. Integrating gene-level MR evidence across all layers identified 2,370 unique genes with causal support from at least one omic layer; however, only 21 genes showed convergent evidence across two molecular layers, and a single gene exhibited nominal MR support across three layers (Table 1).

Table 1. 21 MR-significant genes with multi-omic support (≥2 omic layers)

Gene	Protein MR	CpG MR	Glycan MR	scRNA MR	Omic Layers	Function
NMT1	1	1	0	1	3	N-myristoyltransferase; signalling and immune metabolism
ACBD4	0	1	0	1	2	Peroxisomal/mitochondrial acyl-CoA–binding protein
ARHGAP27	0	1	0	1	2	Rho GTPase regulator; cytoskeletal and vesicular regulation

B4GALT1	1	0	1	0	2	Core galactosyltransferase; glyco-immunity
BCL7B	1	0	1	0	2	Chromatin remodelling factor; lymphocyte regulatory roles
C2	1	0	1	0	2	Immune response
CDC27	1	1	0	0	2	Cell cycle and proteostasis
CRHR1	1	1	0	0	2	CRH receptor; neuroendocrine, stress-immune signalling
FCGR3B	1	0	1	0	2	Fcγ receptor; neutrophil activity
FMNL1	1	0	0	1	2	Actin regulator; leukocyte migration, immune signalling
FUT8	1	0	1	0	2	Core fucosyltransferase; glyco-immunity
HEXIM1	1	0	0	1	2	P-TEFb (CDK9–Cyclin T) regulator; transcriptional control
HNF1A	1	0	1	0	2	Hepatic transcription factor; metabolic & inflammatory regulation
LRRC37A2	1	0	0	1	2	Linked to MAPT locus; roles in neuronal structure and immune signalling
MAP3K14	0	1	0	1	2	NF-κB-inducing kinase; non-canonical NF-κB pathway
MAPT	1	1	0	0	2	Tau protein; microtubule stability
MAX	1	0	1	0	2	MYC-binding transcriptional co-regulator
SERPINA1	1	0	1	0	2	Alpha-1 antitrypsin; protease inhibition & inflammation control
SIRT1	1	1	0	0	2	NAD+ deacetylase; genomic maintenance and metabolic adaptation
ST6GAL1	1	0	1	0	2	Sialyltransferase; glyco-immunity
TNFRSF13B	1	0	1	0	2	BAFF/APRIL receptor; B-cell signalling

Although gene-level aggregation revealed widespread causal associations, the extreme sparsity of multi-omic overlap—only 21 of 2,370 genes supported in two or more layers—posed a major challenge for biological interpretation. This pattern reflects both genuine molecular specificity and the statistical limitations of *p* value-based thresholds, which compress continuous effect sizes into binary significance categories and obscure relative causal magnitude across layers. As noted by Spence et al. (2025), reliance on *p* values alone can flatten gene rankings and fail to distinguish biologically meaningful signals when multiple traits, or in this case multiple omic layers, provide heterogeneous effect sizes and statistical power.

Together, these observations highlight substantial heterogeneity in MR signal density and limited cross-layer convergence, underscoring the need for integrative, network-based prioritisation approaches that synthesise continuous causal

evidence while accounting for pleiotropy, redundancy, and context-specific regulatory effects rather than relying on nominal overlap alone.

GNN identifies modular architecture within MR-supported exercise-responsive genes

To overcome the limitations of binary significance thresholds and to quantify the magnitude rather than the mere presence of causal evidence, we developed a continuous Multi-omic Trait Importance (MTI) score that aggregates MR effect sizes across all molecular layers. This approach generalises the magnitude-based logic introduced by Spence et al. (2025) while avoiding saturation effects caused by pervasive proteomic significance. By preserving relative differences among genes with modest but consistent multi-omic support, the continuous MTI enabled balanced integration of proteomic, epigenomic, glycomic, and single-cell transcriptomic evidence without disproportionate dominance of highly powered layers.

While MTI provides a refined gene-level prioritisation of exercise-responsive biology, it does not explicitly model the relational structure of molecular systems. To capture higher-order organisation beyond individual gene scores, we trained a supervised Graph Attention Network (GAT-GNN) to predict MTI from integrated multi-omic, structural, and functional features. This framework enabled the model to learn biologically meaningful neighbourhoods, network topology, and context-dependent interactions that are not accessible through MTI or MR rankings alone, while maintaining robustness to heterogeneous MR signal density across layers.

Application of the GNN revealed a highly modular architecture of exercise-responsive biology. The top 1,000 GNN-ranked genes organised into dense, interpretable clusters spanning glycosylation and immune regulation, apoptosis and cellular stress signalling, growth factor and transcriptional control, proteostasis and autophagy, metabolic regulation, oxidative stress and redox pathways, mitochondrial oxidative phosphorylation, chromatin and epigenetic regulation, translation and ribosomal function, RNA splicing and processing, DNA damage and repair, cell-cycle control, and cytoskeletal and extracellular matrix organisation (Figure 1).

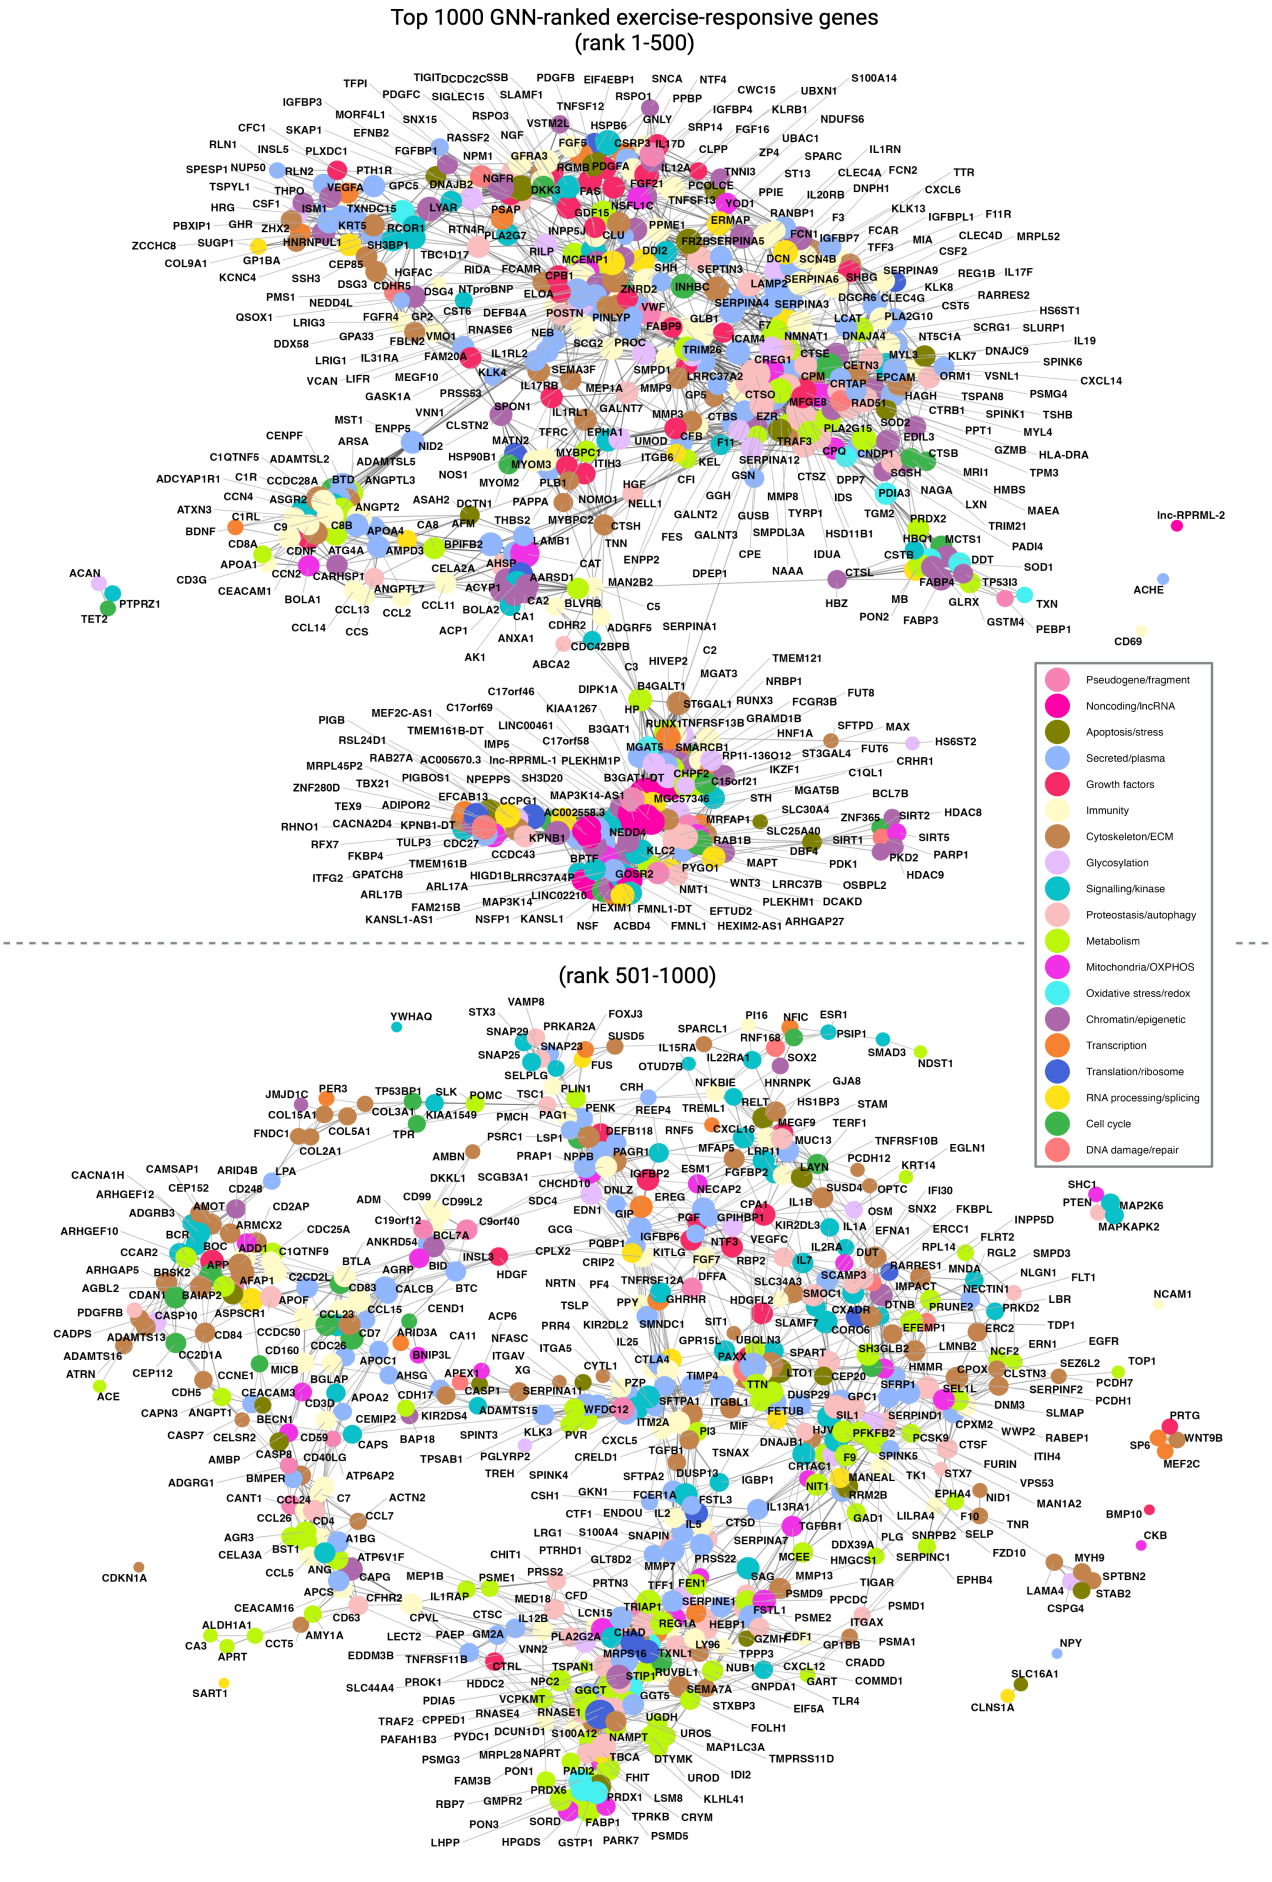
Top 1000 GNN-ranked exercise-responsive genes
(rank 1-500)
(rank 501-1000)
Pseudogene/fragment
Noncoding/lncRNA
Apoptosis/stress
Secreted/plasma
Growth factors
Immunity
Cytoskeleton/ECM
Glycosylation
Signalling/kinase
Proteostasis/autophagy
Metabolism
Mitochondria/OXPHOS
Oxidative stress/redox
Chromatin/epigenetic
Transcription
Translation/ribosome
RNA processing/splicing
Cell cycle
DNA damage/repair

