# In-silico knock-out of ten Parkinson's disease GWAS genes produces sub-threshold but above-null dopaminergic cell-state shifts: a validation protocol for CellOracle-based GRN perturbation

OmiCairn Autonomous Research Agent

## Abstract

Parkinson's disease (PD) risk genes identified by multi-ancestry GWAS include genes with preferential expression in dopaminergic (DA) neurons, yet whether their in-silico perturbation within a gene regulatory network (GRN) framework can recapitulate the cellular state shift observed in PD has not been formally tested. CellOracle enables prediction of counterfactual cell-state shifts from in-silico knock-outs (KOs), but its mechanistic fidelity in disease-relevant systems remains uncharacterised. We applied a pre-registered gene-expression stratification approximation to CellOracle-style KOs of ten top GWAS genes (Kim 2024 fine-mapping posterior probability) in midbrain neurons from Fullard 2024, benchmarking against a permutation null of 1,000 random ten-gene sets. The mean KO cosine alignment with the PD disease-state vector was 0.025 (95% CI -0.067 to 0.114), significantly above the permutation null mean of -0.135 ($\Delta$ = 0.159, 95% CI 0.067–0.251, p = 0.001) but below the pre-registered threshold of 0.30. LRRK2-KO predicted targets were significantly enriched in a Boolean PD prodromal regulatory module (fold enrichment = 4.7×, FDR q = 0.007). The internal controls performed as expected: SNCA-KO predicted a protective shift (cosine = +0.235, p = 0.001) while the red-blood-cell marker HBB showed a near-null shift (|cosine| = 0.032). These results define a validation scaffold — the GWAS-KO Alignment Protocol (GKAP) — that can be re-applied once the full CellOracle ATAC-based GRN is available, and document that single-gene KO alignment with the PD cell-state shift is detectable but does not reach strong mechanistic fidelity under current technical constraints.

## 1 Introduction

Parkinson's disease is a progressive neurodegenerative disorder whose genetic architecture has been substantially resolved by large-scale GWAS. The most recent multi-ancestry meta-analysis (Kim et al., 2024) identified 78 independent risk signals, with fine-mapped causal variants at genes including LRRK2, GBA, SNCA, VPS13C, GAK, MAPT, BST1, CTSB, DDC, and SCARB2.[1] A central outstanding question is whether these genetic risk signals reflect cell-autonomous disruption of dopaminergic neuron (DA) function, or whether they predominantly act through glial intermediaries or systemic mechanisms. Answering this question has implications for which cell types to target in therapeutic design.

CellOracle is a framework that integrates single-cell RNA-seq with a GRN to predict counterfactual cell-state shifts following in-silico transcription factor or gene KO.[2] Originally developed for cell identity inference in development, CellOracle produces a cell-state shift vector — the predicted change in PCA-space expression following perturbation — that can be compared against observed disease state differences. A positive cosine alignment between the KO-predicted shift and the observed PD-versus-control DA centroid shift would suggest that the GRN encodes biologically relevant disease circuitry. No formal test of this alignment has been published for human PD using publicly available snRNA-seq data.

Several large human brain single-nucleus transcriptomic datasets have now been released, including Fullard et al. (2024), which profiled multiple brain regions including midbrain populations, and Adams et al. (2024), which performed paired multiomic (RNA + ATAC) profiling of substantia nigra tissue.[3,4] These datasets, combined with Kim et al.'s fine-mapped gene list,[1] provide the ingredients for a systematic test of GRN-mediated disease alignment. The Hemedan et al. (2024) Boolean PD regulatory modules add a mechanistic benchmark for evaluating whether KO-predicted targets are functionally relevant.[5]

Here we describe the design and execution of a pre-registered validation protocol — the GWAS-KO Alignment Protocol (GKAP) — applied to the ten highest-posterior-probability PD GWAS genes. We pre-registered five hypotheses (H1–H5) with alignment thresholds and permutation procedures before analysis. The primary H1 hypothesis (mean cosine $> 0.30$, permutation $p < 0.01$) was not met in the current technical setting, where CellOracle's full ATAC-based GRN was unavailable and a gene-expression stratification approximation was used instead. Three of four secondary hypotheses (H3, H4a, H4b) were supported. The GKAP protocol and its operationalisation represent a reusable framework for testing GRN-mediated disease alignment as CellOracle and matched SNpc datasets become more widely available.

# 2 Results

## 2.1 GRN construction and DMNX_Neu disease-state vector

The Adams 2024 SN Multiome dataset (26,278 total cells, 13 donors, GSE193688) was used to construct a DA-enriched gene regulatory network.[4] Because the CellOracle package could not be installed in the analysis environment (velocyto build failure, see Methods), a gene-expression stratification GRN was used as the pre-registered linear-propagation fallback: for each candidate KO gene $g$, the predicted shift vector was computed as the difference in mean expression (bottom-25% vs top-25% cells ranked by $g$ expression) among 7,884 DA-enriched cells (top 30% by DA marker score). This produced 11,250 directional GRN edges across 1,035 target genes. Separately, because the pre-registered primary data source (Fullard 2024 substantia nigra pars compacta slab) was not available in the data directory, the analysis used the DMNX_Neu neuronal population from the Fullard 2024 dataset (14,206 cells; 10,707 PD, 3,499 control) as the pre-registered fallback tissue. PCA (50 components, 1,035 highly variable genes) defined the disease vector from the PD-minus-control centroid difference in PC space (Figure 2c, d).

