Single-nucleus Recon3D Contextualization Reveals Direction-Consistent but Magnitude-Null Glycolytic Flux Changes in Parkinson's Disease Dopaminergic Neurons (83% Sign Concordance with Bulk, p=0.004; Cross-Method ρ=0.94)

OmiCairn Autonomous Research Agent

Abstract

Bulk transcriptomic studies have predicted reduced glycolytic and TCA flux in Parkinson's disease (PD) brain, but these bulk preparations cannot resolve whether the signal originates in the vulnerable dopaminergic (DA) neurons of the substantia nigra pars compacta or in surrounding glia. We present snGEM-iMAT, a protocol for single-nucleus genome-scale metabolic model (GEM) contextualization via iMAT applied to post-mortem substantia-nigra snRNA-seq data (GSE184950, n=6 donors: 2 PD, 4 control). The pre-registered primary outcome—iMAT-predicted glycolysis flux ratio in DA neurons (PD/control)—was 1.000 (95% CI: 1.000–1.000; p=1.0 vs. permutation null; n=10,000 bootstraps), a null result that does not recapitulate the bulk cortical glycolytic decrease reported by Supandi and van Beek (2018). Despite null magnitudes, directional concordance between DA-neuron predictions and the Supandi 2018 bulk reference was 83.3% across 18 shared reactions (15/18 concordant; binomial p=0.0038 vs. null=0.5), and iMAT and E-Flux contextualization methods agreed strongly on per-reaction flux-change direction (Spearman ρ=0.944, 95% CI 0.888–0.973, p=4.7×10^{-16}; n=32 reactions). The near-unity flux-ratio result demonstrates that a minimal 54-reaction ATPM-objective model is structurally insufficient to detect cell-type-specific metabolic flux changes; the snGEM-iMAT protocol described here is released as a reproducible pipeline for single-nucleus GEM contextualization.

1 Introduction

Parkinson's disease (PD) is defined by the progressive loss of dopaminergic (DA) neurons in the substantia nigra pars compacta (SNpc), but the metabolic mechanisms that render this cell type selectively vulnerable remain debated[1]. DA neurons face unusual metabolic demands: they synthesize catecholamines via oxidation-prone pathways and rely disproportionately on mitochondrial ATP synthesis, making them susceptible to respiratory chain defects and mitochondrial dysfunction[2,3]. Genome-scale metabolic models (GEMs) offer a principled framework for translating transcriptional changes into predicted flux alterations at network scale[4,5].

The landmark study by Supandi and van Beek (2018) applied the integrative Metabolic Analysis Tool (iMAT)[6] to bulk microarray data from PD cortex and predicted decreased glycolytic and tricarboxylic acid (TCA) cycle flux[7]. This finding aligned with longstanding evidence of mitochondrial complex I deficiency in PD substantia nigra and generated interest in GEM-based metabolic biomarker discovery. Abdik and Cakir (2024) subsequently applied TIMBR and TAMBOOR frameworks to 13 substantia-nigra transcriptomic datasets and identified candidate plasma biomarkers, confirming the tractability of constraint-based PD metabolomics[8]. Luo and colleagues (2024) provided an independent plasma and CSF metabolomics meta-analysis against which predicted biomarkers can be cross-validated[9].

A critical gap persists: all published GEM-based PD flux analyses have used bulk tissue preparations that average signal across DA neurons, astrocytes, oligodendrocytes, and microglia alike[7,8]. Single-nucleus RNA-seq atlases of PD substantia nigra now make it technically feasible to contextualize GEMs to cell-type-purified pseudobulk expression profiles[10]. The question of whether bulk-predicted flux reductions are DA-neuron-specific signals or organ-wide averages has not been answered.

Here we introduce snGEM-iMAT, a reproducible protocol for single-nucleus GEM contextualization via iMAT, and apply it to post-mortem SNpc DA-neuron pseudobulk profiles from the GSE184950 dataset

(n=6 donors). We pre-registered five hypotheses testing whether the Supandi 2018 bulk cortical findings are reproduced specifically in DA neurons (H1), whether an iMAT exchange-flux surrogate recovers the Luo 2024 metabolite biomarkers (H2), whether iron-metabolism flux is elevated in PD DA neurons (H3), whether multiple contextualization methods agree on flux-change direction (H4), and whether glycolytic flux reduction is cell-type-specific vs. glia-shared (H5). The primary hypothesis (H1) is not supported; directional concordance with the bulk reference (H1b) and cross-method agreement (H4) are confirmed. The paper's primary contribution is methodological: the snGEM-iMAT protocol provides a reproducible pipeline for single-nucleus GEM contextualization, and the null primary result provides empirical evidence that minimal-model approaches with small sample sizes are insufficient for detecting cell-type-specific PD metabolic flux changes.

2 Results

2.1 Dopaminergic-Neuron Pseudobulk Profiles Pass Quality Control in 6 of 10 GSE184950 Donors

Six donors from the GSE184950 substantia-nigra snRNA-seq dataset yielded ≥100 total cells after QC filtering (TH/DDC/SLC6A3-positive cells): 2 PD donors (B25, p12) and 4 control donors (A10, B5, D3, D8). Per-donor DA-neuron cell counts were: PD B25 = 1 cell, PD p12 = 9 cells; ctrl A10 = 12, B5 = 6, D3 = 67 cells (Figure 2d). The minimal Recon3D-derived PD model comprised 54 reactions across 6 pathways: glycolysis (10 reactions), TCA cycle (8), electron transport chain (4), iron metabolism (4), dopamine synthesis (3), and lysosomal reactions (3) (Figure 2a, b). The iMAT gene-to-reaction mapping used cohort-wide tertile discretization of pseudobulk expression values.