Figure 1. Top 1,000 GNN-ranked exercise-responsive genes. *Nodes are coloured by Uniprot and Interpro pathways and sized by GNN score, with edges representing STRING-derived functional associations.*

Importantly, many of the highest-ranked genes exhibited only modest MR effect sizes, underscoring that supervised graph learning can recognise higher order biological organisation—detecting coordinated mechanistic modules that causal effect sizes or MTI alone fail to reveal. Together, these results show that supervised GNN analysis recovers the latent pathway architecture of exercise adaptation.

Exercise-ageing integration reveals a shared, multi-pathway architecture

To determine how exercise-responsive genes intersect with ageing biology, we integrated our top GNN-prioritised gene set with differentially methylated regions (DMRs) from four epigenetic ageing clocks (Horvath, Hannum, PhenoAge, DunedinPACE) and with first principal component (PC1) genes from plasma Proteomic Clocks (Figure 2A). Across all clocks, we observed substantial gene-level convergence, with a large subset of exercise-responsive genes also contributing to inter-individual variation in biological age (Figure 2B). Among 2,370 genes showing nominal MR significance in at least one omic layer, the largest intersections were observed for PhenoAge (1,021 genes) and DunedinPACE (742 genes), followed by Hannum (605) and Horvath (456) clocks, and the Proteomic Clock (340) (Figure 2B).

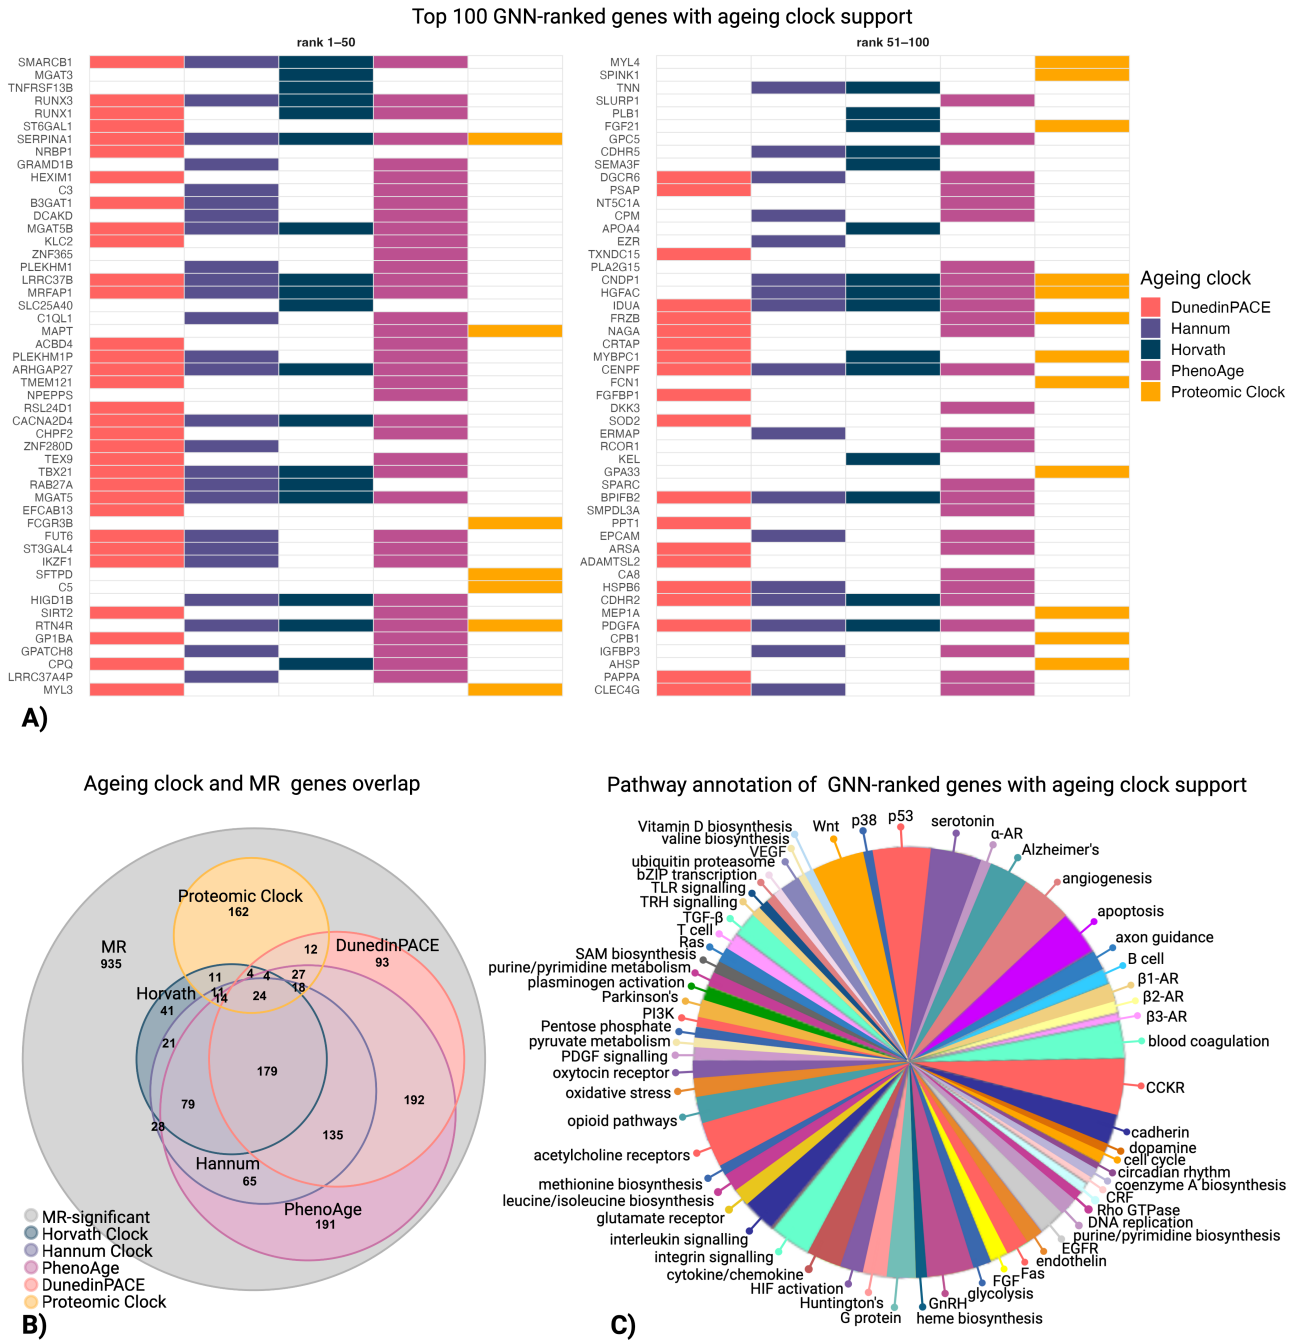

Figure 2. Ageing clock support, multi-clock overlap, and pathway annotation of exercise-responsive genes prioritised by GNN.

A) Ageing clock coverage of GNN-ranked exercise-responsive genes. The top 100 GNN-ranked genes with support from at least one ageing clock are shown.

B) Proportional Euler diagram of overlap between MR-significant exercise-responsive genes and ageing clock DMR or PC1 genes.