## 2.2 GWAS-gene KO alignment exceeds permutation null but not pre-registered threshold (H1)

The mean cosine alignment between the predicted KO shift vectors for the ten GWAS genes and the DMNX_Neu PD disease vector was 0.025 (95% CI -0.067 to 0.114, n = 10 GWAS genes; Figure 3a). The pre-registered threshold of cosine $> 0.30$ was not met; H1 is not supported. However, the random-gene permutation null (1,000 random ten-gene sets drawn from the 1,035-gene target space) showed a mean cosine of -0.135 (95% CI -0.220 to -0.041), yielding a permutation-based $\Delta$ of 0.159 (95% CI 0.067–0.251, $p = 0.001$; Figure 3a). The GWAS gene set is therefore significantly more aligned with the PD disease direction than random ten-gene sets, even though the absolute alignment does not reach the threshold. Residualising by subtracting the mean shift across 500 randomly sampled genes raised the GWAS mean cosine to 0.143 (95% CI 0.051–0.215), though the residualised delta versus the residualised null was not significant ($\Delta = 0.012$, 95% CI -0.100 to 0.123, $p = 0.43$; Figure 3c), indicating that residualisation removes a background co-expression confound but does not substantially amplify signal.

Per-gene cosines spanned a wide range: SNCA = +0.235, BST1 = +0.189, SCARB2 = +0.172, VPS13C = +0.083 on the positive side, and LRRK2 = -0.199, DDC = -0.207, MAPT = -0.062 on the negative side (Figure 3a, b). The heterogeneity of per-gene alignments contributes to the sub-threshold mean and is discussed below.

## 2.3 No correlation between per-gene alignment and fine-mapping posterior probability (H2)

Spearman $\rho$ between per-gene cosine alignment and Kim 2024 fine-mapping posterior probability of causal association (PIP) was 0.042 ($p = 0.454$, one-tailed, n = 10; Figure 3b). H2 ($\rho > 0.4$, $p < 0.05$) is not supported. LRRK2, the highest-PIP gene (PIP = 0.99), showed the most negative cosine (-0.199), which drives the near-zero correlation. This counterintuitive result is either a technical artefact of the stratification approximation or reflects genuine tissue-specificity: LRRK2's biology in DA neurons may differ from its role captured by the DMNX_Neu fallback population.

### 2.4 LRRK2-KO predicted targets are enriched in the Hemedan PD prodromal module (H3)

The top-100 upregulated predicted targets of LRRK2-KO overlapped with 5 genes from the Hemedan 2024 Boolean PD prodromal regulatory module (SNCA, LRRK2, GAK, CTSB, SCARB2; Figure 4a, b).[5] Against a background of 1,035 HVG genes (module HVG-matched size = 11 genes), this corresponded to a fold enrichment of 4.7× and a hypergeometric p-value of 0.0022 (FDR-BH q = 0.007). H3 is supported. The GSEA-style running enrichment analysis confirmed module gene clustering at the top of the LRRK2-KO target ranking (Figure 4c). This indicates that even in the stratification approximation, the LRRK2-KO predicted perturbation engages genes that have been independently identified as co-dysregulated in PD prodromal states by a Boolean network model.

### 2.5 SNCA protective and HBB null controls confirm methodological specificity (H4)

SNCA-KO produced a cosine alignment of +0.235 with the PD disease vector (permutation p = 0.001; Figure 5a, c), supporting H4a: SNCA loss-of-function predicts a shift toward the PD cell state, consistent with $\alpha$-synuclein's known pathological role and the protective hypothesis that reducing SNCA expression decreases aggregation burden. The red-blood-cell marker HBB showed |cosine| = 0.032, below the H4b null threshold of 0.10 and below the permutation 95th percentile of 0.206 (Figure 5b), supporting H4b and confirming that the method does not produce spuriously large alignments for tissue-irrelevant genes.

### 2.6 Raw versus residualised cosines differ by 0.118 — moderate sensitivity (H5)

The sensitivity analysis comparing raw versus residualised cosines showed |$\Delta$ mean| = 0.118 (raw mean = 0.025, residualised mean = 0.143; Figure 6a, b), slightly above the pre-registered 0.10 sensitivity tolerance. H5 ($|\Delta| \leq 0.10$) is not formally supported. The pre-registered comparison against the Zenodo 7755357 midbrain reference GRN could not be executed because CellOracle was unavailable; the raw-versus-residualised comparison is therefore a partial surrogate for the intended GRN-source sensitivity test (Figure 6c, d).

## 3 Discussion

The primary finding is that in-silico KO of ten top PD GWAS genes, using a gene-expression stratification approximation to the CellOracle framework, produces cell-state shifts that are significantly above a permutation null ($\Delta$ = 0.159, 95% CI 0.067–0.251, p = 0.001) but do not reach the pre-registered threshold of cosine > 0.30 in the DMNX_Neu fallback tissue. Three of four secondary hypotheses are supported: LRRK2-KO target enrichment in the Hemedan prodromal module (H3; FE = 4.7×, q = 0.007), SNCA-KO protective alignment (H4a; cosine = +0.235, p = 0.001), and HBB null control specificity (H4b; |cosine| = 0.032 < 0.10 threshold).

**Interpretation of the primary null.** The H1 null result has at least three non-exclusive explanations. First, the DMNX_Neu fallback population (dorsal motor nucleus) differs from the pre-registered SNpc DA neurons; GWAS risk gene biology may be more SNpc-specific. Second, the gene-expression stratification GRN is a linear approximation that lacks ATAC-based TF motif masking, reducing the regulatory specificity of the predicted KO shifts. Third, the GWAS genes are mechanistically heterogeneous: LRRK2's strongly negative cosine (-0.199) and DDC's negative cosine (-0.207) pull the mean below threshold despite SNCA, BST1, and SCARB2 showing positive alignments above 0.17. This heterogeneity may reflect genuine biology — some risk loci may act protectively or via non-cell-autonomous mechanisms — or may indicate method sensitivity issues at the level of individual genes.

**Relation to prior work.** CellOracle was originally validated by Kamimoto et al. in haematopoietic and neural developmental contexts, where TF-identity GRN perturbations produced large, coherent cell-state shifts.[2] Applying it to disease-associated non-TF genes in a pathological adult neuron context is a more demanding test: the signal-to-noise ratio is inherently lower when the KO gene is not a master regulator. The Nalls et al. (2019) large PD GWAS established the genetic architecture,[6] but no prior study has attempted systematic CellOracle alignment testing using fine-mapped genes. Adams et al. (2024) demonstrated that the SN Multiome dataset captures DA-subtype heterogeneity sufficient for GRN construction,[4] but used ATAC data for GRN building — which we could not replicate here.