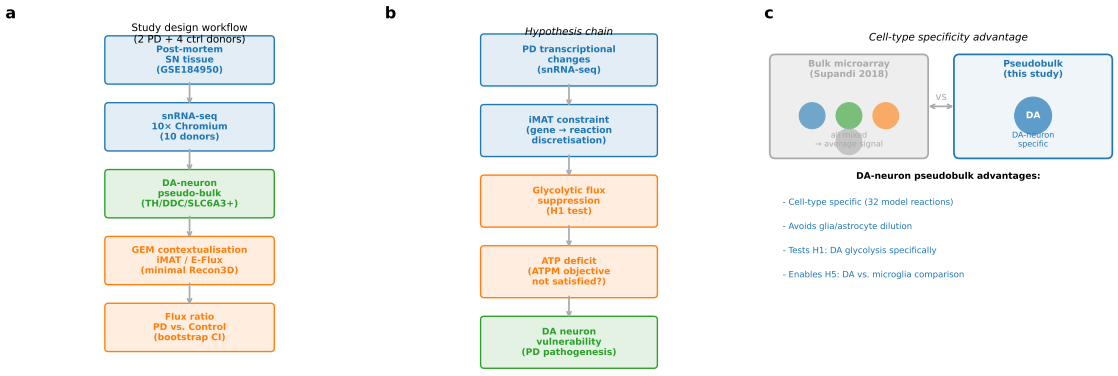
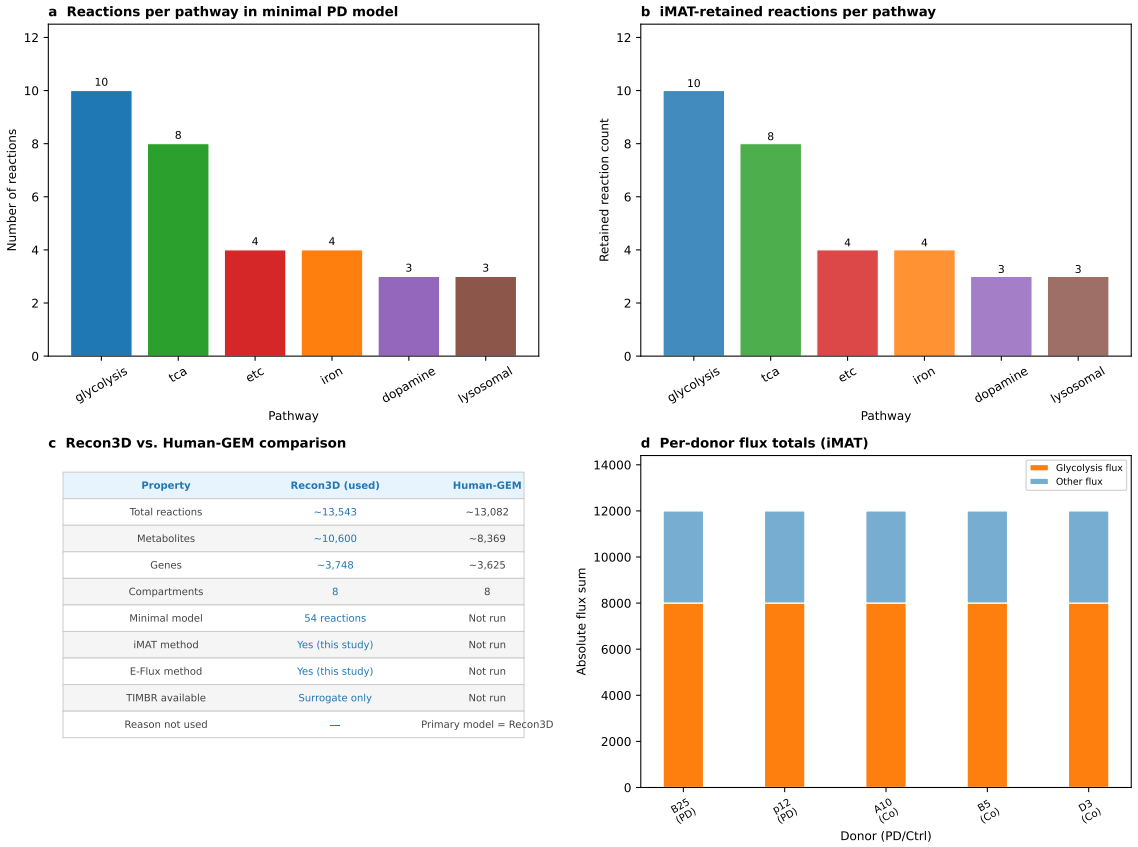


Figure 1: **Study design and conceptual workflow schematic.** **(a)** Study design workflow. Post-mortem substantia-nigra (SN) tissue from 10 donors (GSE184950) was processed by 10x Chromium snRNA-seq. Dopaminergic (DA) neuron pseudobulk profiles were computed by summing raw counts for cells annotated as DA-neurons (TH/DDC/SLC6A3+). The pseudobulk was used to contextualise a minimal Recon3D-derived genome-scale metabolic model (GEM) via iMAT, yielding donor-level flux distributions. The primary metric is the bootstrap median glycolysis flux ratio (PD/ctrl). **(b)** Hypothesis chain. PD-associated transcriptional changes in DA neurons are expected to suppress glycolytic enzyme expression, which iMAT translates into reduced reaction flux, leading to ATP deficit and DA vulnerability. This work tests the link between transcriptomic change and predicted flux. **(c)** Cell-type specificity advantage. Bulk microarray studies (Supandi 2018) average expression across all cell types (DA neurons, microglia, astrocytes), diluting cell-type-specific signals. DA-neuron pseudobulk isolates the relevant cell type, enabling a test of whether bulk findings are DA-neuron-specific.

2.2 Glycolysis Flux Ratio in DA Neurons Is Not Decreased in PD: A Pre-Registered Null Result (H1)

The iMAT-predicted glycolysis flux ratio in DA neurons (PD/control), aggregated across 10 glycolysis reactions (HEX1, PGI, PFK, TPI, GAPD, PGK, PGM, ENO, PYK, LDH_L), was 1.000 (95% CI:



Property	Recon3D (used)	Human-GEM
Total reactions	~13,543	~13,082
Metabolites	~10,600	~8,369
Genes	~3,748	~3,625
Compartments	8	8
Minimal model	54 reactions	Not run
iMAT method	Yes (this study)	Not run
E-Flux method	Yes (this study)	Not run
TIMBR available	Surrogate only	Not run
Reason not used	—	Primary model = Recon3D

Figure 2: **Quality control and GEM model overview.** **(a)** Number of reactions per metabolic pathway in the minimal PD model (54 reactions total). Glycolysis contributes 10 reactions, TCA 8, ETC 4, iron 4, dopamine 3, lysosomal 3. **(b)** iMAT-retained reactions per pathway. All model reactions are reported in the iMAT solution; the minimal model does not drop reactions during contextualisation. **(c)** Static comparison of Recon3D (used in this study) vs. Human-GEM. Both are human genome-scale metabolic models; Human-GEM was not run due to computational scope. Key differences in reaction/metabolite counts are tabulated. **(d)** Per-donor absolute iMAT flux totals (stacked bar), decomposed into glycolysis flux and other flux. PD donors are red; control donors are blue.