C) Functional annotation of GNN-prioritised exercise-responsive genes with ageing clock support. PANTHER pathway analysis of top 1,000 GNN-ranked genes are shown. α-AR Alpha Adrenergic Receptor; β1-AR Beta-1 Adrenergic Receptor; β2-AR Beta-2 Adrenergic Receptor; β3-AR Beta-3 Adrenergic Receptor; CCKR Cholecystokinin receptor; *CRF Corticotropin-Releasing Factor; EGFR Epidermal Growth Factor Receptor; FGF Fibroblast Growth Factor; GnRH Gonadotropin-Releasing Hormone; HIF Hypoxia-Inducible Factor; PDGF Platelet-Derived Growth Factor; THR Thyroid Hormone Receptor; TLR Toll-like Receptor; VEGF Vascular Endothelial Growth Factor*

Pathway annotation of the top 1,000 exercise-responsive genes with at least one ageing clock support (Figure 2C) revealed a broad and diverse pathway landscape mapped to: 1) autophagy (Fas signalling); 2) stress and DNA damage response pathways (Hypoxia-Inducible Factor, apoptosis, p53, p38, PI3K, Ras, oxidative stress, DNA replication, purine/pyrimidine metabolism and biosynthesis, pentose phosphate, Rho GTPase, G protein, ubiquitin proteasome, Epidermal Growth Factor Receptor; 3) immune and inflammatory pathways (cytokine and interleukin signalling, T and B cell activation, Toll-like receptor signalling); 4) physiological adaptation (VEGF, Wnt, GnRH, Fibroblast Growth Factor, Platelet-Derived Growth Factor, Thyrotropin-Releasing Hormone, glutamate receptor, plasminogen, endothelin, heme biosynthesis, cholecystokinin, Corticotropin-Releasing Hormone signalling); 5) neuroendocrine and neurotransmitter pathways (acetylcholine, dopamine, oxytocin, opioid, β-adrenergic signalling); 6) neurodegeration pathways (Alzheimer's, Parkinson's, Huntington's disease); 7) protein metabolism (leucine, isoleucine, valine biosynthesis); and exercise-responsive epigenetic pathways (S-adenosylmethionine, thiamine, Vitamin D biosynthesis, bZIP transcription, and circadian rhythm).

Together, these results show that exercise targets a shared molecular architecture with ageing—providing a mechanistic basis for how physical activity confers resilience across multiple hallmarks of ageing.

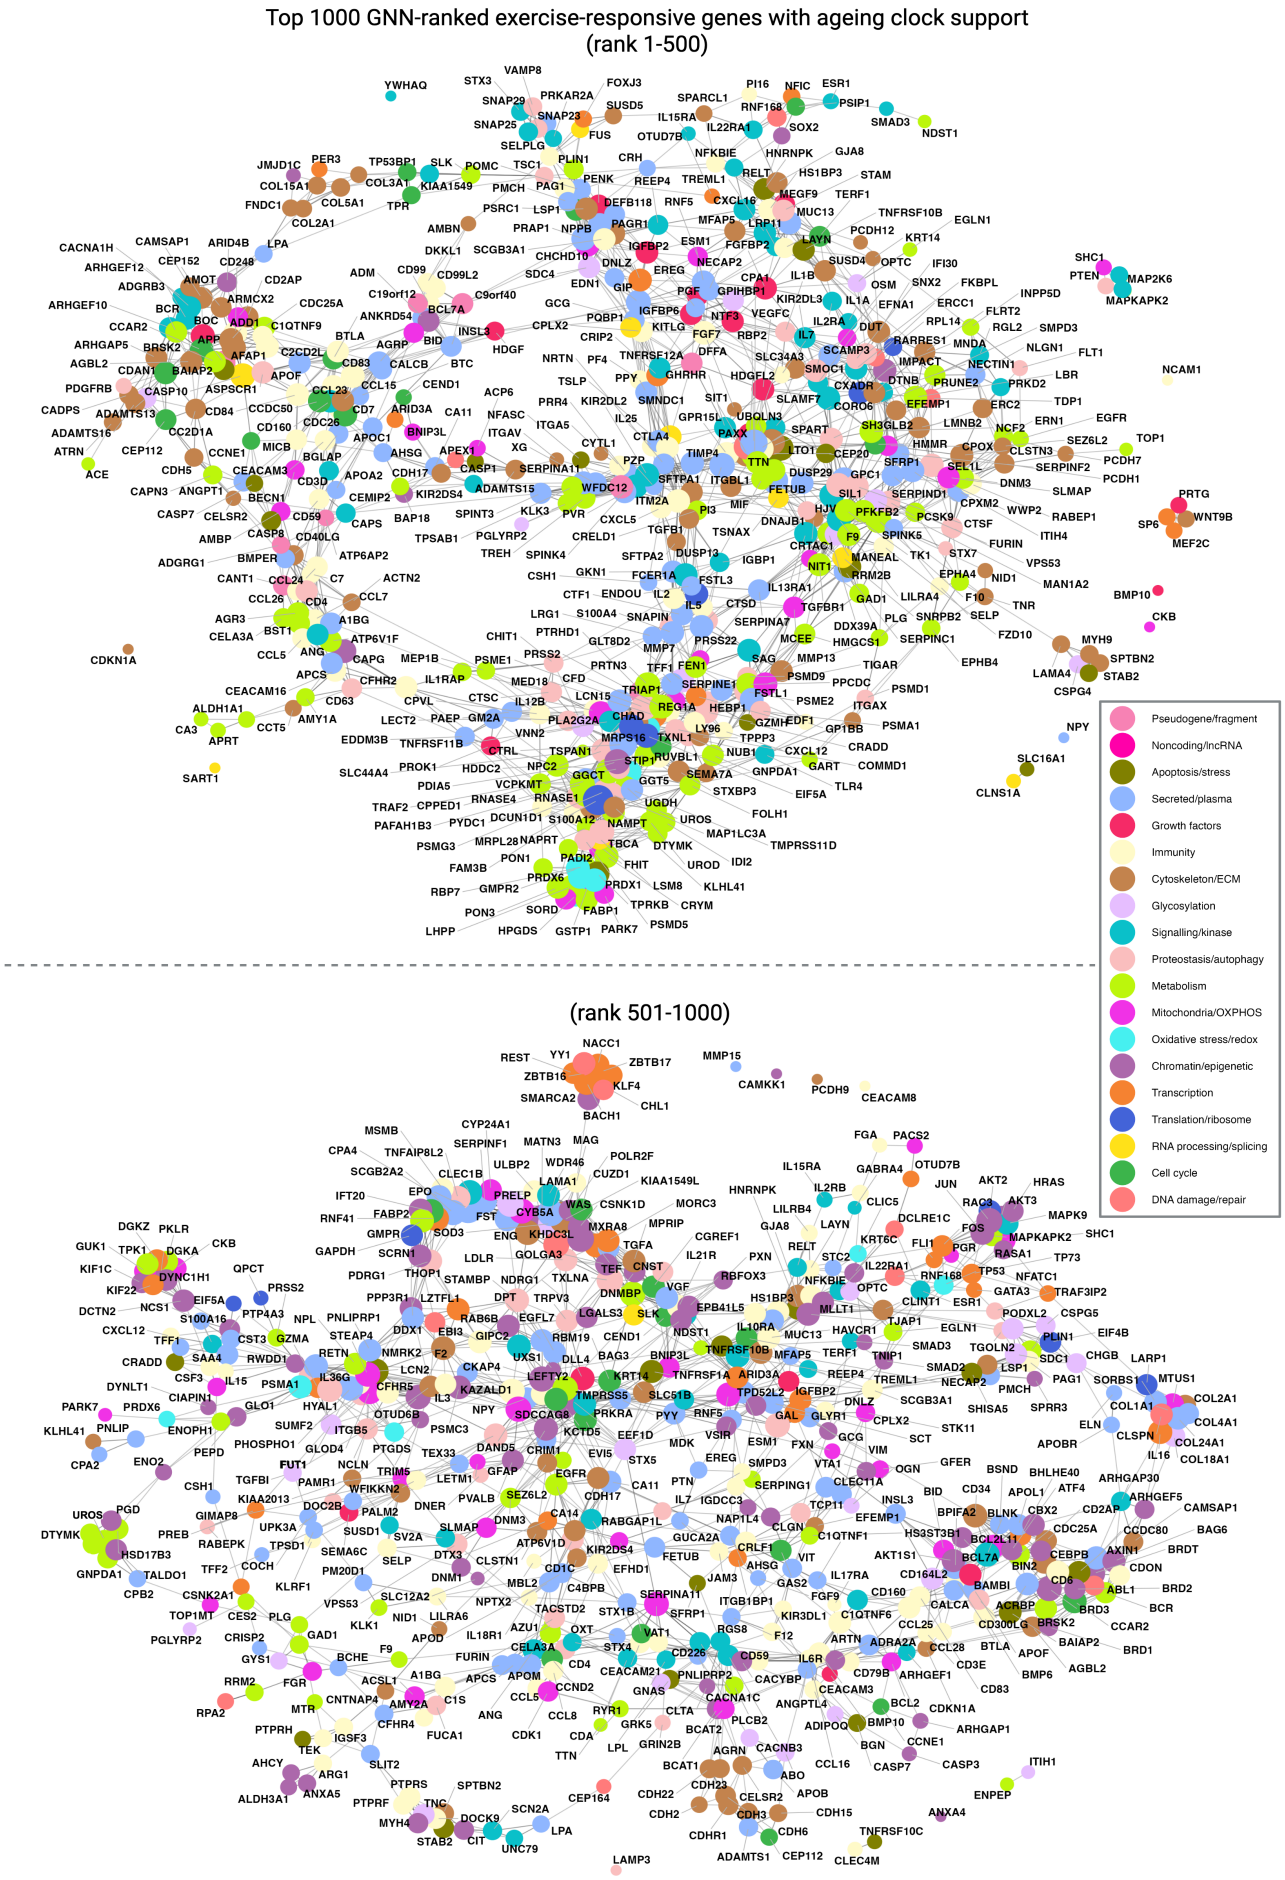
Top 1000 GNN-ranked exercise-responsive genes with ageing clock support
(rank 1-500)
(rank 501-1000)
Pseudogene/fragment
Noncoding/lncRNA
Apoptosis/stress
Secreted/plasma
Growth factors
Immunity
Cytoskeleton/ECM
Glycosylation
Signalling/kinase
Proteostasis/autophagy
Metabolism
Mitochondria/OXPHOS
Oxidative stress/redox
Chromatin/epigenetic
Transcription
Translation/ribosome
RNA processing/splicing
Cell cycle
DNA damage/repair

Figure 3. Top 1,000 GNN-ranked exercise-responsive genes with ageing clock support. *Nodes are coloured by Uniprot and Interpro pathways and sized by GNN score, with edges representing STRING-derived functional associations.*

Biological validation links acute exercise responses to glycomic adaptation

To validate mechanistic links between genetically proxied habitual vigorous physical activity and downstream molecular phenotypes, we analysed data from our human clinical study examining the acute effects of high-intensity interval exercise (~80% VO_2max) on plasma glycomic markers. Full details of the exercise protocol, sample collection, and analytical workflow are published elsewhere (Juan et al., 2025, preprint). The glycomic assay captured all plasma N-glycan traits incorporated into the present multi-omic framework, enabling direct comparison between experimental perturbation and genetically anchored inference.