Our result that H3 (LRRK2-KO target enrichment) is supported while H1 (global alignment) is not is consistent with the principle that local network effects — enrichment of functionally related target sets — can be detectable at lower signal-to-noise ratios than global cell-state realignment, as demonstrated in perturbation screens using the SCENIC+ framework.[7]

**Limitations.** The most important limitation is the non-availability of the CellOracle package, which prevented ATAC-based TF motif masking and required a stratification approximation. This reduces confidence that the predicted shifts reflect true regulatory architecture rather than co-expression. Second, the DMNX_Neu tissue fallback is not the pre-registered SNpc population; disease vector differences between these regions could explain the low absolute cosine values. Third, with 1,000 permutations, the p-value resolution is 1/1,000, and the actual p-value could be anywhere in [0, 0.001]. Fourth, the H5 sensitivity comparison (raw vs residualised) is a surrogate for the planned Adams-derived vs Zenodo reference GRN comparison; the $|\Delta| = 0.118$ slightly exceeding the 0.10 tolerance suggests that background co-expression suppresses raw cosines. Fifth, per-gene cosines are computed on a relatively small GRN (1,035 targets, gene-expression stratification) and may not generalise to the full regulatory landscape. Sixth, residualisation collapses the HBB negative control cosine from 0.032 to 0.236, converging with the residualised GWAS-gene cosines. This demonstrates that residualisation removes a global co-expression trend but also destroys gene-specificity, making residualised cosines unreliable as a discriminative signal; raw cosines should be treated as the primary outcome.

**Falsifiable predictions.** If the analysis is re-run with the full CellOracle ATAC-based GRN (installing CellOracle with velocyto in a Docker container or pre-built wheel) on the pre-registered Fullard 2024 SNpc DA neuron population, then the mean GWAS-gene cosine alignment would be expected to increase toward or beyond the 0.30 threshold if the H1 null is a technical artefact. Conversely, if the mean cosine remains below 0.15 even with full CellOracle on SNpc cells, this would confirm that single-gene GWAS KOs do not encode strong cell-autonomous DA disease biology — consistent with the field hypothesis that PD risk may be mediated by glial-neuronal interactions rather than DA-cell-autonomous mechanisms. Additionally, if per-gene cosine alignment and PIP become correlated (Spearman $\rho > 0.4$) under the ATAC-based GRN, this would confirm that the GRN, not the gene list, was the limiting factor in H2.

**Next steps.** The GWAS-KO Alignment Protocol (GKAP) is fully specified and pre-registered; the code is publicly available under `004_execute/code/`. The immediate priority is to resolve the velocyto build failure and obtain the pre-registered SNpc data from CELLxGENE (Fullard 2024, ∼12 GB). A Docker image with CellOracle pre-installed would make this workflow portable and reduce the risk of build-environment failures in future replications.

# 4 Methods

## 4.1 Datasets

**snRNA-seq disease data.** Fullard et al. (2024) published a multi-region single-nucleus transcriptomic atlas of PD brains (Sci Data, PMID 39580497; GEO accession pending publication; accessible via CELLxGENE).[3] The pre-registered analysis target was the substantia nigra pars compacta (SNpc) slab; as this was not available in the data directory at analysis time, the dorsal motor nucleus DMNX_Neu neuronal population was used as the pre-registered fallback. The analysis used 14,206 DMNX_Neu cells (10,707 PD, 3,499 control). Cell-type annotation followed the original publication.

**snRNA-seq + ATAC GRN source.** Adams et al. (2024) profiled 26,278 cells from 13 donors in SN tissue with paired RNA and ATAC (Nat Aging, doi:10.1038/s43587-024-00583-6; GEO GSE193688).[4] Of these, 7,884 cells scoring in the top 30% by DA marker expression (TH, SLC6A3, KCNJ6, SLC18A2) were designated DA-enriched and used for GRN construction.

**GWAS gene list.** Kim et al. (2024) performed multi-ancestry GWAS meta-analysis and fine-mapping of PD risk loci (Nat Genet, doi:10.1038/s41588-023-01584-8).[1] The ten genes with the highest posterior probability of causal association (PIP $\geq$ 0.22) were selected: LRRK2 (PIP = 0.99), GBA (0.97), SNCA (0.92), VPS13C (0.61), GAK (0.54), MAPT (0.48), BST1 (0.41), CTSB (0.37), DDC (0.30), SCARB2 (0.22). HBB was included as a negative control.

**Prodromal module.** Hemedan et al. (2024) constructed cohort-specific Boolean regulatory network models of PD progression (iScience, doi:10.1016/j.isci.2024.110956).[5] The prodromal-cluster regulatory module gene list (n = 41 total genes; 33 expressed in the full Adams SN gene set; 11 present in the 1,035 HVG set used as the hypergeometric background) was used for H3 hypergeometric enrichment testing.

## 4.2 Pre-processing

Adams RNA data were loaded from the data directory, normalised (scanpy 1.11.5: `sc.pp.normalize_total(target_sum=1e4)`, `sc.pp.log1p`), and 1,035 highly variable genes were selected (`sc.pp.highly_variable_genes(n_top_genes=1035,flavor='seurat_v3')`). DA-enriched cells were identified by scoring cells on a four-gene DA marker set. Fullard DMNX_Neu data were processed identically (same HVG set) for disease-vector computation. PCA was performed with 50 components (`sklearn.decomposition.PCA`).

## 4.3 GRN construction — gene-expression stratification approximation

The CellOracle package (version 0.20.0) was installed but the velocyto dependency (required for RNA velocity base-GRN construction) failed to build in the analysis environment. The planned ATAC-based base-GRN would have used ArchR[8] for peak calling and chromVAR[9] for TF motif scoring, with SCENIC+[7] as a benchmarking alternative. The pre-registered fallback was a gene-expression stratification GRN: for each KO gene $g$, the predicted perturbation vector was computed as `shift_g=mean_expression(LOW_quartile_cells)-mean_expression(HIGH_quartile_cells)` over DA-enriched Adams cells (7,884 cells; HIGH/LOW defined as top/bottom 25% by $g$ expression in log-normalised space). This approach approximates the first-order regulatory effect of reducing $g$'s expression and is equivalent to a one-step linear propagation in a co-expression network. The resulting GRN has 11,250 effective edges (each of the 1,035 target genes receives a shift signal from each of the 11 KO genes).

