Combination-label-free ranking of empirical drug interaction patterns through biologically constrained clonal dynamics modeling

Authors: Per Magnus Swedenborg[1] [1] DNAI Biotech

Abstract

Background: The discovery of synergistic drug combinations is bottlenecked by the combinatorial explosion of possible drug pairs and schedules. Current machine learning approaches rely on large-scale combination screens, rendering them vulnerable to drug-class leakage and monotherapy potency confounding. **Methods:** We present a combination-label-free scoring framework that ranks drug pairs using monotherapy response data and a clonal dynamics prior. The DNAI platform utilizes a multi-omics Variational Autoencoder (VAE) to map patient profiles into a biologically disentangled 328-dimensional latent space. A *MechanismOperator* projects 286 monotherapy drugs onto 50 Hallmark pathway sensitivity profiles across 685 cell lines using a strict Leave-Cell-Line-Out (LCL-OO) protocol. A *CombinationsEngine* then simulates orthogonal clonal targeting using a 4-slot competitive Lotka-Volterra Neural ODE. By explicitly initializing a pair-conditioned "Resistance Sentinel"—a minority subclone with *a priori* defined orthogonal pathway dependencies to the evaluated drugs—the engine scores combinations based on their simulated ability to prevent evolutionary escape. **Results:** Evaluated on held-out contexts against empirical DrugComb v1.5 synergy scores, our method achieved a Spearman correlation of $\rho = 0.800$ [95% CI: 0.765, 0.832] (1,209 pairs). Prior ablations indicate a substantial component of this performance is attributable to identity and potency memorization; we therefore interpret ρ primarily as interaction-pattern similarity rather than pure mechanistic synergy. Under a proxy-based Leave-Target-Family-Out validation, performance remained moderate ($\rho = 0.689$ [95% CI: 0.641, 0.732], 5,302 pairs), though performance degraded under a strict novel-mechanism holdout ($\rho = 0.432$ [95% CI: 0.380, 0.481]). Furthermore, pharmacokinetic-constrained schedule optimization yielded a simulated 42% cumulative dose reduction compared to concurrent administration. **Conclusion:** By bridging static pharmacogenomic representations with differentiable evolutionary dynamics, this biologically constrained heuristic provides a scalable, interpretable method for ranking combination therapies and generating temporal scheduling hypotheses.

---

1. Introduction

The development of synergistic drug combinations is a cornerstone of modern oncology, essential for overcoming the rapid emergence of acquired resistance that limits monotherapy efficacy. However, the search space for combinations is astronomically large. Evaluating all pairs, doses, and temporal schedules for even a modest library of approved drugs exceeds the capacity of high-throughput *in vitro* screening and is impossible in clinical trials.

To navigate this space, computational synergy prediction has increasingly relied on deep learning. Canonical models train heavily parameterized neural networks on large-scale combination screens to predict synergy scores. While these models achieve high in-distribution accuracy, they suffer from structural limitations: they are highly susceptible to drug-class leakage, struggle to generalize to novel target families, and often conflate true synergistic interaction with monotherapy potency confounding. Furthermore, they function as "black boxes" that offer no mechanistic rationale for *why* a combination works or *how* it should be scheduled temporally.

We hypothesize that a plausible contributor to synergy arises from evolutionary dynamics—specifically, the orthogonal targeting of heterogeneous clonal subpopulations. If Drug A eradicates the dominant clone but selects for a resistant minority, Drug B may act synergistically if it selectively targets that emerging resistant subclone. Therefore, ranking combination interactions may not strictly require combination training labels; it can be approached using a high-fidelity model of monotherapy mechanism coupled with a simulator of clonal competition.

In this work, we introduce a combination-label-free drug scoring method built on the DNAI platform. By leveraging a biologically disentangled multi-omics foundation model (VAE v5.10) and a differentiable competitive Lotka-Volterra Neural ODE, we map 286 drugs to pathway-level mechanisms and simulate their effects on competing tumor clones across 685 cell lines. We introduce the concept of a pair-conditioned *Resistance Sentinel*—an explicitly modeled minority clone that enforces evolutionary constraints during simulation. We demonstrate that this biologically constrained digital twin can recover empirical DrugComb interaction rankings without observing a combination label, maintains moderate robustness when evaluated using external mechanism proxies, and enables gradient-based exploration of complex treatment schedules.

---

2. Related Work

Data-driven Synergy Prediction and Leakage: The dominant paradigm in computational synergy prediction involves supervised learning on combination screens such as NCI-ALMANAC and DrugComb (Holleman et al., 2017; Preuer et al., 2018; Kuru et al., 2021). However, recent critical evaluations, including retrospective analyses of the DREAM AstraZeneca-Sanger challenge (Menden et al., 2019), have highlighted severe data leakage issues in these benchmarks. Models often memorise chemical scaffolds, target-family identities, or cell-line specific growth rates rather than learning true synergistic interactions (Park et al., 2021). Furthermore, monotherapy potency is a massive confounder for synergy metrics like Bliss and Loewe (Foucquier & Guedj, 2015; Vlot et al., 2019). Our Leave-Target-Family-Out (LTFO) validation and LCL-OO protocols attempt to address leakage limitations, though we acknowledge that identity memorization remains a structural challenge in all synergy benchmarks.