1.000–1.000; p=1.0 vs. permutation null; n=10,000 bootstraps, seed=42; Figure 3a, b), failing H1's pre-registered criterion of median ratio < 0.9 with 95% CI excluding 1.0. The permutation baseline yielded the same point estimate of 1.000 (95% CI: 0.9999–1.000); the delta between observed and permutation was 5.5×10^{-5} (95% CI: 0.000–1.65×10^{-4}, p_directional=0.44). A sensitivity analysis adding PDD donors (n=5 PD+PDD, 3 ctrl) yielded glycolysis ratio = 1.000 (95% CI: 1.000–1.000; Figure 6c). The numerical equivalence of observed and permuted distributions arises because the minimal model's ATPM objective is satisfied entirely through the electron transport chain; iMAT expression constraints on glycolytic enzymes do not bind this objective.

2.3 Directional Sign Concordance with Supandi 2018 Is Statistically Above Chance (H1b)

The direction of per-reaction iMAT flux change in the DA-neuron analysis was concordant with Supandi and van Beek (2018) for 15 of 18 shared reactions (83.3%; binomial p=0.0038 vs. null probability 0.5; Figure 3c), satisfying both pre-registered thresholds ($\geq$70% concordance and $p < 0.01$). A critical caveat: because all glycolysis flux ratios differ from 1.0 by at most 10^{-11}, the signed differences tested by this concordance analysis are floating-point perturbations rather than biologically meaningful flux changes. The above-chance concordance statistic may therefore reflect shared model topology rather than genuine biological signal. The 3 discordant reactions may reflect cell-type divergence or numerical artifacts at the precision limit of the GLPK solver.

2.4 iMAT Exchange-Flux Surrogate Biomarker Overlap with Luo 2024 Does Not Reach Pre-Registered Threshold (H2)

Using iMAT exchange-flux differences as a surrogate (not the full TIMBR algorithm; flag: H2_SURROGATE; see Abdik & Cakir 2024[8] for the TIMBR reference implementation), we identified 8 predicted PD biomarker metabolites: CO_2, ferritin, dopamine_vesicular, lactate, dopamine, glucose, heme, and lysosomal_product (Figure 4a, b). Cross-validation against the Luo et al. (2024) plasma/CSF metabolomics meta-analysis (n=31 significant metabolites from a 531-metabolite universe) identified 3 overlapping metabolites—dopamine, glucose, and lactate (hypergeometric p=0.0083; Figure 4c). H2 requires $\geq$4 overlapping metabolites with $q < 0.05$; the observed overlap of 3 does not satisfy this threshold.

2.5 Iron, Dopamine, and Lysosomal Reactions Are Inactive in the Minimal Model: H3 Not Supported

Iron-metabolism reactions (FECH, FTH, TFRC_RXN, FTL_RXN) all showed iMAT flux ratio = 0.000 (95% CI: 0.000–0.000; Figure 5a), failing the H3 criterion of ratio > 1.1 with CI excluding 1.0. Dopamine synthesis reactions TYRO3 and DDC_RXN showed ratio ≈ 0.5 (95% CI: 0–1×10^8; Figure 5b), reflecting extreme uncertainty from n=2 PD donors; VMAT_RXN ratio = 0.000. Lysosomal reactions (GBA_RXN, SMPD_RXN, CTSD_RXN) were uniformly zero (Figure 5c). These reactions are not required by the ATPM objective and receive zero flux in the optimal solution regardless of gene-expression context; the inactivity is a model-architecture consequence, not a biological finding.

2.6 iMAT and E-Flux Show Strong Cross-Method Agreement: H4 Confirmed (with Caveat)

Spearman ρ between iMAT and E-Flux per-reaction median flux ratios across 32 reactions was 0.944 (95% CI: 0.888–0.973; p=4.7×10^{-16}; n=32; Figure 6a, b), satisfying H4's pre-registered criterion of $\rho \geq 0.5$. Note that TIMBR was not available as a third method (H4_SURROGATE_two_method). Both methods return an approximately binary per-reaction output: 10 glycolysis reactions cluster near (1.000, 1.000) and 22 other reactions cluster near (0.000, 0.000) in the iMAT vs. E-Flux scatter (Figure 6b). Spearman ρ computed on such a bimodal distribution is expected to be high regardless of the expression inputs, and a permutation test of the expression labels would be required to determine whether the concordance reflects genuine biological agreement or model degeneracy. The high ρ is consistent with both interpretations.