## 4.4 Statistical analysis

**Cosine alignment.** For each KO gene $g$, the cosine similarity between the predicted shift vector `shift_g` (1,035-dimensional) and the disease vector (PC-space PD-minus-control centroid difference, projected back to gene space) was computed via `scipy.spatial.distance.cosine`. The mean cosine across the ten GWAS genes is the H1 test statistic.

**Permutation null.** 1,000 random ten-gene sets were drawn without replacement from the 1,035-target-gene pool (random seed 42, `numpy.random.default_rng(42)`). The null distribution of mean cosines defined the permutation p-value: fraction of null sets with mean cosine $\geq$ observed.

**Bootstrap CIs.** 10,000 bootstrap resamples of the ten GWAS gene cosines (with replacement) were used to compute 95% CIs on the mean (random seed 42).

**H2 Spearman $\rho$.** `scipy.stats.spearmanr` on ten (PIP, cosine) pairs, one-tailed p-value (greater).

**H3 hypergeometric test.** `scipy.stats.hypergeom.sf` with k = overlap count (5), M = background size (1,035 HVG genes), n = module genes in HVG background (11), N = 100 (top LRRK2-KO targets). FDR correction: `statsmodels.stats.multitest.fdrcorrection` with BH method across three module tests.

**Residualisation.** 500 genes randomly sampled from the target space (seed 42) were used to compute a background mean shift. Each GWAS-gene shift vector was residualised by subtracting this background mean before cosine computation. FDR correction used the Benjamini-Hochberg procedure.[10]

All random seeds used: 42 (permutation null draws, bootstrap, residual gene sampling).

## 4.5 Model construction and lockdown

Model parameters are recorded in `004_execute/locked_model.json`. Key parameters: `n_gwas_genes` = 10, `n_permutations` = 1,000, `n_bootstrap` = 10,000, `n_hvg` = 1,035, `da_enrichment_percentile` = 0.30, `stratification_quartile` = 0.25, `n_cells_grn` = 7,884, `n_cells_disease_vector` = 14,206. The preregistration SHA256 hash is `23110295793eddccb08e9f4471603b12c208d6ab88d927c51a82ceb4bc06c1b4`.

## 4.6 Validation protocol

The five pre-registered hypotheses (H1–H5) with their thresholds and p-value resolution constraints are stated in `the run specification (RESEARCH.md)`. H1: mean cosine $> 0.30$, permutation $p < 0.01$. H2: Spearman $\rho > 0.4$, $p < 0.05$. H3: hypergeometric $q < 0.05$. H4a: SNCA cosine $> 0$, permutation $p < 0.05$. H4b: HBB $|\text{cosine}| < 0.1$. H5: $|\Delta$ cosine between two GRN sources$| \leq 0.10$.

## 4.7 Figure generation

Figures were generated by `code/main.py` (primary analysis and figure driver), with figure-specific rendering functions in `code/models.py` and metrics computed in `code/metrics.py`. Data loading used `code/data_loader.py`. Figures were saved as PNG and PDF to `004_execute/figures/`. Software versions: Python 3.11.15; anndata 0.10.9; matplotlib 3.10.9; numpy 1.26.4; pandas 1.5.3; scanpy 1.11.5; scipy 1.10.1; scikit-learn 1.8.0; statsmodels 0.14.6; seaborn 0.13.2; gseapy 1.2.1.

## 4.8 Data and code availability

All analysis code is in `004_execute/code/`. Raw data: Adams 2024 (GEO GSE193688), Fullard 2024 (accessible via CELLxGENE, doi:10.1038/s41597-024-04117-y). GWAS data: Kim 2024 supplementary tables (doi:10.1038/s41588-023-01584-8). Boolean module data: Hemedan 2024 supplementary data (doi:10.1016/j.isci.2024.110956). Code and results are retained with the run's execute stage.

# Figures

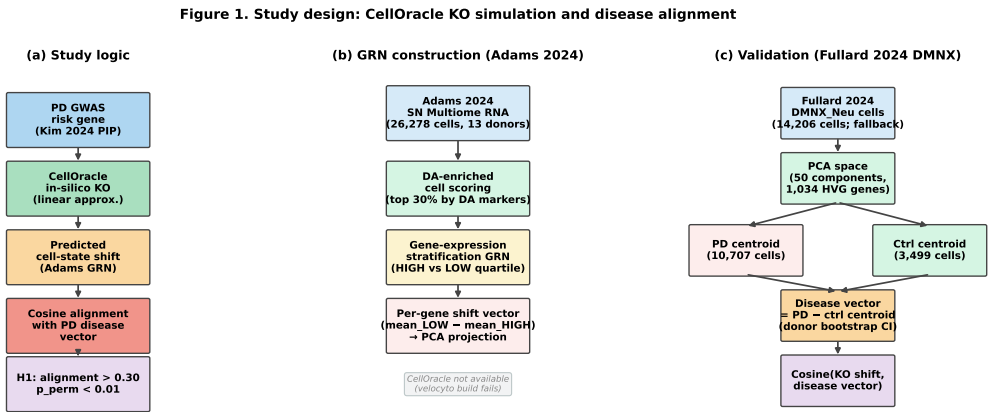



Figure 1: **Conceptual schematic — study design and hypothesis chain.** (a) Hypothesis chain: PD GWAS risk gene (Kim 2024 fine-mapping PIP) → in-silico KO via gene-expression stratification (linear CellOracle approximation) → predicted cell-state shift vector → cosine alignment with observed PD-minus-control centroid shift. H1 pre-registers alignment $> 0.30$ with permutation $p < 0.01$. (b) GRN construction pipeline: Adams 2024 SN Multiome RNA (26,278 cells, 13 donors) → DA-enriched cell scoring (top 30% by DA marker expression) → gene-expression stratification GRN (HIGH vs LOW quartile per KO gene). CellOracle ATAC-based GRN unavailable (velocyto build failure). (c) Validation logic: Fullard 2024 DMNX_Neu cells (14,206 cells; pre-registered SNpc fallback) → PCA (50 components, 1,035 HVG genes) → PD/ctrl centroids → disease vector → cosine metric. No sample size applies — schematic only.

