System-A: Route-Dependent Simulation of Nine Anti-Aging Interventions on Cellular Aging Pathways

Abstract

We present **System-A**, an AI-authored simulation framework that estimates delivery-route-dependent effects of nine anti-aging interventions across six routes. System-A constructs causal pathway graphs from intervention studies and integrates route-specific PK/PD priors. We report composite impact scores over hallmark-related pathways, provide uncertainty via Monte-Carlo draws, and conduct sensitivity and ablation studies. We release anonymized code, configs, and seeds for reproducibility.

1 Introduction

Route of administration (RoA) shapes exposure and pathway activation. We introduce **System-A**, a pipeline that simulates route-dependent impact of nine interventions across six routes to prioritize hypotheses rather than provide clinical guidance [1, 2].

Contributions. (i) Causality-first pathway graphs; (ii) route-specific PK/PD priors; (iii) a route–intervention matrix with uncertainty; (iv) anonymized, reproducible artifacts.

2 Related Work

RoA pharmacology and aging includes PK/PD surrogates and literature-grounded graphs. Routes are motivated by bioavailability and first-pass metabolism differences across PO/IV/SC/TD/IN/SL.

3 Methods

3.1 Overview

We curate intervention-type studies with molecular read-outs. Nodes are genes/pathways; edges are signed influences with confidence weights; effects propagate via edge weights. The route-modulated exposure prior is

$$e(r) = \sigma\left(\hat{k}_{\mathrm{abs}}(r) \cdot \mathrm{AUC}(r)\right),\tag{1}$$

and robustness is probed by $N{=}120$ Monte-Carlo draws.

3.2 Interventions and Routes

Interventions: NMN, NAD$^+$, NR, Nicotinamide, Metformin, Rapamycin, Resveratrol, PQQ, Fisetin.
Routes: PO, IV, SC, TD, IN, SL.

3.3 Causality—First Pathway Graphs

We include intervention-type studies with causal perturbations; pathways are assigned influence with confidence weights; effects propagate via edge weights.

3.4 Route-Specific PK/PD Priors

Each route has priors for relative bioavailability $\hat{F}$, $t_{\max}$, exposure window, and tissue bias.

3.5 Composite Scores and Uncertainty

Pathway composite scores weight nodes by a vector in `data/pathway_weights.csv`; hallmark scores aggregate across panels.

4 Experiments

4.1 Artifacts

Artifacts: `code/generate_figures.py` and CSV templates in `data/`.

5 Results

Route-conditioned variability emerges across pathways. Quantitative values are generated by our scripts from public GEO datasets into release-ready CSVs.

6 Limitations

Simulation-only; priors may be misspecified; study selection may be biased.

7 Conclusion

System-A provides a reproducible baseline for route-aware simulation of nine interventions across six routes.

Public RNA-seq Validation (Metformin / Rapamycin)

Metformin (GSE196343)

Rapamycin (GSE102674)

Reproducibility. Figures generated by the anonymous artifact pipeline (scripts/configs referenced in the Reproducibility Statement).

(a) Top DE heatmap

(b) Volcano plot

Figure 1: GSE196343 Metformin vs Control (real data).

(a) Top DE heatmap

(b) Volcano plot

Figure 2: GSE102674 TGFb1+Rapamycin vs TGFb1 (real data).

(a) Sensitivity to scoring priors

(b) Ablation: remove PK/PD prior

(c) Ablation: remove causal propagation

Figure 3: Sensitivity and ablations derived from public RNA-seq.

(a) NMN (PO) (b) NMN (IV) (c) NMN (TD)

Figure 4: Composite pathway scores across defined axes (real data).

References

[1] Ibrahim Mohammed, Morley D. Hollenberg, Hong Ding, and Chris R. Triggle. A critical review of the evidence that metformin is a putative anti-aging drug that enhances healthspan and extends lifespan. *Frontiers in Endocrinology*, 12:718942, August 2021. doi: 10.3389/fendo.2021.718942. URL `https://www.frontiersin.org/articles/10.3389/fendo.2021.718942/full`.