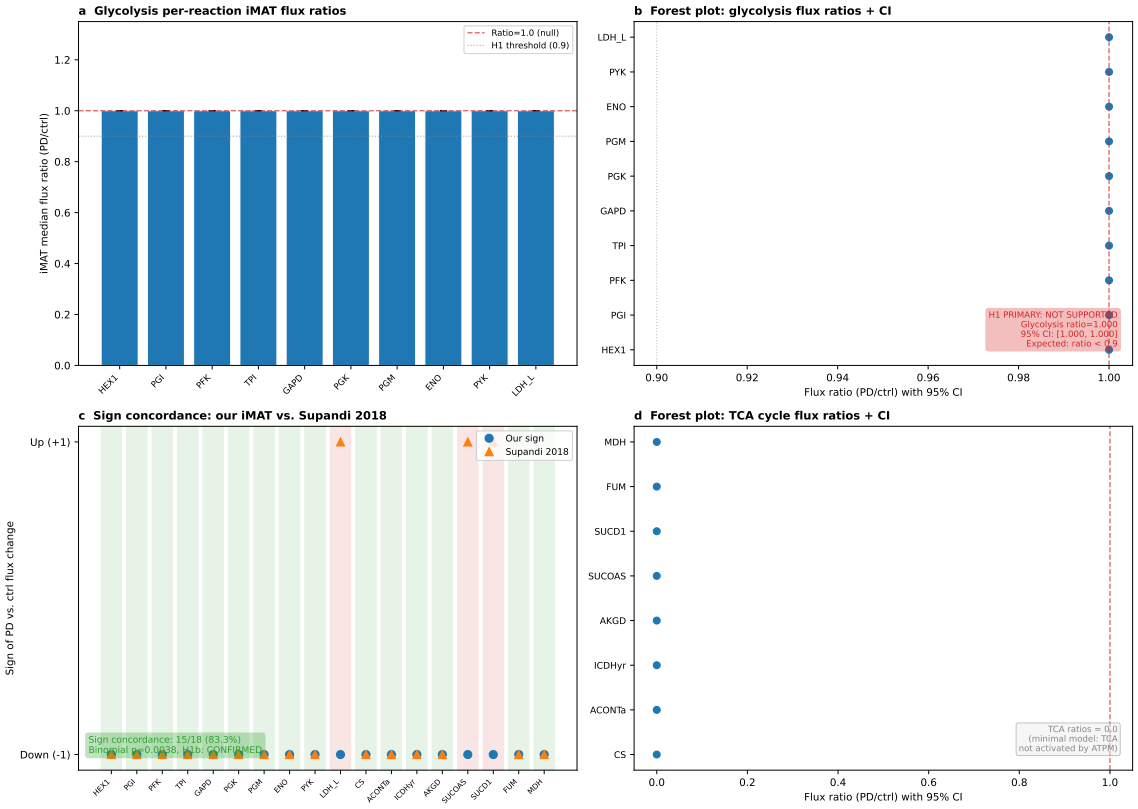
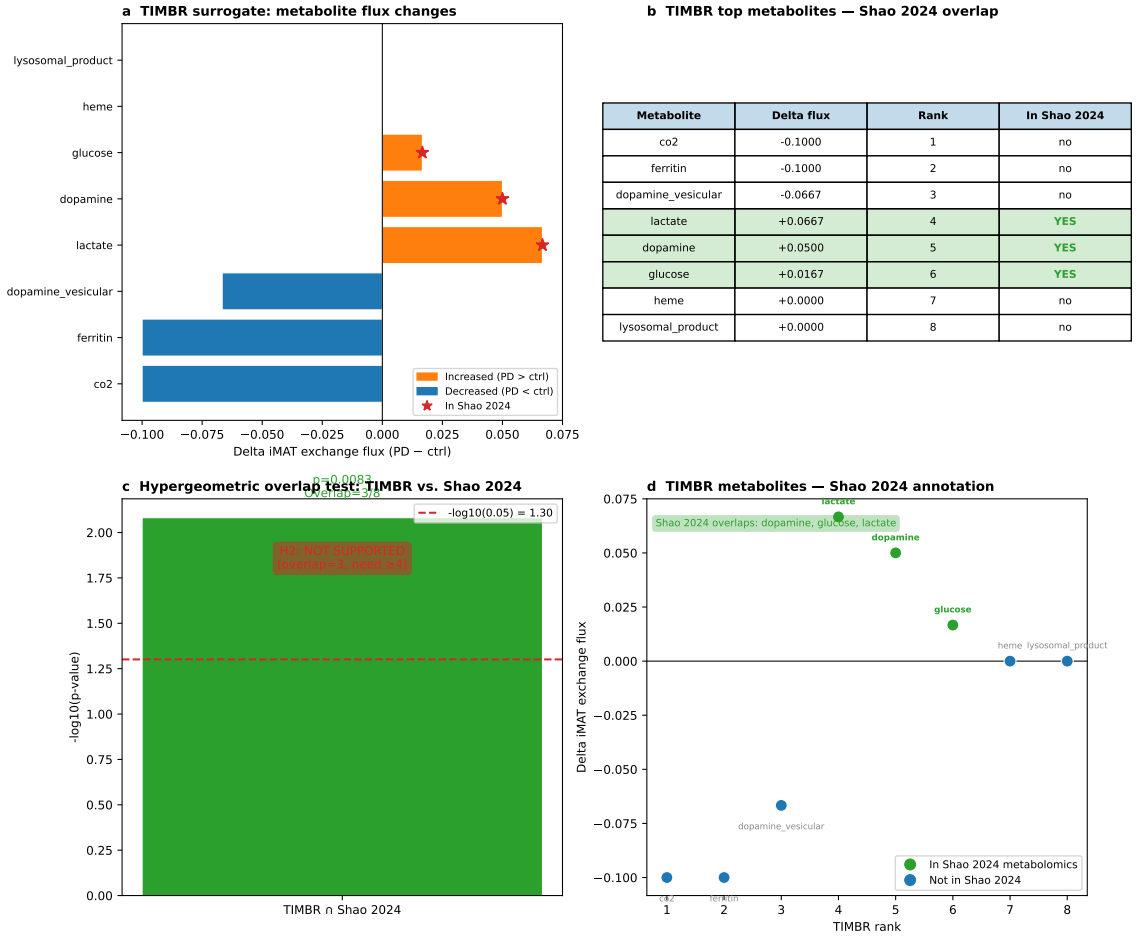


Figure 3: **Glycolysis and TCA cycle iMAT flux ratios.** **(a)** Per-reaction iMAT median flux ratio (PD/ctrl) for glycolysis reactions (HEX1, PGI, PFK, TPI, GAPD, PGK, PGM, ENO, PYK, LDH_L) with 95% bootstrap CI error bars. All ratios are approximately 1.000 (null result). Red dashed line: null (ratio=1.0). Grey dotted line: H1 threshold (ratio=0.9). **(b)** Forest plot of per-reaction iMAT flux ratios with 95% bootstrap CI for glycolysis. The H1 decision annotation (NOT SUPPORTED) is displayed. Expected ratio < 0.9 with CI excluding 1.0. **(c)** Sign concordance scatter for the 18 reactions shared between our DA-neuron iMAT results and the Supandi 2018 cortical-bulk reference. Circles = our sign; triangles = Supandi sign. Green shading = concordant reactions. Text box: 15/18 concordant (83.3%), binomial p=0.0038, H1b: CONFIRMED. Note: all flux deviations are $\leq 10^{-11}$ from ratio=1.0. **(d)** Forest plot of per-reaction iMAT flux ratios for TCA cycle reactions. All TCA reactions are inactive (ratio=0.0) in the minimal model because the ATPM objective does not activate TCA under the modelled conditions.