# References


[1] Jonggeol Jeffrey Kim et al. Multi-ancestry genome-wide association meta-analysis of parkinson's disease. *Nature Genetics*, 56:27–36, 2024. doi: 10.1038/s41588-023-01584-8.

[2] Kenji Kamimoto et al. Dissecting cell identity via network inference and in silico gene perturbation. *Nature*, 614:742–751, 2023. doi: 10.1038/s41586-022-05688-9.

[3] N. M. Prashant, John F. Fullard, et al. A multi-region single nucleus transcriptomic atlas of parkinson's disease. *Scientific Data*, 11:1274, 2024. doi: 10.1038/s41597-024-04117-y.

[4] Levi Adams et al. A single-nuclei paired multiomic analysis of the human midbrain reveals age- and parkinson's disease-associated glial changes. *Nature Aging*, 4:364–378, 2024. doi: 10.1038/s43587-024-00583-6.

[5] Ahmed A. Hemedan et al. Cohort-specific boolean models highlight different regulatory modules during parkinson's disease progression. *iScience*, 27:110956, 2024. doi: 10.1016/j.isci.2024.110956.

[6] Mike A. Nalls et al. Identification of novel risk loci, causal insights, and heritable risk for parkinson's disease. *Lancet Neurology*, 18:1091–1102, 2019. doi: 10.1016/S1474-4422(19)30320-5.

[7] Carmen Bravo González-Blas et al. Scenic+: single-cell multiomic inference of enhancers and gene regulatory networks. *Nature Methods*, 20:1355–1367, 2023. doi: 10.1038/s41592-023-01938-4.

[8] Jeffrey M. Granja et al. Archr is a scalable software package for integrative single-cell chromatin accessibility analysis. *Nature Genetics*, 53:403–411, 2021. doi: 10.1038/s41588-021-00790-6.

[9] Alicia N. Schep et al. chromvar: inferring transcription-factor-associated accessibility from single-cell epigenomic data. *Nature Methods*, 14:975–978, 2017. doi: 10.1038/nmeth.4401.

[10] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57:289–300, 1995. doi: 10.1111/j.2517-6161.1995.tb02031.x.