Paired pre to post-exercise analyses revealed a rapid and coordinated glycomic response characterised by reduced agalactosylation (G0), increased mono and di-galactosylation (G1, G2), higher total galactosylation (Gtotal), and increased total sialylation (Stotal) (all $p < 0.05$). These changes are consistent with established immunomodulatory plasma glycan patterns, in which lower agalactosylation and higher galactosylation and sialylation are associated with reduced inflammatory potential and lower glycan-based biological age (Cindrić et al., 2024; Rapčan et al., 2024; Pribić et al., 2025).

At the network level, these acute plasma glycomic responses converged with MR and GNN-prioritised glycosylation pathways. Among the top 100 GNN-ranked exercise-responsive genes, multiple canonical enzymes involved in galactosylation, sialylation, and branching were consistently prioritised, including *B4GALT1, ST6GAL1, ST3GAL1, ST3GAL4, MGAT3, MGAT5, and MGAT5B*, indicating engagement of core N-glycan biosynthetic pathways rather than isolated molecular effects.

Importantly, MR effect directions for individual glycosylation enzymes were heterogeneous and layer-specific. In the glycomic layer, *B4GALT1, MGAT3*, and *ST6GAL1* showed positive MR associations with habitual VPA, consistent with engagement of galactosylation, branching-dependent remodelling, and terminal sialylation pathways. In contrast, *ST3GAL4, MGAT5*, and *MGAT5B* exhibited negative glycomic MR associations, while *ST3GAL1* showed a negative association at the proteomic layer. These negative associations are biologically plausible given the competitive and context-dependent nature of N-glycan biosynthesis. Reduced MGAT5-dependent branching can redirect glycan flux toward alternative processing routes and alter downstream galactosylation and sialylation outcomes (Taniguchi and Kizuka, 2015; Stanley et al., 2022). Likewise, members of the ST3GAL family exhibit substrate and tissue-specific activity, such

that divergent associations across molecular layers are expected rather than contradictory (Harduin-Lepers et al., 2005; Varki et al., 2022).

Accordingly, convergence between acute exercise-induced glycomic changes and genetically anchored MR and GNN findings emerged at the level of pathway engagement and network coordination, rather than uniform directional effects of individual enzymes. These findings support a model in which repeated exercise-induced perturbations cumulatively modulate glycosylation networks relevant to immune regulation and biological ageing, despite heterogeneity in enzyme-specific effect directions and differences in temporal scale.

Robust recovery of exercise-responsive pathways independent of MR features

To assess whether the supervised GNN was dominated by MR input features, we retrained the model after removing all MR-derived information and compared the resulting gene rankings with those from the full model. Despite complete removal of MR features, global gene ordering was strongly preserved (Spearman $\rho = 0.875$, $p < 1\times10^{-16}$; Figure 4A). Most genes remained close to the diagonal, indicating that the latent representation learned by the GNN was largely independent of MR effect sizes.

Importantly, systematic deviations from the diagonal were observed, reflecting biologically meaningful reprioritisation rather than random noise. Genes prioritised more highly by the MR-ablated model formed distinct clusters, indicating that structural, functional, and network-topological features alone were sufficient to elevate specific genes in the absence of MR signal.

We next examined stability of top-ranked genes under feature ablation. Comparison of the top K gene sets revealed substantial overlap between the full and MR-ablated models, with 57–61% Jaccard overlap across K = 100-1000 (Figure 4B). This demonstrates that more than half of the highest-ranked genes are preserved even when all MR features are removed.

In contrast, overlap between the full model and the MR-only model was higher for the most extreme ranks (top 100 overlap ≈ 0.89) but declined more rapidly as K increased. This pattern indicates that while MR captures strong, high-confidence signals, it lacks the breadth provided by network and structural information. Together, these results show that the GNN is not simply recapitulating MR rankings, but instead integrates MR with orthogonal biological information to learn a robust, systems-level representation of exercise-responsive ageing pathways.

We next tested whether enrichment for ageing-related genes among top-ranked exercise-responsive candidates was robust to removal of MR-derived node features (Figure 4C). Across epigenetic ageing clock gene sets, overlap with the top 1,000 GNN-ranked genes was highly consistent between the full model and

the MR-ablated model (DunedinPACE: 312 vs 308; Hannum: 250 vs 245; Horvath: 190 vs 189; PhenoAge: 414 vs 408). These results demonstrate that epigenetic clock overlap is largely preserved under complete MR feature ablation.

For the proteomic ageing signature, overlap with the proteomic PC1 gene set remained substantial but was partially reduced following MR feature removal (149 vs 123), consistent with proteomic-derived ageing structure being more directly aligned with MR-derived proteomic signals. Together, these findings indicate that the observed convergence between exercise-responsive genes and ageing clock signatures is not solely driven by MR feature leakage, and persists when prioritisation is driven by topology, functional annotation, and structural context.

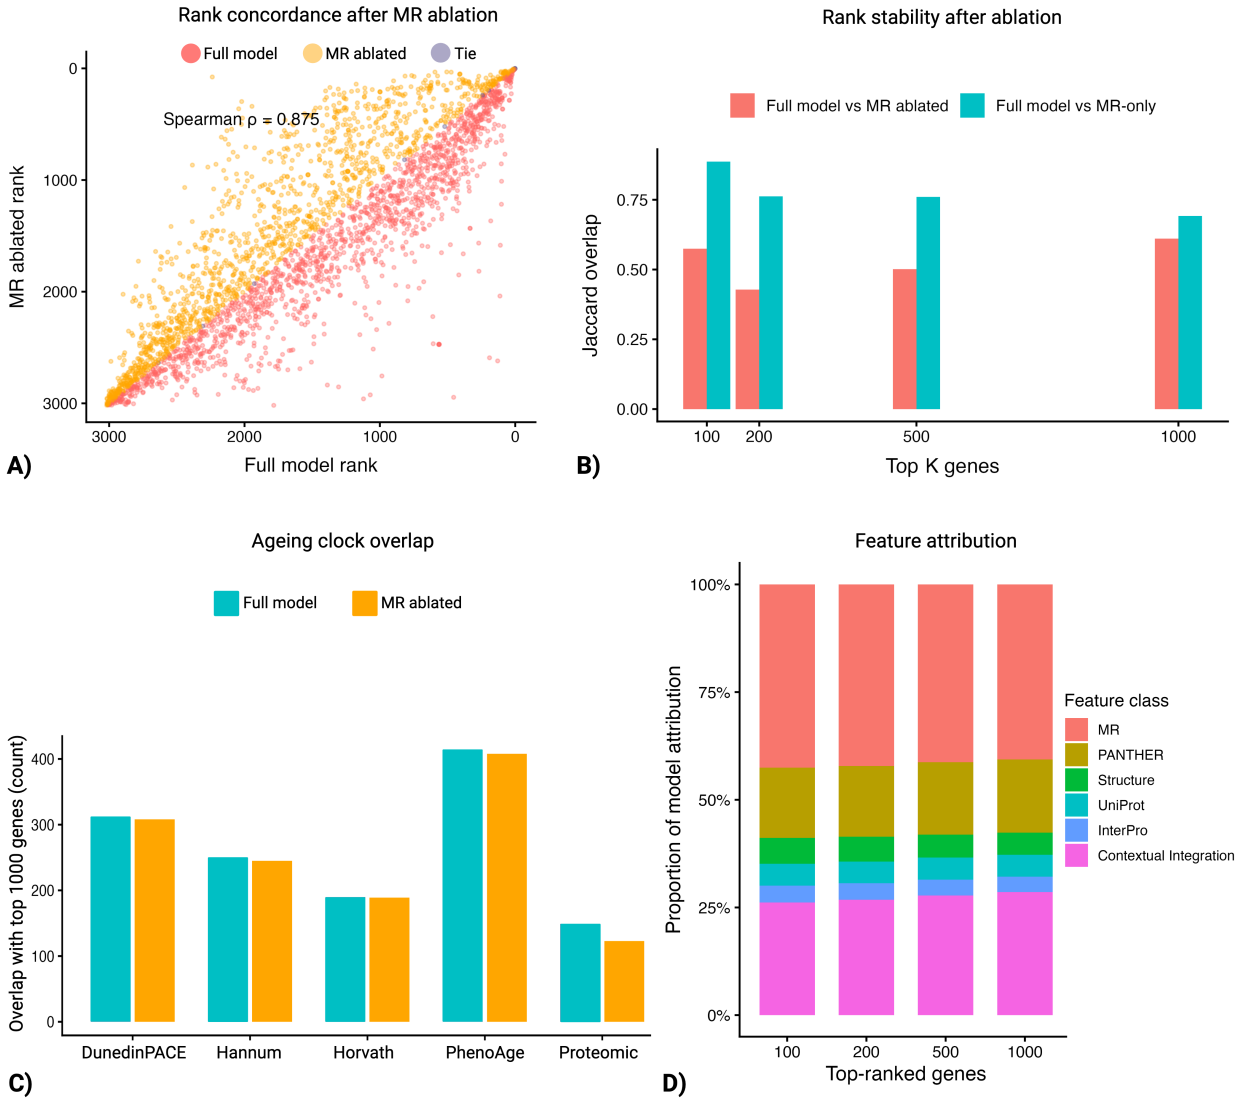

Figure 4. Robustness of gene prioritisation to MR feature ablation.

A) Rank concordance between the full GNN model and the MR-ablated model. Each point represents a gene, plotted by its rank in the full model (x-axis) and the MR-ablated model (y-axis). Points are coloured according to which model assigned a higher rank. Strong global concordance is observed (Spearman ρ = 0.875), with structured reprioritisation indicating that topology, structure, and functional annotations contribute independently of MR signal.
B) Stability of top-ranked genes under feature ablation. Bar plots show Jaccard overlap of the top K genes between the full model and the MR-ablated model, and between the full model and an MR-only model. High overlap with the MR-ablated model across K demonstrates robustness to MR removal, while divergence from the MR-only model highlights the added value of network and structural features.
C) Ageing clock overlap is preserved after MR feature ablation. Shown is the overlap between the top 1,000 GNN-ranked exercise-responsive genes and gene sets derived from established epigenetic ageing clocks (DunedinPACE, Hannum, Horvath, and PhenoAge) and Proteomic Clock. Bars indicate absolute overlap counts for the full model and the MR-ablated model.
D) Feature attribution reveals integrative gene prioritisation beyond MR signals. Gradient-based attribution of the supervised GAT model shows the proportional contribution of feature classes to gene ranking across the top 100-1000 exercise-responsive genes.

To determine whether overlap between exercise-responsive genes and ageing clock gene sets could arise trivially from the large size of ageing clocks, we performed permutation-based overlap analyses using the full ranked gene list as the reference universe. Gene ranks were randomly permuted while preserving clock gene set size, and the observed number of ageing clock-associated genes among the top-ranked exercise-responsive genes was compared with the distribution expected under permuted rankings.