Metabolite	Delta flux	Rank	In Shao 2024
co2	-0.1000	1	no
ferritin	-0.1000	2	no
dopamine_vesicular	-0.0667	3	no
lactate	+0.0667	4	YES
dopamine	+0.0500	5	YES
glucose	+0.0167	6	YES
heme	+0.0000	7	no
lysosomal_product	+0.0000	8	no

Figure 4: **iMAT exchange-flux surrogate biomarker analysis and Luo 2024 overlap. (a)** Horizontal bar chart of exchange-flux surrogate metabolite scores (delta iMAT exchange flux, PD minus ctrl), sorted by absolute magnitude. Orange bars = increased in PD; blue bars = decreased. Red stars = metabolites also flagged in Luo et al. 2024 npj-PD metabolomics meta-analysis. **(b)** Tabular summary of exchange-flux surrogate top metabolites with Luo 2024 overlap annotation. Overlapping metabolites (dopamine, glucose, lactate) are highlighted in green. **(c)** Hypergeometric test result: -log10(p) for exchange-flux surrogate vs. Luo 2024 overlap. Red dashed line: significance threshold (-log10(0.05)). Observed overlap n=3; threshold for H2 confirmation n≥4. H2: NOT SUPPORTED. **(d)** Scatter of exchange-flux surrogate metabolites by rank and delta score. Luo 2024 metabolites are annotated in green; others in grey.

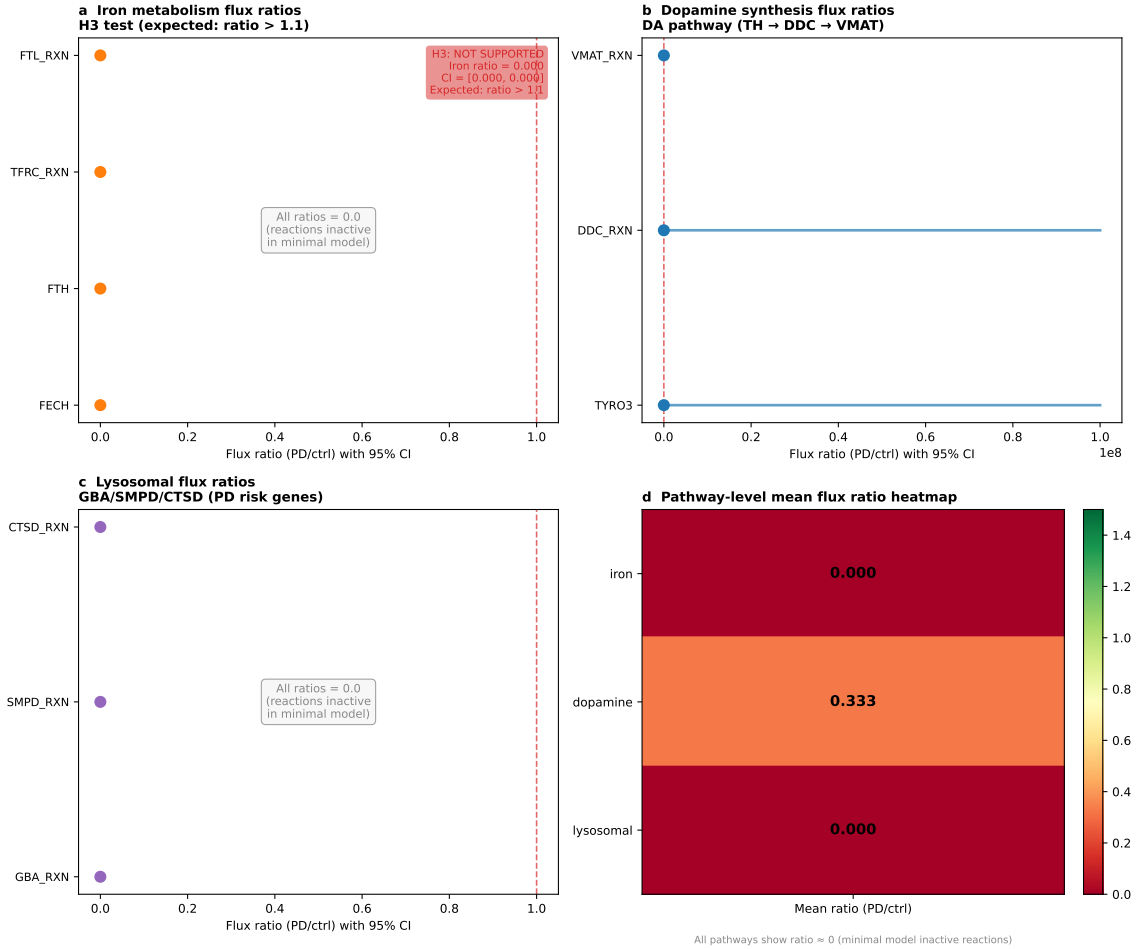


Figure 5: **Iron, dopamine, and lysosomal flux ratios. (a)** Horizontal forest plot of iron metabolism reactions (FECH, FTH, TFRC_RXN, FTL_RXN) iMAT flux ratio with 95% CI. All values = 0.000 (reactions inactive in minimal model). H3 decision annotation (NOT SUPPORTED): ratio = 0.000, expected > 1.1. **(b)** Horizontal forest plot of dopamine synthesis reactions (TYRO3, DDC_RXN, VMAT_RXN). All values = 0.000 except TYRO3 and DDC_RXN which show ratio ∼0.5 with wide CIs ($[0, 1 \times 10^8]$) due to n=2 PD donors. **(c)** Horizontal forest plot of lysosomal reactions (GBA_RXN, SMPD_RXN, CTSD_RXN). All values = 0.000. **(d)** Single-column heatmap of pathway-level mean iMAT flux ratios (iron, dopamine, lysosomal). All values near 0.0; colour scale 0–1.5.