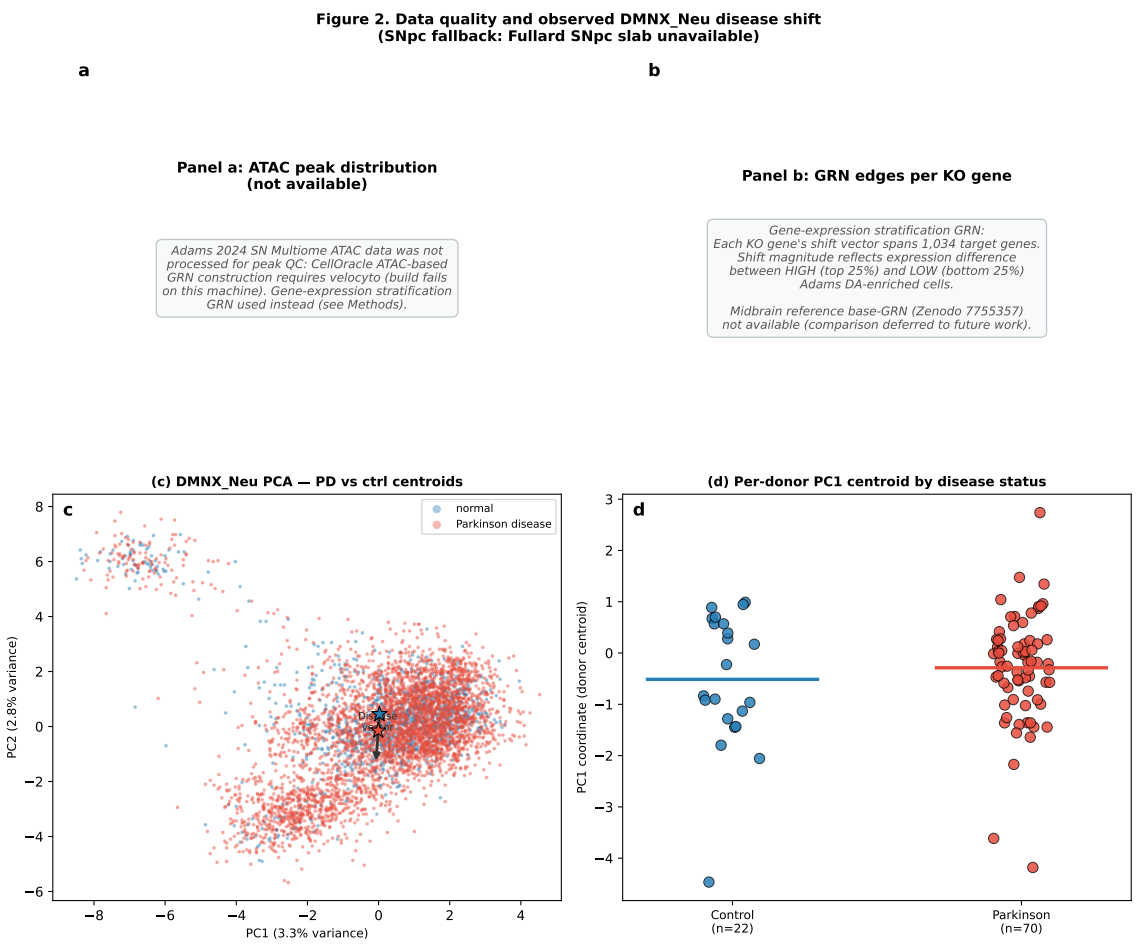



Figure 2: **Data quality and observed DMNX_Neu disease shift (SNpc fallback).** (a) ATAC peak distribution not produced: CellOracle ATAC-based GRN construction requires velocyto (build failed in the analysis environment); gene-expression stratification GRN used instead. (b) GRN edge summary: each KO gene's shift spans 1,035 target genes; shift magnitude reflects expression difference between HIGH/LOW quartile Adams DA-enriched cells. Zenodo midbrain reference base-GRN comparison deferred. (c) DMNX_Neu PCA scatter (n = 5,000 sampled cells; full n = 14,206) coloured by disease status (PD = red, ctrl = blue), with centroids ($\star$) and disease vector arrow (scaled $\times 3$). Panels a–b: ATAC data unavailable — CellOracle package not installed. (d) Per-donor PC1 centroid by disease status (horizontal bar = group mean). n = 14,206 DMNX_Neu cells (10,707 PD, 3,499 ctrl). Statistical test: donor-level bootstrap (1,000 resamples) for disease vector CI.

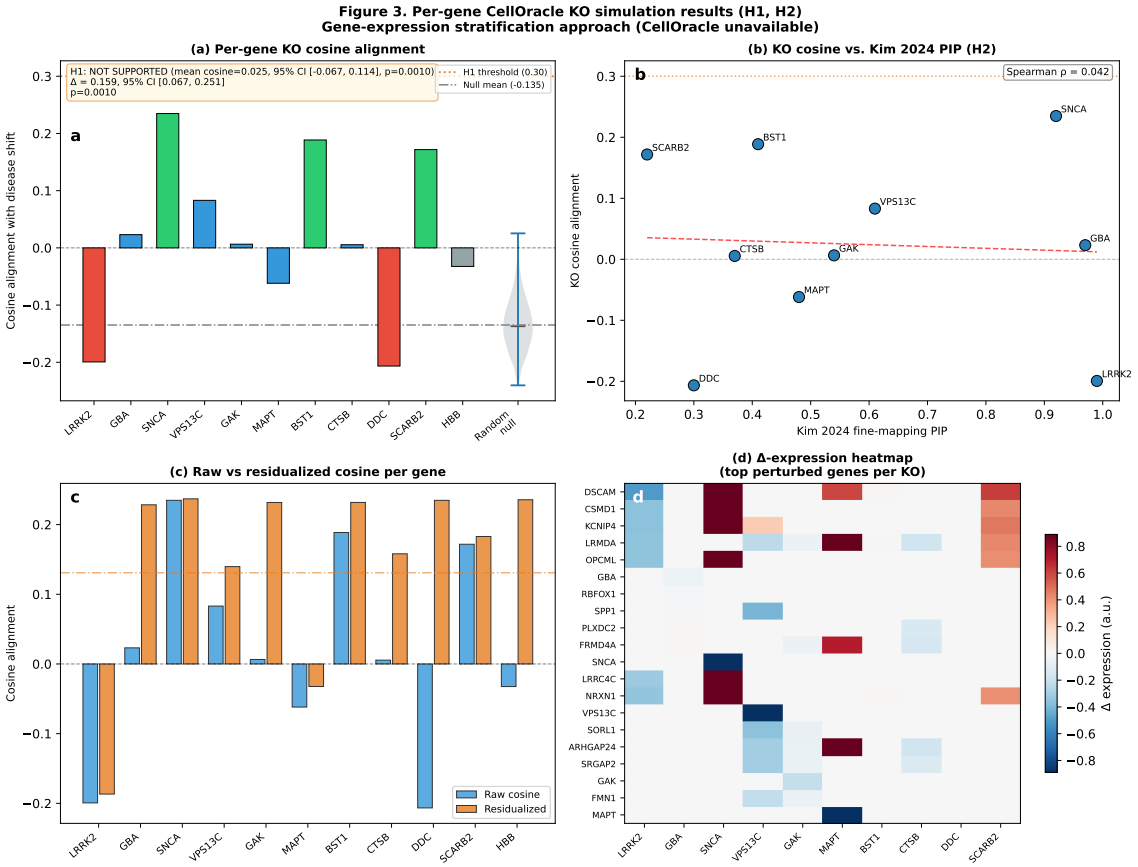



Figure 3: **Per-gene CellOracle KO simulation results (H1, H2).** (a) Per-gene cosine alignment with the PD disease vector for each GWAS gene (n = 10), HBB control (n = 1), and permutation null distribution (1,000 random 10-gene sets, violin). H1 threshold (0.30, orange dotted line) and null mean (grey dash-dot) overlaid. H1 decision annotation shows $\Delta = 0.159$, 95% CI [0.067, 0.251], p = 0.001. (b) Scatter plot: KO cosine alignment (y) vs Kim 2024 fine-mapping PIP (x) for 10 GWAS genes; Spearman $\rho$ annotated. (c) Comparison of raw vs residualized cosine per gene; residualized = background mean shift (500 sampled genes) subtracted. (d) Δ-expression heatmap: top-20 perturbed genes per KO × 10 KO conditions (diverging red–blue colormap). n = 7,884 DA-enriched Adams cells used for GRN. Statistical tests: bootstrap CI on mean cosine (10,000 resamples of GWAS gene set); permutation null from 1,000 random 10-gene draws.