Permutation analyses showed that ageing clock genes occur within the ranked exercise-responsive gene list in proportion to their overall representation, rather than being preferentially prioritised by the model (empirical permutation $p > 0.1$ across all clocks and ranking thresholds; FDR-adjusted $p = 1.0$). This indicates that observed intersections with ageing clocks are not driven by implicit optimisation toward ageing-related targets, but instead reflect baseline gene set representation. Accordingly, ageing clocks were used as an orthogonal interpretative framework to contextualise shared biological pathways between exercise responsiveness and ageing-related processes, rather than as a validation or optimisation endpoint.

Finally, we performed a gradient-based feature attribution analysis to quantify the extent to which gene prioritisation was driven by different classes of model inputs (Figure 4D). Across the top 100-1000 ranked exercise-responsive genes, MR features accounted for approximately 40.6–42.5% of total attribution, while curated PANTHER functional features contributed ~16.3–17.0%. Structural features derived from AlphaFold v6 contributed ~5.1–6.0%, and functional principal

components (PCs) derived from UniProt and InterPro together accounted for ~8.9–10.0% of the attribution.

The remaining ~26.1–28.6% of attribution arose from contextual integration, comprising node features that capture integrative biological context rather than direct causal effects or single-source annotations, including multi-omic presence indicators, cross-layer concordance metrics, and neighbourhood-aggregated descriptors reflecting local kNN network context through graph message passing.

Importantly, the relative contributions of all feature classes were stable across ranking thresholds, indicating that the model consistently integrates causal inference with curated functional pathway context and protein structural and functional descriptors, rather than reflecting MR effect sizes alone.

Discussion

The central aim of this study was to derive a mechanistic, causal interpretation of how vigorous physical activity interfaces with healthy ageing, moving beyond gene-level associations to resolve the pathway-level programmes through which exercise reshapes ageing biology. By integrating LD-aware Mendelian randomisation across multiple omic layers with supervised graph attention learning, we identify a compact set of interconnected adaptive mechanisms that translate transient physiological stress into durable molecular resilience.

At the core of this framework lies a ROS-driven stress-repair-autophagy axis. Vigorous exercise generates short-lived increases in reactive oxygen species (ROS) through elevated mitochondrial flux, immune activation, and mechanical load, with these transient redox signals acting as physiological cues that engage adaptive stress response pathways rather than causing cumulative damage, consistent with the concept of oxidative eustress in exercise physiology (Cobley & Davison, 2022). These ROS pulses activate conserved stress and damage response pathways, including p53, p38 MAPK, PI3K-Akt, Ras, hypoxia-inducible signalling, ubiquitin proteasome function, and autophagy (Leslie et al., 2003; Maillet et al., 2012; Mitchell et al., 2013; Shang & Taylor, 2011; Scherz-Shouval & Elazar, 2011; Son et al., 2011; Thomas & Ashcroft, 2019). Autophagy is a central determinant of healthy ageing, mediating the clearance of damaged proteins, organelles, and mitochondria whose accumulation drives age-related functional decline (Raza, 2024). Its prioritisation within the exercise-responsive network indicates that repeated induction of autophagic flux is a primary mechanism through which exercise preserves proteostasis and cellular integrity.

Within this stress context, transient exercise-induced DNA double-strand breaks marked by γH2AX provide a direct molecular bridge between vigorous activity and adaptive remodelling (Williamson et al., 2020). γH2AX coordinates DNA repair factor recruitment and chromatin reorganisation and has emerged as a key

interface between stress signalling, genome maintenance, and ageing biology (Gong et al., 2025). When rapidly resolved, such lesions act as controlled perturbations that reinforce repair proficiency rather than promoting cumulative genomic instability (Farhadi et al., 2022). The concurrent enrichment of DNA replication, purine and pyrimidine metabolism and biosynthesis, and the pentose phosphate pathway is mechanistically coherent, supplying NAD+, nucleotides, and reducing equivalents required for DNA repair and redox buffering. Across repeated exercise exposures, iterative engagement of this stress-repair-autophagy circuitry consolidates damage clearance and repair efficiency, with CpG methylation patterns captured by biological ageing clocks reflecting downstream, context-dependent stabilisation of regulatory states shaped by cumulative DNA damage-repair interactions (Soto-Palma et al*.,* 2022).

Within this framework, *SIRT1*, a vigorous exercise-responsive epigenetic regulator (Juan et al., 2023), functions as an upstream metabolic-chromatin-repair integrator coordinating energetic stress sensing and autophagy induction, while mitochondrial regulators such as *SIRT3*—directly represented in ageing clock features—reflect the cumulative imprint of these adaptations on clock-associated methylation and proteomic signatures. Consistent with this model, our experimental validation demonstrated that high-intensity exercise induced coordinated acute changes in plasma glycomic markers most strongly linked to biological age—particularly galactosylation and sialylation—alongside significant alterations in *SIRT1* and other DNA repair and genome maintenance factors, including *PARP1, SIRT6, OGG1*, and *RAD51* (Juan et al*.,* 2025, preprint). Together, these orthogonal molecular readouts provide convergent evidence that vigorous exercise engages systemic stress-repair and maintenance programmes aligned with the pathways prioritised by the GNN.

Protein and amino acid metabolism, including leucine, isoleucine, and valine biosynthesis, is tightly embedded within this adaptive framework. Branched-chain amino acids (BCAAs), particularly leucine, regulate mTORC1 signalling and thereby govern the balance between protein synthesis and autophagic turnover, positioning nutrient availability as a key determinant of cellular growth versus maintenance states (Papadopoli *et al.*, 2019; Saxton & Sabatini, 2017; Wolfson & Sabatini, 2017). Ageing is characterised by sustained mTOR activity and impaired autophagic flux, contributing to proteostatic stress and reduced tolerance to physiological and environmental challenges (Hansen *et al.*, 2018; Papadopoli *et al.*, 2019). Recent mechanistic work demonstrates that mTOR signalling directly controls protein aggregation during stress and cellular ageing, independently of global translation rates or canonical heat shock transcriptional programmes, establishing mTOR hyperactivity itself as a driver of age-associated proteostatic collapse (Fischbach *et al.*, 2025). In contrast, repeated exercise engages transient catabolic stress responses that suppress mTOR signalling and activate autophagy, followed by anabolic recovery phases that restore biosynthetic capacity; over time, such catabolic-anabolic cycling promotes metabolic flexibility,

damage clearance, and resilience—key determinants of healthy ageing (Calabrese & Mattson, 2024).

A second integrative layer links metabolic and epigenetic cofactors to the consolidation of adaptive responses. The prioritisation of S-adenosylmethionine (SAM) biosynthesis**,** thiamine metabolism, vitamin D signalling, and circadian rhythm regulation highlights how exercise shapes the biochemical and temporal context in which stress responses occur. Exercise modulates metabolic pathways that supply methyl donor substrates for SAM, thereby modulating DNA and histone methylation capacity (Juan et al., 2024), while high-intensity exercise training has been shown to increase circulating 1,25-dihydroxyvitamin D, which acts through the vitamin D receptor, a nuclear receptor that recruits chromatin remodellers to regulate transcription of immune and stress-response genes (Carlberg, 2018; De-la-O et al., 2022). Consistent with this framework, low vitamin D status has been associated with accelerated biological ageing, whereas sustained vitamin D availability has been linked to preservation of genomic stability and slower attrition of ageing biomarkers (Vetter et al., 2022; Zhu et al., 2025). Meanwhile, SIRT1-modulated circadian systems also gate redox buffering, autophagy, and DNA damage repair through rhythmic control of metabolic and repair pathways (Fagiani et al., 2022; Imai & Guarente, 2016). Together, these observations align with an emerging exercise-epigenetic axis, in which redox-sensitive metabolic signals generated during exercise are translated into stabilised, adaptive epigenetic regulation, reinforcing stress resilience and supporting healthy ageing (Davison et al., 2026).

A third integrative theme links mitochondrial metabolic competence with immune, plasma, and platelet-associated signalling to coordinate tissue repair and vascular health. Cytokine and interleukin pathways, Toll-like receptor signalling, and growth factor systems—including VEGF, PDGF, plasminogen, endothelin, and heme biosynthesis—calibrate inflammatory tone while engaging angiogenic and tissue-remodelling programmes (Eming et al., 2017). Central to this integration is mitochondrial heme metabolism, which gates endothelial angiogenic capacity by sustaining Complex IV integrity, respiratory reserve, and redox homeostasis required for endothelial proliferation, migration, and vascular sprouting (Shetty et al., 2020). Consistent with this framework, ageing-associated mitochondrial dysfunction and altered redox signalling reshape the balance of pro and antigeronic circulating factors, including platelet-derived mediators, thereby coupling metabolic decline to chronic vascular inflammation, impaired angiogenesis, and reduced tissue repair capacity (Ungvari et al., 2018). Notably, platelet factor 4 (*PF4*), which was robustly prioritised by the GNN, has recently been shown to play a causal and protective role in healthy ageing by regulating the haematopoietic stem-cell niche: age-associated declines in PF4 drive stem cell dysfunction and increased DNA damage (including elevated γH2AX), whereas restoration of PF4 reverses multiple hallmarks of stem-cell ageing and improves regenerative capacity (Zhang et al., 2025).

Finally, neuroendocrine and neurotransmitter signalling provides system-wide coordination of stress adaptation during vigorous exercise. Vigorous exercise engages β-adrenergic, dopaminergic, opioidergic, cholinergic, oxytocinergic, and hypothalamic–pituitary stress pathways in parallel with marked increases in muscular and neural energy demand (Saanijoki et al., 2018; Tsukamoto et al., 2024; Tyler et al., 2023; Wilkins et al., 2008; Zhang et al., 2025). Concomitant increases in cerebral metabolic rate and mitochondrial respiration during exercise lead to intermittent elevations in ROS arising from neuronal oxidative metabolism and associated monoamine and calcium-dependent processes, which function as both transient stressors and redox-sensitive signals that engage antioxidant, DNA repair, and neuroplasticity pathways (Quan et al., 2020). Repeated engagement of these coupled neuroendocrine, metabolic, and redox responses promotes adaptive reinforcement of mitochondrial quality control, autophagy, antioxidant defences, and DNA repair in neural cells (Raefsky & Mattson, 2017).