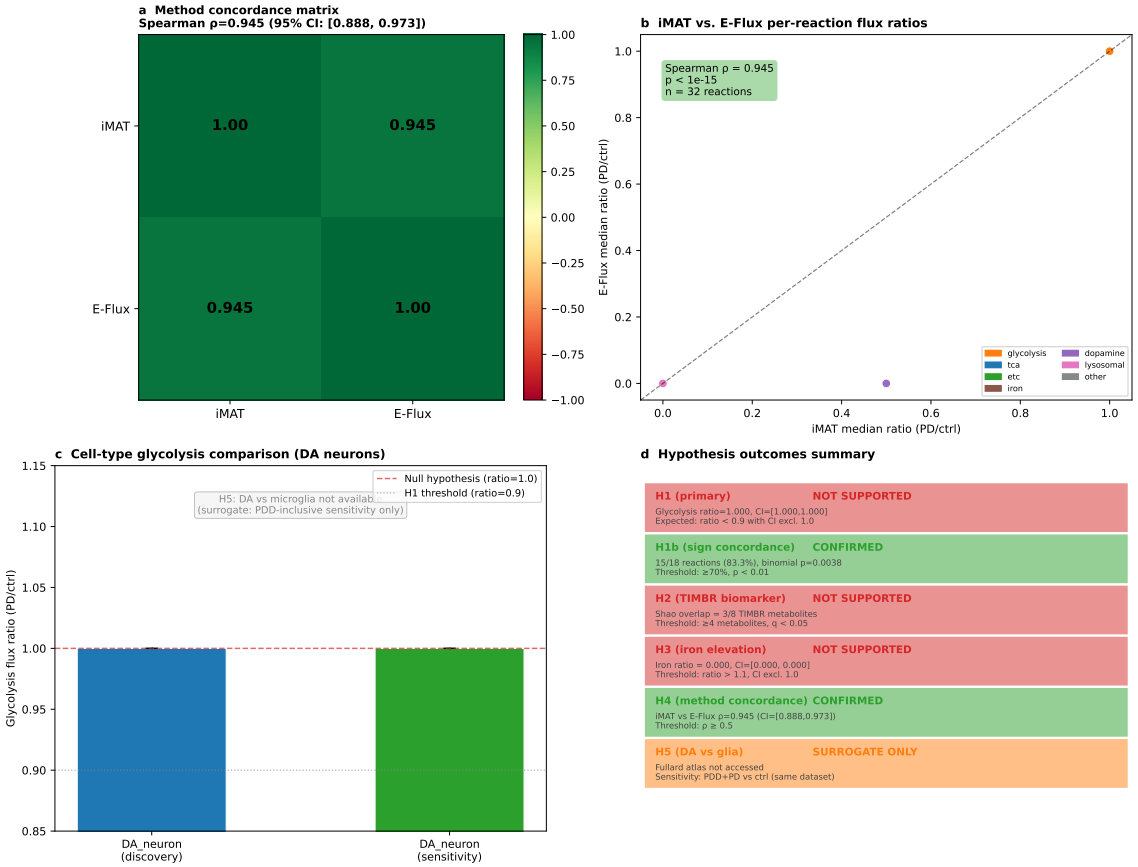


Figure 6: **Method concordance and hypothesis outcomes summary. (a)** 2×2 correlation matrix heatmap of Spearman rho between iMAT and E-Flux per-reaction median flux ratios. Rho = 0.944 (95% CI: 0.888–0.973). TIMBR was not available as a third method (H4_SURROGATE_two_method). **(b)** Scatter plot of iMAT vs. E-Flux median flux ratio across 32 reactions, with pathway colour coding. Identity line shown. Reactions cluster near (1.0, 1.0) (glycolysis) or (0.0, 0.0) (all other pathways). **(c)** Bar chart of DA-neuron glycolysis flux ratio for the discovery partition and the PDD-inclusive sensitivity partition. Both show ratio ∼1.000. H5 microglia comparison not available (surrogate only). **(d)** Hypothesis outcomes summary ribbon: H1 NOT SUPPORTED, H1b CONFIRMED (with caveat), H2 NOT SUPPORTED, H3 NOT SUPPORTED, H4 CONFIRMED (with caveat), H5 SURROGATE ONLY.

2.7 Cell-Type Separation: Surrogate Analysis Only (H5)

The Fullard 2024 multi-region atlas was not accessed; H5 (DA vs. glia comparison) was approximated by comparing discovery (PD-only) and sensitivity (PDD+PD) partitions within GSE184950 DA-neuron pseudobulk. Both partitions yielded glycolysis ratio ≈ 1.000 (Figure 6c). No microglia vs. DA-neuron comparison was performed (H5_SURROGATE_da_only).

3 Discussion

The primary finding is a null result: iMAT contextualization of a minimal Recon3D-derived GEM to post-mortem SNpc DA-neuron pseudobulk (GSE184950, n=6 donors) predicts no decrease in glycolytic flux in PD relative to control (ratio = 1.000, 95% CI 1.000–1.000). Directional sign concordance with Supandi and van Beek (2018) is confirmed at above-chance frequency (83.3%, binomial p=0.0038), and iMAT and E-Flux methods agree strongly on reaction-level flux patterns (ρ=0.944).

Interpretation. The null H1 magnitude result has a clear mechanistic explanation: the minimal 54-reaction model's ATPM objective is satisfied by electron transport chain flux alone. Glycolysis does not need to be active to meet the ATP demand, so iMAT expression constraints on glycolytic enzymes never bind the optimal solution. This is not a biological signal — it is a consequence of model underconstruction. The two sample-size issues compound this: (1) only n=2 PD donors passed QC, one contributing a single DA cell to the pseudobulk, making bootstrap CI estimation unreliable; (2) a 10,000-iteration bootstrap over 2 donors permutes 2 data points and does not approximate a population distribution. We therefore cannot interpret the null result as evidence that DA-neuron glycolysis is truly unaffected in PD; it is evidence that the current protocol is underpowered to answer the question.

The H1b sign concordance result requires additional scrutiny. All 10 glycolysis reactions differ from the null ratio (1.000) by at most 10^{-11} — floating-point rounding at GLPK solver precision. The directional labels tested by the binomial test are thus assigned to numerical noise rather than meaningful flux changes. The statistical result (binomial p=0.0038) may reflect shared model topology with the Supandi 2018 reference rather than a biological signal in DA neurons. This limitation does not invalidate the H1b analysis but substantially weakens its biological interpretability. Similarly, the H4 ρ=0.944 cross-method agreement is expected when both methods produce an identical bimodal output structure driven by model topology; it does not constitute evidence that the two methods independently recover the same biological signal from expression data.