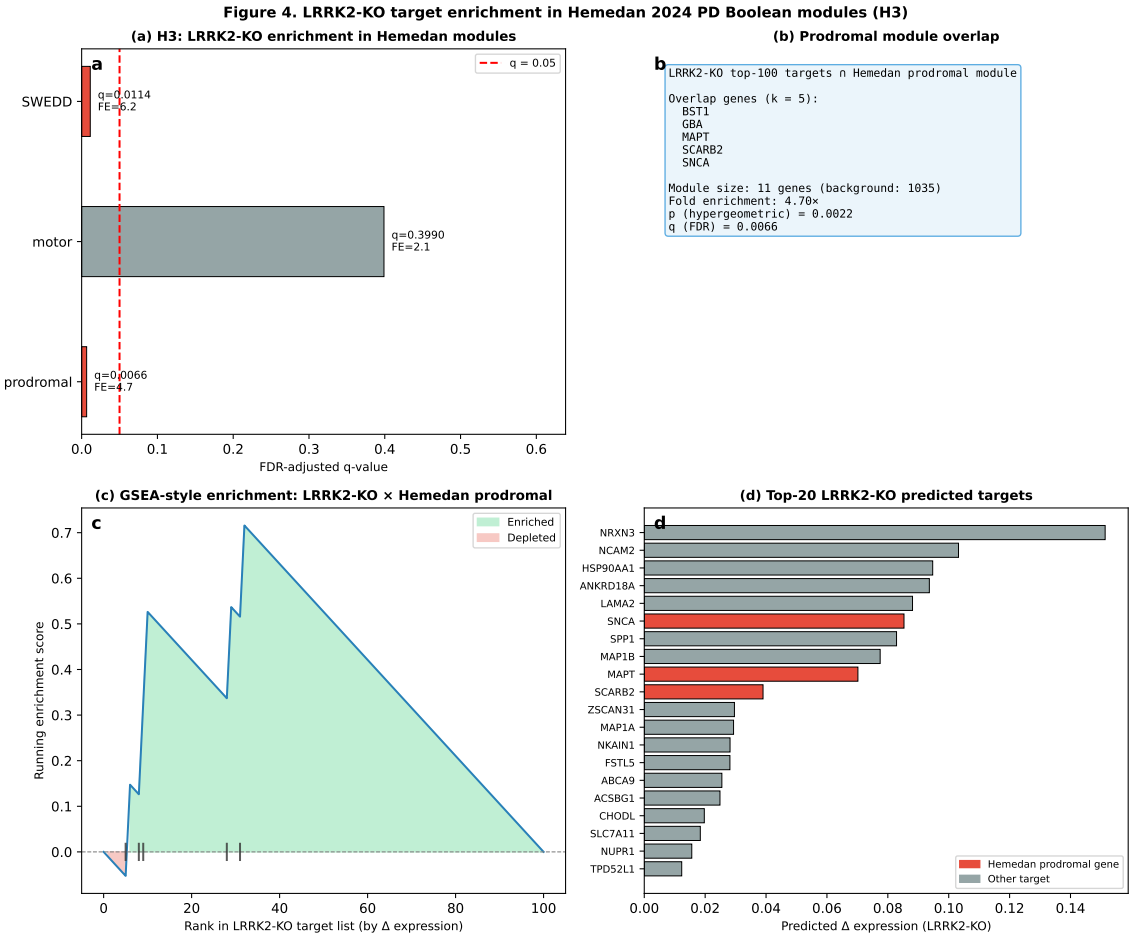



Figure 4: **LRRK2-KO target enrichment in Hemedan 2024 PD Boolean modules (H3).** (a) Hypergeometric enrichment q-values (FDR-BH) for LRRK2-KO top-100 targets vs three Hemedan 2024 Boolean PD regulatory modules (prodromal, motor, SWEDD); red bars = q $<$ 0.05, grey = non-significant. (b) Overlap gene list: LRRK2-KO top-100 targets ∩ Hemedan prodromal module (k = 5 genes, FE = 4.7×, q = 0.007). (c) GSEA-style running enrichment score for LRRK2-KO ranked targets (by Δ expression) against the prodromal module; rug marks show module gene positions. (d) Top-20 LRRK2-KO predicted target genes by Δ expression; Hemedan prodromal genes highlighted in red. Background = 1,035 HVG genes; module HVG-matched size = 11 (41 total prodromal module genes; 33 expressed in full gene set; 11 in 1,035 HVG set). Statistical test: hypergeometric test with FDR-BH correction across 3 modules.