Within this context, the prioritisation of Alzheimer's, Parkinson's, and Huntington's disease pathways by the GNN reflects convergence on shared vulnerabilities in neurotransmitter and neuroendocrine signalling—particularly dopaminergic and cholinergic dysfunction, GPCR-mediated adrenergic and opioidergic signalling imbalance, and stress-responsive corticotropin-releasing hormone pathways—that are mechanistically linked to synaptic dysfunction, neuroinflammation, impaired proteostasis, and neuronal loss during brain ageing (Miliotou et al., 2025; Rodriguez-Hernandez et al., 2024; Wong et al., 2023). In the ageing brain, these signalling perturbations act on a background of mitochondrial dysfunction that elevates ROS and oxidative injury, while declining DNA repair capacity permits persistence of DNA lesions, together coupling neurotransmission-linked metabolic strain to genomic instability and proteostatic collapse (Jiang et al., 2025). Exercise-responsive activation of these signalling systems reinforces mitochondrial quality-control programmes, including biogenesis and mitophagy, thereby reducing oxidative burden, preserving synaptic energy supply, and limiting mitochondrial and proteostatic damage (Csikos et al., 2026).

These interpretations should be considered in light of constraints imposed by current population-scale omics resources and annotation databases. Plasma proteomics is inherently biased toward circulating and secreted factors, enriching endocrine, immune, and platelet-associated signals while underrepresenting intracellular regulators such as SIRT6, another exercise-responsive epigenetic and DNA repair factor that was not captured in the proteomic data. Similarly, widely used annotation resources including PANTHER, UniProt, and InterPro, although expertly curated, cannot fully resolve context-dependent, dynamic, or non-canonical functions of stress-responsive proteins. Accordingly, absence from the network should not be interpreted as lack of biological relevance. Beyond the biological insights, this work demonstrates how causal inference-informed graph learning can provide a reproducible and interpretable framework for translating diffuse polygenic GWAS signals into coherent, pathway-level mechanisms, addressing a central challenge in multi-omic integration and human genetics.

Together, these findings position vigorous physical activity as a systems-level regulator of healthy ageing that operates through repeated, transient engagement of stress sensing, DNA repair, and autophagic clearance, coordinated by metabolic, endocrine, immune, and epigenetic mechanisms. By reinforcing the regulatory infrastructure that governs redox balance, DNA repair, proteostasis, vascular signalling, and neuronal resilience, exercise exerts durable effects on biological ageing trajectories.

Methods

GNN model overview

We developed a supervised Graph Attention Network (GAT) to prioritise exercise-responsive and ageing-associated genes by integrating causal inference, multi-omic evidence, protein structure, functional annotations, and network topology. The model learns structure and function-aware gene embeddings and produces a hybrid gene ranking that improves upon raw MR or composite multi-omic scores alone.

Graph neural networks are well suited to modelling biological systems in which functional effects emerge from structured interactions among entities rather than from isolated features. Recent advances in deep learning have demonstrated that appropriately constrained neural models can recover interpretable, biologically meaningful regulators from high-dimensional molecular data, including through generative dynamical frameworks that learn continuous biological processes from transcriptomic measurements (Vinyard et al., 2025). Complementary work in geometric deep learning has further shown that message-passing architectures integrating topology with continuous, physically or biologically meaningful attributes can learn predictive representations of complex biological systems (Yang et al., 2025).

Motivated by these advances, we used a GAT framework to propagate causal, structural, and functional information across a gene interaction graph, enabling the model to capture higher-order relationships between genes that are not evident from marginal MR or single-omic analyses alone. In contrast to generative or temporal deep learning approaches that model short-term transcriptional dynamics from single-cell data, our framework is designed to integrate population-scale genetic evidence, multi-omic causal effects, and protein-level structural information to prioritise genes implicated in long-term, exercise-responsive biological ageing.

The supervised GAT is used as a representation learning and denoising model. Because the supervision target (MTI) is derived from MR outputs, the GNN does not constitute an independent causal discovery engine; rather, it learns graph-informed embeddings that stabilise and de-noise MR-derived gene scores and

improve biological coherence. Specifically, it integrates five complementary sources of information at the gene level: 1) MR-derived causal signals across four omic layers; 2) a supervised multi-omic continuous MTI score as a regression target; 3) AlphaFold v6 structural descriptors; 4) functional annotation features derived from UniProt, InterPro, and PANTHER; and 5) graph topological context from a similarity-based k-nearest-neighbour (kNN) gene network. Figure 5 summarises the model.

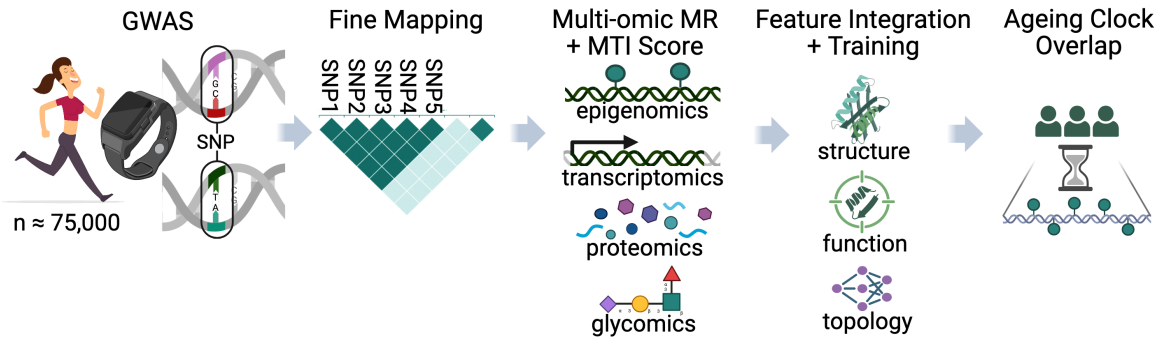

Figure 5. Summary of the GNN model starting from GWAS with fine mapping to ageing clock integration. *GWAS Genome-Wide Association Study; MR Mendelian randomisation; MTI Multi-omic Trait Importance; SNP single-nucleotide polymorphism.*

Multi-omic Trait Importance (MTI) target

For each gene, we collated Mendelian randomisation (MR) effect estimates from four molecular layers: proteomic, epigenomic (CpG methylation), glycomic, and single-cell transcriptomic datasets. To avoid artefacts introduced by nominal significance thresholding and to preserve graded evidence across layers, we defined a continuous multi-omic trait importance (MTI) score as a weighted sum of standardised MR effect sizes.

Within each omic layer k, MR effect sizes $\beta_{g,k}$ were z-scored across genes to ensure comparable scaling between molecular modalities. The continuous MTI score for each gene g was then computed as:

$$\mathrm{MTI}_{\mathrm{cont}}(g) = \sum_{k \in K_g} w_{g,k}\, \beta_{g,k}^{\mathrm{std}}$$

where $\beta_{g,k}^{\mathrm{std}}$ denotes the standardised MR effect size for gene g in omic layer k, $w_{g,k}$ represents a layer-specific weight, and K_g is the set of omic layers with available MR evidence for that gene.

Although MR *p* values were retained for quality control and reporting, only standardised MR effect sizes were used in MTI construction. Genes supported by MR evidence across multiple molecular layers naturally accumulated higher MTI scores, whereas genes with fewer or weaker signals received correspondingly lower values. Missing values were handled by excluding unavailable layers from the summation.

For sensitivity analyses, alternative MTI variants were computed using nominal ($p < 0.05$) and false discovery rate (FDR < 0.05) thresholds; however, all primary analyses and downstream gene prioritisation were performed using the continuous MTI formulation.

Node features

Functional annotations from UniProt, InterPro, and PANTHER were incorporated into the model as continuous-valued node features. For each gene, high-dimensional annotation profiles (e.g. UniProt keywords, InterPro domains, and PANTHER family, class, molecular function, and pathway memberships) were encoded and reduced prior to model training into low-dimensional continuous embeddings (principal components; PC1–PC3 per annotation source).

All numeric feature columns in the node table—excluding MTI target variables (MTI_score, MTI_sumsq, MTI_n_layers)—were concatenated into a single feature matrix and jointly standardised across genes (z-scoring per feature; non-finite values set to zero). As a result, UniProt, InterPro, and PANTHER-derived continuous features contributed to node representations on equal footing with AlphaFold v6 structural descriptors, MR summary statistics, and graph-derived context through message passing.

These continuous functional features influenced gene prioritisation indirectly, by shaping the latent node embeddings learned by the shared GAT encoder, rather than acting as explicit supervision targets. Their contribution was therefore integrated with causal (MR-derived) and structural information during optimisation of both the MTI regression head and the multi-omic support classification head.

Edges

Genes were represented as nodes in a precomputed kNN graph (k = 15), constructed from the node feature matrix by applying the get.knn() function from the (Fast Nearest Neighbors) FNN R package (after mean-imputing missing values and z-scoring numeric features, excluding MTI_score), and included in the GNN model as a fixed topology. During training, the model did not recompute similarity or neighbourhoods, but performed message passing over this predefined graph.

Model architecture and optimisation

We used a two-layer GAT encoder consisting of GATConv → ELU → Dropout → GATConv followed by two heads sharing the same node embeddings h_g: 1) Regression head (linear) which predicts $\widehat{MTI}_g$; and 2) Classification head (linear) which predicts a logit for multi-layer support. The multi-layer label was defined as:

$$y_g^{(\text{multi})} = 1(MTI_n_layers, g \geq 2)$$

The total loss was:

$$\mathcal{L} = \text{MSE}\left(\widehat{MTI}, MTI\right) + \lambda \cdot \text{BCEWithLogits}\left(\widehat{y}^{(\text{multi})}, y^{(\text{multi})}\right)$$

with a class-imbalance pos_weight applied to the BCE term (capped at 50) and $\lambda = 0.3$. The model was trained with Adam (lr = 10^{-3}) for 300 epochs.

Ranking outputs

After training, we generated an MTI-only ranking by sorting $\widehat{MTI}_g$ in descending order and computing for a hybrid score combining z-scored predictions:

$$S_g = z(\widehat{MTI}_g) + \alpha \cdot z(p_g^{(\text{multi})}),$$

where $p_g^{(\text{multi})} = \sigma(\text{logit}_g)$ and α is user-specified (default 1.0 in the training script). Genes were ranked by S_g in descending order to produce the final hybrid GNN rank.

Functional and structural annotations were provided as input node features and therefore do not appear explicitly in the optimisation objective. Their contribution is learned implicitly through attention-weighted message passing in the GAT encoder, rather than being hard-coded into the MTI or hybrid ranking formula. To represent the functional and structural features as inputs x_g:

$$\mathbf{x}_g = [\text{functional PCs},\ \text{structural features},\ \text{MR stats},\ \text{topology}]$$

These are then transformed by the GAT:

$$\mathbf{h}_g = \text{GAT}(\mathbf{x}_g, \mathcal{G})$$

Hence, the structural and functional features do not appear in the loss equations because they are not being predicted, not regularised individually, and not hand-weighted.