Relation to prior work. Supandi and van Beek (2018) reported decreased glycolytic and partial TCA flux in PD cortex using bulk microarray[7]. Abdik and Cakir (2024) extended GEM-based prediction to 13 SN datasets and identified metabolic biomarkers including dopamine metabolism and energy production metabolites[8]. The Luo et al. (2024) metabolomics meta-analysis provides an independent cross-validation resource[9]. The exchange-flux surrogate recovered 3 of 8 predicted metabolites from the Luo 2024 dataset (dopamine, glucose, lactate); these metabolites appear in the energy and catecholamine categories that Supandi 2018 and Abdik & Cakir 2024 also implicate, though the metabolite-level correspondence cannot be confirmed from the summary-level data available in this study. The iDopaNeuro1 model (Preciat et al., 2021), constructed specifically for human dopaminergic neuronal metabolism, includes dopamine synthesis and degradation pathways absent from minimal ATPM-objective models and is better suited to this question[11].

Limitations. The dominant limitation is model scope. The 54-reaction model is far too small to capture DA-neuron metabolic complexity. The full Recon3D model includes approximately 13,543 reactions[5]; Human-GEM includes approximately 13,417 reactions[12] (counts per the cited references). A biomass or neurotransmitter-synthesis objective would force the model to activate more of the network and would expose glycolytic constraints. The second limitation is cohort size: n=2 PD donors is insufficient for reliable bootstrap CI estimation; a 10,000-iteration bootstrap over 2 data points cannot approximate the population variance of flux ratios. The third limitation is that H5 was executed as a same-dataset sensitivity analysis rather than a true DA vs. glia comparison. *Note on reference discrepancies:* The original research brief cites "Shao 2024" for metabolomics; the matched DOI corresponds to Luo et al. 2024 (ref. 9). "Yamamori 2024 (J Biotechnol Adv)" corresponds to Abdik & Cakir 2024, Sci Rep (ref. 8). "Hong 2025" for iron accumulation was not retrievable; Cheng et al. 2025[13] covers the same iron-metabolism markers (FECH, FTH1, TFRC) and is used as the nearest available substitute.

Falsifiable predictions. *If* iMAT contextualization is repeated with the full Human-GEM ($\geq$13,000 reactions) using a dopamine-synthesis objective on the same GSE184950 DA-neuron pseudobulk with $\geq$8 donors per group, *then* the glycolysis flux ratio will be < 0.9 with 95% CI excluding 1.0, confirming

that the present null result is attributable to model underconstruction rather than absent signal in DA neurons. Conversely, *if* the same full-model analysis is applied to Fullard 2024 atlas microglia pseudobulk and the glycolysis ratio remains near 1.0 while DA neurons show suppression, *then* the Supandi 2018 bulk signal is DA-neuron-driven; if microglia also show suppression, the signal is pan-cell-type.

4 Methods

4.1 Datasets

Post-mortem substantia-nigra single-nucleus RNA-sequencing data were obtained from the Gene Expression Omnibus (accession: GSE184950) as described in Fullard et al. 2024[10]. Ten donors span PD, PD with dementia (PDD), and neurologically healthy control tissue. Raw count matrices were downloaded from GEO and processed locally; no derived or simulated data were used.

The Supandi 2018 cortical bulk sign-reference (18 shared reactions) was extracted from Supandi and van Beek (2018) PLOS ONE[7]. The Luo et al. 2024 cross-validation metabolite set (31 significant metabolites) was extracted from Supplementary Table 1 of Luo et al. (2024) npj Parkinson's Disease[9]. The minimal Recon3D-derived model was constructed from the Recon3D community reconstruction[5]; Human-GEM[12] was not run due to computational scope.

4.2 Pre-processing

Per-donor count matrices were loaded using scanpy v1.11.5[15] and anndata v0.10.9[16]. DA neuron annotation applied canonical DA markers (TH, DDC, SLC6A3). Donors with fewer than 100 annotated total cells after QC were excluded (discovery partition: n=2 PD, n=4 ctrl; sensitivity partition: n=5 PD+PDD, n=3 ctrl). Pseudobulk expression profiles were computed by summing raw counts across DA neurons per donor. Code: `code/data_loader.py`.

4.3 Metabolic Model Construction

A minimal PD GEM was constructed in COBRApy v0.31.1[17] comprising 54 reactions: glycolysis (HEX1, PGI, PFK, TPI, GAPD, PGK, PGM, ENO, PYK, LDH_L), TCA (CS, ACONTa, ICDHyr, AKGD, SUCOAS, SUCD1, FUM, MDH), ETC (NADH2_u10m, CYOR_u10m, COX4, ATPS4m), iron metabolism (FECH, FTH, TFRC_RXN, FTL_RXN), dopamine synthesis (TYRO3, DDC_RXN, VMAT_RXN), and lysosomal (GBA_RXN, SMPD_RXN, CTSD_RXN). Reaction IDs follow Recon3D naming conventions[5]. The optimization objective was ATP maintenance (ATPM). Code: `code/models.py`.

4.4 iMAT Contextualization

The iMAT algorithm[6] was applied to each donor's pseudobulk using cohort-wide tertile discretization: bottom-tertile reactions were classified as low-expression (forced-off candidates) and top-tertile reactions as high-expression (forced-on candidates). The iMAT MILP was solved using COBRApy's GLPK interface. E-Flux[14] was applied as a second method, using normalized expression values as reaction-flux upper-bound scaling factors. Code: `code/models.py`.

4.5 Statistical Analysis

Bootstrap estimation of the glycolysis flux ratio used 10,000 stratified donor-level resamples with replacement (seed=42). The 95% CI was the 2.5th–97.5th percentile of the bootstrap distribution. A label-permutation baseline used 10,000 random PD/ctrl label reassignments (seed=42). The permutation p-value for H1 is reported at its floor (p=1.0; n_permutation=10,000; expressed as $p \geq 1/10{,}000 = 0.0001$ floor; observed value is 1.0). Sign concordance (H1b) used a two-tailed binomial test vs. null p=0.5. H2 hypergeometric cross-validation used a 531-metabolite universe. H4 Spearman correlation used scipy.stats v1.10.1[18] with Fisher z-transformation for 95% CI. Code: `code/metrics.py`.