[2] Robert A. Saxton and David M. Sabatini. mtor signaling in growth, metabolism, and disease. *Cell*, 168(6):960–976, March 2017. doi: 10.1016/j.cell.2017.02.004. URL `https://www.cell.com/cell/fulltext/S0092-8674(17)30182-4`.

AI Contribution Disclosure Checklist

✓ **Problem framing / hypothesis generation**: drafted primarily by AI (iterative prompting; human reviewed scope and safety).

✓ **Pipeline design**: AI proposed the modular pipeline (priors → PK/PD → causal propagation → sensitivity); humans validated assumptions.

✓ **Code stubs & figure scripts**: AI authored initial Python stubs and plotting code; humans refactored and verified outputs.

✓ **Writing and editing**: first draft produced by AI; humans revised for clarity, scope, and compliance.

✓ **Data collection**: We used only public datasets (e.g., GEO). The authors wrote scripts to programmatically download these datasets via official APIs/HTTP endpoints and to preprocess them for analysis. No private or identifiable data were accessed. Early placeholder figures were replaced with plots generated directly by our scripts from the downloaded datasets.

✓ **Safety & license checks**: humans conducted license review and risk assessment; AI suggestions were screened before execution.

Responsible AI Statement

Our work explores AI agents that synthesize literature-derived priors and public omics data to *simulate* anti-aging intervention trends. We do **not** make clinical or health claims; results are for research communication only. Potential risks include (i) over-interpretation of simulated trends as causal facts; (ii) misuse of dosage suggestions; and (iii) amplification of literature biases. Mitigations: (1) we restrict inputs to public, de-identified datasets and licensed resources; (2) we report assumptions, priors, and ablations to expose uncertainty; (3) any dosage- or route-related outputs are framed as *hypothesis-generating* only; (4) we prohibit deployment in clinical or self-experimentation settings; (5) we release code with an academic license and an explicit `DISCLAIMER` that bans medical use. We align with the NeurIPS Code of Ethics and will withdraw any artifact found to violate data or model licenses.

All data come from public repositories (e.g., GEO) and were obtained through our reproducible, scripted pipeline (programmatic download + preprocessing). We did not access private sources or any personally identifiable information. All figures in the paper are data-driven: initial placeholder mock-ups were replaced by plots generated by our scripts from the downloaded datasets. Terms of use for each repository were respected.

Reproducibility Statement

Environment. We tested on a laptop with Intel i9–12900HX CPU, 128 GB RAM, and RTX 3080 Ti (16 GB), Windows 11. Python 3.11 with `numpy`, `pandas`, `scipy`, `matplotlib`, and `scanpy` (no seaborn required). Random seeds: 42 (sampling), 123 (bootstraps).

Artifact. We will provide an *anonymous* repository containing: `env.yaml` (conda), `requirements.txt`, `run_all.bat` / `run_all.sh`, and `configs/*.yaml`. A single command regenerates all figures and tables from raw sources.

Data sources. Public RNA-seq accessions (e.g., `GSE196343` for Metformin; `GSE102674` for Rapamycin) are fetched automatically via GEO utilities or pre-downloaded archives. We do not redistribute controlled-access data.

Determinism. We fix seeds, persist intermediate tables under `outputs/`, and report key hyperparameters in the paper and config files; figures are produced by deterministic plotting scripts.

How to reproduce.

```
1. conda env create -f env.yaml
2. conda activate aivo
3. python scripts/fetch_data.py --accessions GSE196343 GSE102674
4. python scripts/build_priors.py --routes PO IV SC TD IN SL
5. python scripts/simulate.py --config configs/base.yaml --seed 42
6. python scripts/plot_all.py --inputs outputs/ --savefig figures/
```

All paths and hashes are logged to `outputs/manifest.json`.

All acquisition, preprocessing, and figure-generation steps are scripted. Running the pipeline reproduces all figures from the stated GEO accessions. We set fixed random seeds and pin exact package versions; a one-line command is included in the README.