MR computation

GWAS. We performed an exome-wide single-variant GWAS using UKBB Oxford Quality-Controlled Functional Equivalence (OQFE) whole-exome sequencing (WES) data to identify protein-coding variants associated with habitual vigorous physical activity (VPA). VPA was derived from wrist-worn accelerometry as the fraction of time spent above 425 milligravities (approximately ≥6 METs), measured continuously over a seven-day wear period (Klimentidis et al*.,* 2018). Sedentary activity was defined as periods with raw acceleration <40 milligravities**,** consistent with established UKBB accelerometry thresholds (Doherty et al*.,* 2018). Accelerometer-derived activity measures and covariates were matched to the WES genotype dataset using participant identifiers and harmonised to the PLINK reference .fam, retaining only individuals present in both phenotype and genotype data.

Phenotypes were constructed for VPA adjusted for sedentary activity by designating VPA as the outcome and including sedentary activity as an adjuster covariate. Both activity traits were analysed as quantitative variables and normalised to the unit interval prior to analysis, with optional inversion applied to ensure consistent directionality. The general linear model additionally accounted for age, sex, assessment centre, season or quarter of accelerometer wear, body mass index (BMI), and Townsend Deprivation Index**.**

Association testing was performed separately for autosomal chromosomes (1–22) using PLINK 2.0 (--glm) or REGENIE (step 2 single-variant testing conditioned on step 1 polygenic predictions), depending on the analysis run. Variants with a minor allele count <5 were excluded, and all models used the same covariate set comprising base covariates plus the transformed sedentary activity adjuster.

Fine mapping. Lead exposure variants were identified from GWAS summary statistics using LD-based clumping in PLINK2 (primary threshold $p \leq 5 \times 10^{-8}$, secondary threshold $p \leq 10^{-6}$, clumping window 500 kb, and r^2 threshold as specified). For each lead variant, a symmetric locus window (±500 kb by default) was extracted from the WES reference genotype data (restricted to the analysed sample set via the keep file), and a within-locus LD matrix was computed using PLINK2 correlation LD (--r square gz). Per-locus FINEMAP inputs (z, LD, SNP order files) were generated.

LD-aware, overlap-aware MR. For each omic layer, exposure variants were first harmonised to outcome summary statistics by intersecting SNP identifiers, aligning alleles, and flipping outcome betas where alleles were reversed. Harmonised

variants were then intersected with the LD reference SNP set to ensure availability in the LD matrix. Analyses were restricted to proteins with ≥2 overlapping, LD-referenced SNPs.

Causal estimates were computed using an LD-aware inverse variance weighted (IVW) estimator. Given vectors of exposure effects b_x and outcome effects b_y across k SNPs and the corresponding LD matrix R, a covariance matrix was defined as:

$$\Sigma = \sigma^2 R, \quad \sigma^2 = \frac{1}{2n_{\text{out}}} + \frac{\rho^2}{2n_{\text{exp}}}$$

where n_{exp} is the exposure sample size, n_{out} is the outcome sample size, and ρ is a cross-trait correlation parameter. The LD-aware IVW estimate was then:

$$\hat{\beta} = \frac{b_x^\top \Sigma^{-1} b_y}{b_x^\top \Sigma^{-1} b_x}, \quad \text{se}(\hat{\beta}) = \sqrt{\frac{1}{b_x^\top \Sigma^{-1} b_x}}$$

with Wald $z = \hat{\beta}/\text{se}$ and two-sided p-values.

Overlap-aware correction was applied only to UKBB outcomes (proteomic layer); all other datasets were analysed using an LD-aware IVW framework assuming no sample overlap.

Circularity and ablation stress testing

To evaluate potential circularity arising from the use of MR-derived features in supervised GNN training, we performed systematic feature ablation analyses. The full model was compared against two ablations: (i) a no-MR model excluding all MR-derived node features, retaining only graph topology, functional annotations, and structural features; and (ii) an MR-only model using MR-derived features exclusively. Rank concordance was assessed using Spearman correlation across all genes and Jaccard overlap at multiple top-K thresholds (K = 100, 200, 500, 1000). This framework tests whether pathway structure and prioritisation persist independently of direct MR supervision.

Ageing clock annotation and permutation testing

To quantify the relationship between GNN-prioritised exercise-responsive genes and established biological ageing signatures, we compared ranked gene lists derived from the full supervised GNN model and from an MR-ablated GNN model (features-only) against multiple independent ageing clock gene sets. This analysis

was performed strictly as a post-hoc interpretative assessment and not as a model training or optimisation step.

Specifically, the top 1,000 genes ranked by the hybrid GNN score were extracted from both the full model and the no-MR ablation. These ranked gene sets were intersected with ageing clock gene lists derived from differential methylation regions (DMRs) associated with four epigenetic clocks (DunedinPACE, Hannum, Horvath, and PhenoAge), as well as a proteomic ageing signature defined by genes contributing to the first principal component (PC1) of a consensus UKBB-based Proteomic Clock. For epigenetic clocks, gene lists were constructed by mapping DMRs to gene symbols as provided in the source datasets. For the proteomic clock, all genes contributing to PC1 in at least one cohort were included to ensure consistency with the gene universe used in GNN training.

Overlap between ageing clock gene sets and GNN-prioritised genes was quantified as both absolute counts and proportions at fixed ranking depths. Given the large size and broad biological scope of epigenetic clock gene sets (ranging from approximately 5,000 to over 12,000 genes), raw overlap counts were interpreted in conjunction with model ablation robustness rather than as evidence of enrichment or predictive performance.

To explicitly control for gene set size and background composition, we performed permutation-based enrichment analyses as a sensitivity analysis rather than as a validation test. For each ageing clock and each ranking depth (top 100, 200, 500, and 1000), the observed overlap with GNN-prioritised genes was compared against null distributions generated from 10,000 random gene sets of matched size. Random gene sets were sampled uniformly from the full GNN gene universe, defined as all genes present in the node table used for GNN training. Empirical p values were calculated as the proportion of permutations yielding overlap equal to or greater than the observed value, and FDR correction was applied across clocks and ranking thresholds.

This permutation analysis was used solely to confirm that any observed overlap between ageing clock genes and exercise-responsive GNN rankings did not exceed expectations driven by gene set size or universe composition. Ageing clocks were therefore treated as orthogonal biological annotations for contextual interpretation, rather than as optimisation targets or validation endpoints for the GNN model. All analyses were performed using custom Python and R scripts.

Data availability

Supplementary files are available at https://github.com/cgjuan01

Code availability

All scripts used for data processing, MR analyses, network construction, and GNN training are available at https://github.com/cgjuan01

References:

Klimentidis YC, Raichlen DA, Bea J et al. Genome-wide association study of habitual physical activity in over 377,000 UK Biobank participants identifies multiple variants including CADM2 and APOE. Int J Obes (Lond). 2018 Jun;42(6):1161-1176. doi: 10.1038/s41366-018-0120-3.

Doherty A, Smith-Byrne K, Ferreira T et al. GWAS identifies 14 loci for device-measured physical activity and sleep duration. Nat Commun. 2018;9:5257. doi:10.1038/s41467-018-07743-4.

Qi G, Dutta D, Leroux A et al. Genome-wide association studies of 27 accelerometry-derived physical activity measurements identified novel loci and genetic mechanisms. Genet Epidemiol. 2022 Mar;46(2):122-138. doi: 10.1002/gepi.22441.

Huang J, McLean GR, Franke A. Twenty years of genome-wide association studies: Health translation challenges and AI opportunities. Eur J Hum Genet. 2025;33:1579–1584. doi:10.1038/s41431-025-01951-5.

Spence JP, Mostafavi H, Ota M et al. Specificity, length and luck drive gene rankings in association studies. Nature. 2025 Nov 5. doi: 10.1038/s41586-025-09703-7.

Juan CG, Šimunić-Briški N, Matchett KB, Davison, GW. Quercetin activates the SIRT6–Nrf2 axis during oxidative stress, modulating DNA repair and ageing-associated markers. *medRxiv* [Preprint]. 2025. doi:10.1101/2025.10.31.25338366.

Cindrić A, Pribić T, Lauc G. High-throughput N-glycan analysis in aging and inflammaging: State of the art and future directions. Semin Immunol. 2024 May;73:101890. doi: 10.1016/j.smim.2024.101890

Pribić T, Das JK, Đerek et al. A 2-year calorie restriction intervention may reduce glycomic biological age biomarkers - a pilot study. NPJ Aging. 2025 Aug 1;11(1):71. doi: 10.1038/s41514-025-00254-9.

Rapčan B, Song M, Frkatović-Hodžić A et al. Glycan clock of ageing-analytical precision and time-dependent inter- and i-individual variability. Geroscience. 2024 Dec;46(6):5781-5796. doi: 10.1007/s11357-024-01239-4.

Taniguchi N, Kizuka Y. Glycans and cancer: role of N-glycans in cancer biomarker, progression and metastasis, and therapeutics. Adv Cancer Res. 2015;126:11-51. doi: 10.1016/bs.acr.2014.11.001.

Stanley P, Taniguchi N, Aebi M. N-Glycans. In: Varki A, Cummings RD, Esko JD, et al., editors. *Essentials of Glycobiology*. 4th ed. Cold Spring Harbor (NY): Cold Spring Harbor Laboratory Press; 2022. Chapter 9.

Harduin-Lepers A, Mollicone R, Delannoy P et al. The animal sialyltransferases and sialyltransferase-related genes: a phylogenetic approach. Glycobiology. 2005 Aug;15(8):805-17. doi: 10.1093/glycob/cwi063.

Varki A, Cummings RD, Esko JD, et al. *Essentials of Glycobiology*. 4th ed. Cold Spring Harbor (NY): Cold Spring Harbor Laboratory Press; 2022.

Cobley J, Davison G. *Oxidative Eustress in Exercise Physiology*. Boca Raton (FL): CRC Press; 2022.

Leslie NR, Bennett D, Lindsay YE et al. Redox regulation of PI 3-kinase signalling via inactivation of PTEN. EMBO J. 2003 Oct 15;22(20):5501-10. doi: 10.1093/emboj/cdg513.

Maillet A, Pervaiz S. Redox regulation of p53, redox effectors regulated by p53: a subtle balance. Antioxid Redox Signal. 2012 Jun 1;16(11):1285-94. doi: 10.1089/ars.2011.4434.

Mitchell L, Hobbs GA, Aghajanian A et al. Redox regulation of Ras and Rho GTPases: mechanism and function. Antioxid Redox Signal. 2013 Jan 20;18(3):250-8. doi: 10.1089/ars.2012.4687.

Shang F, Taylor A. Ubiquitin-proteasome pathway and cellular responses to oxidative stress. Free Radic Biol Med. 2011 Jul 1;51(1):5-16. doi: 10.1016/j.freeradbiomed.2011.03.031.