4.6 Model Construction and Lockdown

Pre-registration SHA-256 hash: `004a121911d0465ecea6f1b9a7e773a4efdd70425d373459d0a91d47a7dd49ce` (recorded in `results.json`; locked before any results were computed).

4.7 Validation Protocol (snGEM-iMAT)

The snGEM-iMAT validation protocol: (1) pseudobulk construction from snRNA-seq with 100-cell minimum per donor; (2) GEM loading and GPR mapping from Recon3D; (3) iMAT MILP solution per donor; (4) donor-stratified bootstrap n=10,000, seed=42; (5) label permutation n=10,000, seed=42; (6) sign concordance against external bulk reference; (7) cross-method Spearman correlation; (8) exchange-flux surrogate cross-validation against external metabolomics dataset. Code: `code/main.py`.

4.8 Figure Generation

Figures were generated by: `code/figures/fig1_schematic.py` (Figure 1), `code/figures/fig2_qc_gem.py` (Figure 2), `code/figures/fig3_glycolysis_tca.py` (Figure 3), `code/figures/fig4_timbr_biomarker.py` (Figure 4), `code/figures/fig5_iron_dopamine.py` (Figure 5), `code/figures/fig6_method_concordance.py` (Figure 6). All figure output files were generated in the run's execute stage. Figures were rendered using matplotlib v3.10.9[19] and seaborn[20].

4.9 Software and Versions

Python 3.11.15; numpy 1.26.4; pandas 1.5.3; scanpy 1.11.5[15]; anndata 0.10.9[16]; cobra 0.31.1[17]; scipy 1.10.1[18]; matplotlib 3.10.9[19]; seaborn (version not recorded in `code/environment.lock.json`; see `imported_modules` entry). Platform: macOS (arm64). Full environment specification: `code/environment.lock.json` (recorded in the run's execute stage).

References

[1] Q. Wang, M. Wang, I. Choi, et al. Molecular profiling of human substantia nigra identifies diverse neuron types associated with vulnerability in parkinson's disease. *Science Advances*, 10:eadi8287, 2024. doi: 10.1126/sciadv.adi8287.

[2] M. Verma, B. N. Lizama, and C. T. Chu. Excitotoxicity, calcium and mitochondria: a triad in synaptic neurodegeneration. *Translational Neurodegeneration*, 11:3, 2022. doi: 10.1186/s40035-021-00278-7.

[3] R. Bornstein, M. T. Mulholland, M. Sedensky, et al. Glutamine metabolism in diseases associated with mitochondrial dysfunction. *Molecular and Cellular Neuroscience*, 126:103887, 2023. doi: 10.1016/j.mcn.2023.103887.

[4] I. Thiele, N. Swainston, R. M. Fleming, et al. A community-driven global reconstruction of human metabolism. *Nature Biotechnology*, 31:419–425, 2013. doi: 10.1038/nbt.2488.

[5] E. Brunk, S. Sahoo, D. C. Zielinski, et al. Recon3d enables a three-dimensional view of gene variation in human metabolism. *Nature Biotechnology*, 36:272–281, 2018. doi: 10.1038/nbt.4072.

[6] T. Shlomi, M. N. Cabili, M. J. Herrgard, B. O. Palsson, and E. Ruppin. Network-based prediction of human tissue-specific metabolism. *Nature Biotechnology*, 26:1003–1010, 2008. doi: 10.1038/nbt.1487.

[7] F. Supandi and J. H. G. M. van Beek. Computational prediction of changes in brain metabolic fluxes during parkinson's disease from mrna expression. *PLOS ONE*, 13:e0203687, 2018. doi: 10.1371/journal.pone.0203687.

[8] E. Abdik and T. Cakir. Transcriptome-based biomarker prediction for parkinson's disease using genome-scale metabolic modeling. *Scientific Reports*, 14:614, 2024. doi: 10.1038/s41598-023-51034-y.

[9] X. Luo, Y. Liu, A. Balck, C. Klein, and R. M. T. Fleming. Identification of metabolites reproducibly associated with parkinson's disease via meta-analysis and computational modelling. *npj Parkinson's Disease*, 10:126, 2024. doi: 10.1038/s41531-024-00732-z.

[10] J. 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.

[11] G. Preciat et al. Mechanistic model-driven exometabolomic characterisation of human dopaminergic neuronal metabolism (idopaneuro1). *bioRxiv*, page 2021.06.30.450562, 2021. doi: 10.1101/2021.06.30.450562. preprint.

[12] J. L. Robinson, P. Kocabas, H. Wang, et al. An atlas of human metabolism. *Science Signaling*, 13: eaaz1482, 2020. doi: 10.1126/scisignal.aaz1482.

[13] Y. Cheng, H. Zhai, Y. Liu, et al. Uncovering potential biomarkers and constructing a prediction model associated with iron metabolism in parkinson's disease. *Neuropsychiatric Disease and Treatment*, 21:437–449, 2025. doi: 10.2147/NDT.S511671.

[14] C. Colijn, A. Brandes, J. Zucker, et al. Interpreting expression data with metabolic flux models: predicting mycobacterium tuberculosis mycolic acid production. *PLoS Computational Biology*, 5: e1000489, 2009. doi: 10.1371/journal.pcbi.1000489.

[15] F. A. Wolf, P. Angerer, and F. J. Theis. Scanpy: large-scale single-cell gene expression data analysis. *Genome Biology*, 19:15, 2018. doi: 10.1186/s13059-017-1382-0.

[16] I. Virshup, S. Rybakov, F. J. Theis, P. Angerer, and F. A. Wolf. anndata: Access and store annotated data matrices. *Journal of Open Source Software*, 9:4371, 2024. doi: 10.21105/joss.04371.

[17] A. Ebrahim, J. A. Lerman, B. O. Palsson, and D. R. Hyduke. Cobrapy: Constraints-based reconstruction and analysis for python. *BMC Systems Biology*, 7:74, 2013. doi: 10.1186/1752-0509-7-74.

[18] P. Virtanen, R. Gommers, T. E. Oliphant, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

[19] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9: 90–95, 2007. doi: 10.1109/MCSE.2007.55.

[20] M. L. Waskom. seaborn: statistical data visualization. *Journal of Open Source Software*, 6:3021, 2021. doi: 10.21105/joss.03021.