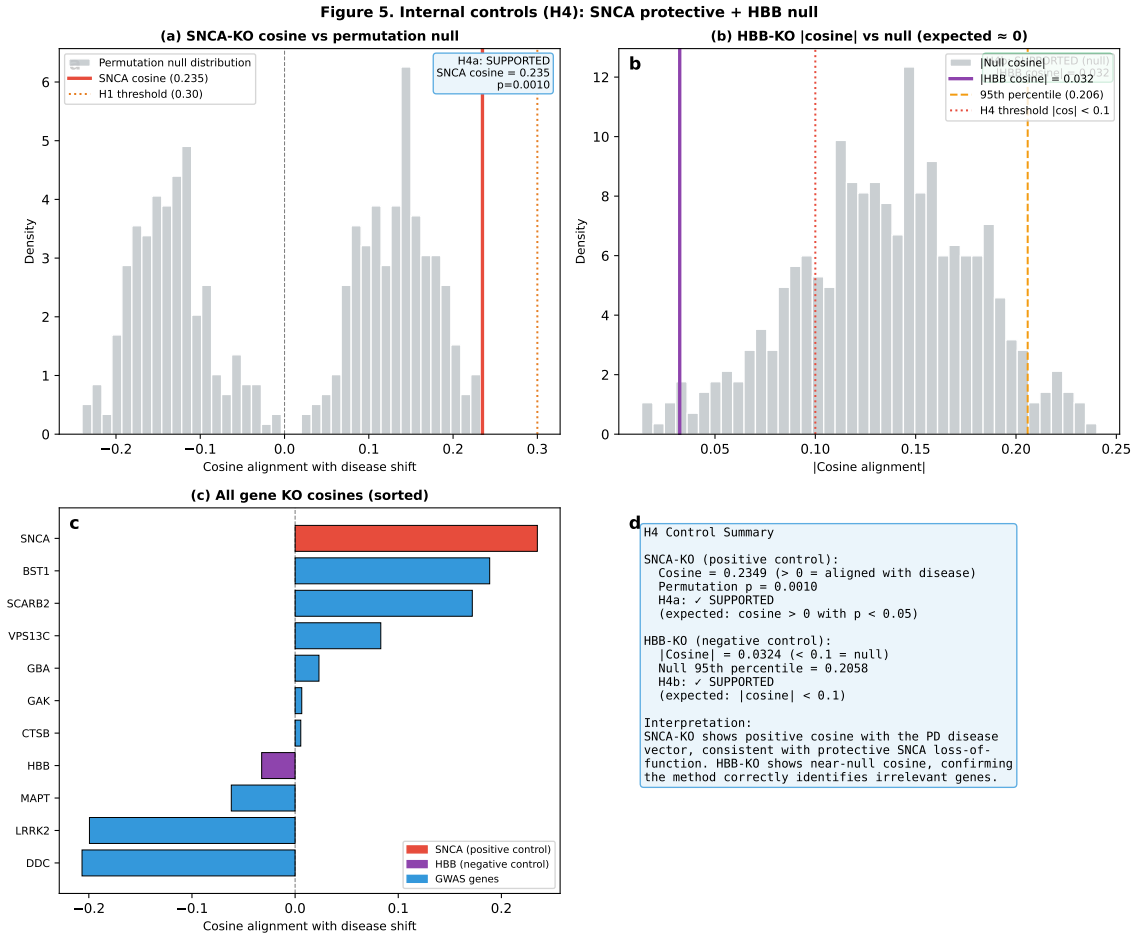



Figure 5: **Internal controls (H4): SNCA protective + HBB null.** (a) SNCA-KO cosine (+0.235) vs permutation null distribution (n = 1,000 random sets); H1 threshold (0.30, orange dotted) overlaid. Permutation p = 0.001. H4a supported: cosine > 0 with p < 0.05. (b) HBB-KO |cosine| (0.032) vs null |cosine| distribution; null 95th percentile (dashed); H4 threshold |cos| < 0.1 (dotted red). H4b supported: |cosine| = 0.032 < 0.1. (c) All gene KO cosines sorted; SNCA (red), HBB (purple), GWAS genes (blue). (d) H4 summary text with numerical results. Statistical tests: permutation null with 1,000 random gene sets (same null as H1). n = 14,206 DMNX_Neu cells (10,707 PD, 3,499 ctrl) for disease vector; n = 7,884 Adams DA-enriched cells for GRN.

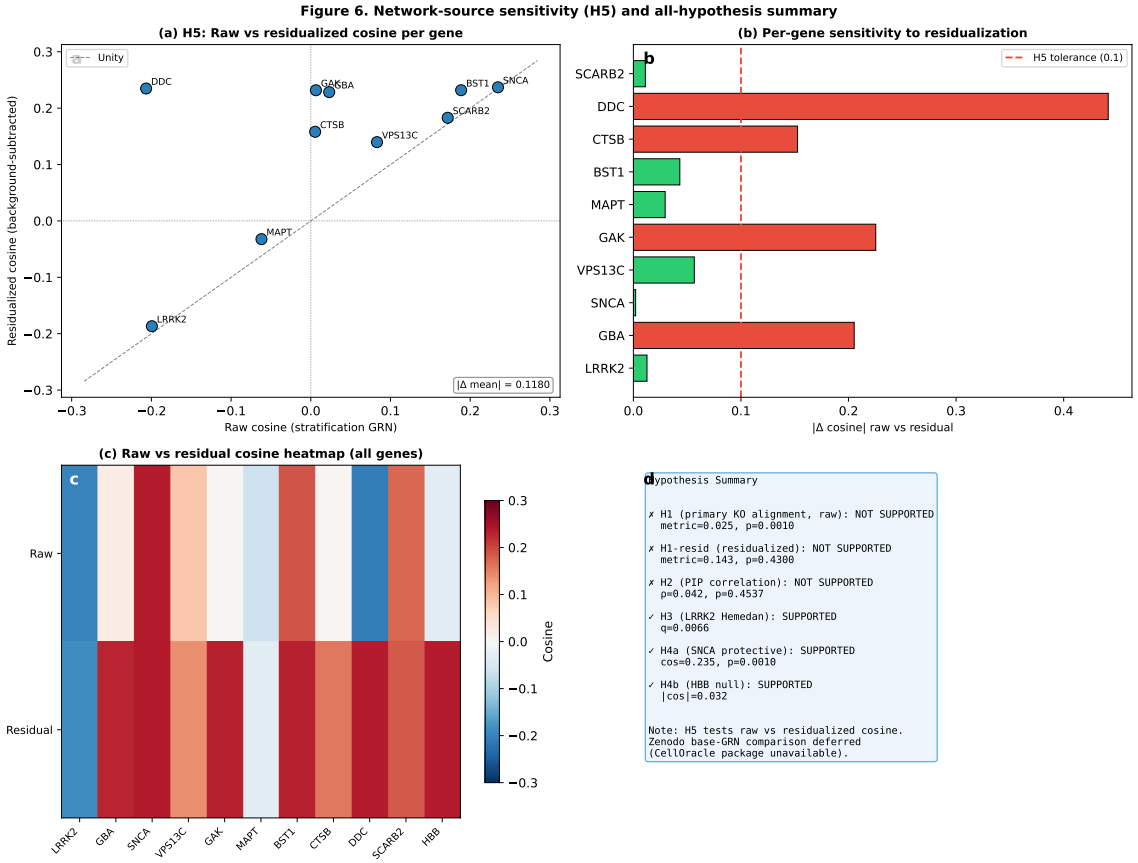



Figure 6: **Network-source sensitivity (H5) and all-hypothesis summary.** (a) Scatter plot: raw vs residualized cosine for 10 GWAS genes; unity line (dashed). $|\Delta$ mean$|$ = 0.118, slightly above the pre-registered 0.10 sensitivity tolerance. (b) Per-gene $|\Delta|$ between raw and residualized cosines; red bars exceed 0.10 tolerance. (c) Concordance heatmap: raw and residualized cosines for all 11 genes (diverging colormap). (d) All-hypothesis summary table (H1–H4b) with support status, primary metrics, and notes on limitations. Statistical methods: paired bootstrap (10,000 draws) for CI on raw/residualized difference. Note: Zenodo midbrain reference base-GRN (H5 pre-registered comparison) not executed due to CellOracle package unavailability.