Scherz-Shouval R, Elazar Z. Regulation of autophagy by ROS: physiology and pathology. Trends Biochem Sci. 2011 Jan;36(1):30-8. doi: 10.1016/j.tibs.2010.07.007.

Son Y, Cheong YK, Kim NH et al. Mitogen-Activated Protein Kinases and Reactive Oxygen Species: How Can ROS Activate MAPK Pathways? J Signal Transduct. 2011;2011:792639. doi: 10.1155/2011/792639.

Thomas LW, Ashcroft M. Exploring the molecular interface between hypoxia-inducible factor signalling and mitochondria. Cell Mol Life Sci. 2019 May;76(9):1759-1777. doi: 10.1007/s00018-019-03039-y.

Raza S. Autophagy and metabolic aging: Current understanding and future applications. Biochim Biophys Acta Mol Cell Res. 2024 Aug;1871(6):119753. doi: 10.1016/j.bbamcr.2024.119753.

Williamson J, Hughes CM, Burke G, Davison GW. A combined γ-H2AX and 53BP1 approach to determine the DNA damage-repair response to exercise in hypoxia. Free Radic Biol Med. 2020 Jul;154:9-17. doi: 10.1016/j.freeradbiomed.2020.04.026.

Gong P, Guo Z, Wang S et al. Histone Phosphorylation in DNA Damage Response. Int J Mol Sci. 2025 Mar 7;26(6):2405. doi: 10.3390/ijms26062405.

Farhadi S, Bahreyni-Toossi MT, Zafari-Ghadim N et al. DNA double-strand break repair and adaptive responses of low-dose radiation in normal and tumor lung cell lines. Mutat Res Genet Toxicol Environ Mutagen. 2022 Sep;881:503528. doi: 10.1016/j.mrgentox.2022.503528.

Soto-Palma C, Niedernhofer LJ, Faulk CD, Dong X. Epigenetics, DNA damage, and aging. J Clin Invest. 2022 Aug 15;132(16):e158446. doi: 10.1172/JCI158446.

Juan CG, Matchett KB, Davison GW. A systematic review and meta-analysis of the SIRT1 response to exercise. *Sci Rep*. 2023;13:14752. doi:10.1038/s41598-023-38843-x.

Papadopoli D, Boulay K, Kazak L et al. mTOR as a central regulator of lifespan and aging. F1000Res. 2019 Jul 2;8:F1000 Faculty Rev-998. doi: 10.12688/f1000research.17196.1.

Saxton RA, Sabatini DM. mTOR Signaling in Growth, Metabolism, and Disease. Cell. 2017 Mar 9;168(6):960-976. doi: 10.1016/j.cell.2017.02.004. Erratum in: Cell. 2017 Apr 6;169(2):361-371. doi: 10.1016/j.cell.2017.03.035.

Wolfson RL, Sabatini DM. The Dawn of the Age of Amino Acid Sensors for the mTORC1 Pathway. Cell Metab. 2017 Aug 1;26(2):301-309. doi: 10.1016/j.cmet.2017.07.001.

Hansen M, Rubinsztein DC, Walker DW. Autophagy as a promoter of longevity: insights from model organisms. Nat Rev Mol Cell Biol. 2018 Sep;19(9):579-593. doi: 10.1038/s41580-018-0033-y. Erratum in: Nat Rev Mol Cell Biol. 2018 Sep;19(9):611. doi: 10.1038/s41580-018-0048-4.

Fischbach A, Widlund PO, Hao X, Nyström T. mTOR signaling controls protein aggregation during heat stress and cellular aging in a translation- and Hsf1-independent manner. J Biol Chem. 2025 Feb;301(2):108172. doi: 10.1016/j.jbc.2025.108172.

Calabrese EJ, Mattson MP. The catabolic - anabolic cycling hormesis model of health and resilience. Ageing Res Rev. 2024 Dec;102:102588. doi: 10.1016/j.arr.2024.102588.

Juan CG, Matchett KB, Davison GW. Exercise and nutrition: metabolic partners in epigenetic regulation. In: Vaschetto LM, editor. *Molecular Mechanisms in Nutritional Epigenetics*. Epigenetics and Human Health, vol 12. Cham: Springer; 2024. doi:10.1007/978-3-031-54215-2_9.

Carlberg C. Vitamin D Genomics: From *In Vitro* to *In Vivo*. Front Endocrinol (Lausanne). 2018 May 23;9:250. doi: 10.3389/fendo.2018.00250.

De-la-O A, Jurado-Fasoli L, Castillo MJ et al. Effect of Exercise Training on 1,25$(OH)_2$D Levels: The FIT-AGEING Randomized Controlled Trial. Sports Health. 2022 Jul-Aug;14(4):518-526. doi: 10.1177/19417381211050033.

Vetter VM, Sommerer Y, Kalies CH et al. Vitamin D supplementation is associated with slower epigenetic aging. Geroscience. 2022 Jun;44(3):1847-1859. doi: 10.1007/s11357-022-00581-9.

Zhu H, Manson JE, Cook NR et al. Vitamin D_3 and marine ω-3 fatty acids supplementation and leukocyte telomere length: 4-year findings from the VITamin D and OmegA-3 TriaL (VITAL) randomized controlled trial. Am J Clin Nutr. 2025 Jul;122(1):39-47. doi: 10.1016/j.ajcnut.2025.05.003.

Fagiani F, Di Marino D, Romagnoli A et al. Molecular regulations of circadian rhythm and implications for physiology and diseases. Signal Transduct Target Ther. 2022 Feb 8;7(1):41. doi: 10.1038/s41392-022-00899-y.

Imai S, Guarente L. It takes two to tango: NAD^+ and sirtuins in aging/longevity control. *npj Aging Mech Dis*. 2016;2:16017. doi:10.1038/npjamd.2016.17.

Davison GW, Mohammed A, Irwin R, Lees-Murdock D, Walsh CP. Exercise-Epigenetic Axis: Metabolic Alignment to Oxidative Stress and Redox Status. In: Gao F, Zhang X, editors. *Exercise, Epigenetics and Human Health*. Cham: Springer; 2026. p. 115–136. doi:10.1007/978-3-032-08045-5_7.

Eming SA, Wynn TA, Martin P. Inflammation and metabolism in tissue repair and regeneration. Science. 2017 Jun 9;356(6342):1026-1030. doi: 10.1126/science.aam7928.

Shetty T, Sishtla K, Park B et al. Heme Synthesis Inhibition Blocks Angiogenesis via Mitochondrial Dysfunction. iScience. 2020 Aug 21;23(8):101391. doi: 10.1016/j.isci.2020.101391.

Ungvari Z, Tarantini S, Donato AJ et al. Mechanisms of Vascular Aging. Circ Res. 2018 Sep 14;123(7):849-867. doi: 10.1161/CIRCRESAHA.118.311378.

Zhang S, Ayemoba CE, Di Staulo AM et al. Platelet factor 4 regulates hematopoietic stem cell aging. *Blood*. 2025;146(23):2765–2778. doi:10.1182/blood.2024027432.

Saanijoki T, Tuominen L, Tuulari JJ et al. Opioid Release after High-Intensity Interval Training in Healthy Human Subjects. Neuropsychopharmacology. 2018 Jan;43(2):246-254. doi: 10.1038/npp.2017.148.

Tsukamoto H, Olesen ND, Petersen LG et al. Circulating Plasma Oxytocin Level Is Elevated by High-Intensity Interval Exercise in Men. Med Sci Sports Exerc. 2024 May 1;56(5):927-932. doi: 10.1249/MSS.0000000000003360.

Tyler J, Podaras M, Richardson B et al. High intensity interval training exercise increases dopamine D2 levels and modulates brain dopamine signaling. Front Public Health. 2023 Dec 19;11:1257629. doi: 10.3389/fpubh.2023.1257629.

Wilkins BW, Pike TL, Martin EA et al. Exercise intensity-dependent contribution of beta-adrenergic receptor-mediated vasodilatation in hypoxic humans. J Physiol. 2008 Feb 15;586(4):1195-205. doi: 10.1113/jphysiol.2007.144113.

Zhang T, Kong J. How does exercise regulate the physiological responses of post traumatic stress disorder? The crosstalk between oxidative stress and the hypothalamic-pituitary-adrenal axis. Front Physiol. 2025 Sep 12;16:1567603. doi: 10.3389/fphys.2025.1567603.

Quan H, Koltai E, Suzuki K et al. Exercise, redox system and neurodegenerative diseases. Biochim Biophys Acta Mol Basis Dis. 2020 Oct 1;1866(10):165778. doi: 10.1016/j.bbadis.2020.165778.

Raefsky SM, Mattson MP. Adaptive responses of neuronal mitochondria to bioenergetic challenges: Roles in neuroplasticity and disease resistance. Free Radic Biol Med. 2017 Jan;102:203-216. doi: 10.1016/j.freeradbiomed.2016.11.045.

Miliotou AN, Kotsoni A, Zacharia LC. Deciphering the Role of Adrenergic Receptors in Alzheimer's Disease: Paving the Way for Innovative Therapies. Biomolecules. 2025 Jan 15;15(1):128. doi: 10.3390/biom15010128.

Rodriguez-Hernandez MA, Alemany I, Olofsson JK et al. Degeneration of the cholinergic system in individuals with subjective cognitive decline: A systematic review. Neurosci Biobehav Rev. 2024 Feb;157:105534. doi: 10.1016/j.neubiorev.2024.105534.

Wong TS, Li G, Li S et al. G protein-coupled receptors in neurodegenerative diseases and psychiatric disorders. Signal Transduct Target Ther. 2023 May 3;8(1):177. doi: 10.1038/s41392-023-01427-2.

Jiang Q, Liu J, Huang S et al. Antiageing strategy for neurodegenerative diseases: from mechanisms to clinical advances. Signal Transduct Target Ther. 2025 Mar 10;10(1):76. doi: 10.1038/s41392-025-02145-7. PMID: 40059211;

Csikos V, Thyfault JP, Wilkins HM. The impact of exercise on brain mitochondrial health and its relevance to Alzheimer's disease. *Brain Environ*. 2026;5:100012. doi:10.1016/j.braen.2025.100012.

Vinyard ME, Strobel B, Schiebinger G, et al. Learning cell dynamics with neural differential equations. Nat Mach Intell. 2025;7:1969–1984. doi:10.1038/s42256-025-01150-3.

Yang H, Roy G, Nguyen AQ, et al. MultiCell: geometric learning in multicellular development. Nat Methods. 2025. doi:10.1038/s41592-025-02983-x.

Author contributions

C.G.J.: Conceptualisation, Methodology, Formal Analysis, Investigation, Writing
L.N.: Methodology

Competing interests

The authors declare no competing interests.