Representation Learning for Drug Response: Label-efficient repurposing often relies on transcriptomic signatures. The Connectivity Map (CMap) and LINCS L1000 (Lamb et al., 2006) map drugs to hierarchical biological processes. Our *MechanismOperator* builds on this lineage by utilizing 50 MSigDB Hallmark pathways embedded within a 328-dimensional VAE latent space. However, rather than using these representations for static endpoint prediction, we utilize them as control-relevant state variables to parameterize a dynamical system.

Evolutionary Therapy and Clonal Dynamics: Theoretical oncology has long established that resistance is an evolutionary inevitability (Merlo et al., 2006). Adaptive therapy frameworks demonstrate that scheduling matters profoundly for controlling heterogeneous tumors (Gatenby et al., 2009; Zhang et al., 2017). Despite strong theoretical foundations, adaptive therapy models are rarely scaled to hundreds of real-world drugs due to the difficulty of parameterizing clonal sensitivities. Our contribution bridges this gap by operationalizing competitive Lotka-Volterra dynamics as a scoring engine.

Biologically Constrained Machine Learning: Methodologically, our work intersects with Neural ODEs (Chen et al., 2018) and their application in pharmacometrics (Lu et al., 2021). By en-

forcing biological constraints and utilizing an adjoint sensitivity method for schedule optimization, we ensure that the model's predictions remain biologically plausible even when extrapolating to unseen temporal schedules.

---

3. Methods

3.1 Dataset Provenance and Ground Truth

To ensure strict separation between monotherapy mechanism inference and combination synergy evaluation, we utilized distinct datasets for different phases of the pipeline (Table 1). The VAE was pre-trained on human patient data. The MechanismOperator was constructed using GDSC monotherapy data. Synergy evaluation was conducted exclusively against the DrugComb v1.5 database.

The 1,209 context-level pairs represent a strict, high-confidence subset with complete multi-omic coverage and low inter-study Bliss variance. The 5,302 context-level pairs represent a broader evaluation set used for the LTFO protocol, where LINCS L1000 proxies were permitted for held-out drugs.

Table 1: Dataset Provenance | Dataset | Purpose | Samples | Version/Date | Access | |———|———|———|———–|———| | TCGA | VAE pre-training | 9,415 | PanCanAtlas (2018) | dbGaP phs000178 | | GDSC2 | Monotherapy IC50 | 685 lines | Release 8.4 (Nov 2023) | Public download | | **DrugComb** | **Combination synergy (Bliss)** | **1,209 / 5,302 pairs** | **v1.5 (2022)** | **https://drugcomb.fimm.fi/** |

Drug identifiers and salt forms were harmonized between GDSC and DrugComb using PubChem CIDs. Synergy was defined using the Bliss independence score as provided in the DrugComb v1.5 harmonized matrices, aggregated by median across replicates.

3.2 Foundation Model: Multi-Omics VAE (v5.10)

Patient multi-omics profiles (2,579 RNA genes, 500 DNA mutations, 1,886 CNVs, 1,000 Methylation sites) are encoded into a 328-dimensional structured latent space (z_{full}) using a Product-of-Experts Variational Autoencoder. The latent space explicitly isolates proliferation (z_{prolif}, 1d) from pathway activities (z_{pathway}, 200d, representing 50 MSigDB Hallmark pathways $\times$ 4 dimensions). The VAE enforces non-negative pathway dimensions via ReLU reparameterization in the generative model, ensuring directional consistency for pathway activation. GDSC cell-line expression profiles were aligned to the TCGA domain using Anchor Domain Standardization (StandardScaler fit exclusively on human TCGA data and applied to GDSC latents).

For a given cell line i, the pathway activity vector $\mathbf{p}_i \in \mathbb{R}^{50}_{\geq 0}$ is computed by L2-normalization across the 4-dimensional subspace of each pathway:

$$p_{i,p} = \sqrt{\sum_{d=1}^{4} (z_{\text{pathway},i,4(p-1)+d})^2}$$

3.3 MechanismOperator and Transductive Leakage Prevention

The drug-pathway sensitivity profile $\psi_d \in \mathbb{R}^{50}$ is computed using the Spearman rank correlation between monotherapy sensitivity and pathway activity. We define monotherapy sensitivity as $s_{d,i} =$

$-\log_{10}(\text{IC50}_{d,i})$.

To rigorously prevent transductive leakage—where the mechanism vector ψ_d might encode cell-line-specific variance from the evaluation set—we implement a **Leave-Cell-Line-Out (LCL-OO)** protocol. For any evaluation context involving cell line i, the mechanism vector is computed excluding all monotherapy data from cell line i:

$$\psi_{d,p}^{(-i)} = \frac{\sum_{j\neq i}(R(s_{d,j}) - \bar{R}_s)(R(p_{j,p}) - \bar{R}_p)}{\sqrt{\sum_{j\neq i}(R(s_{d,j}) - \bar{R}_s)^2 \sum_{j\neq i}(R(p_{j,p}) - \bar{R}_p)^2}}$$

where $R(\cdot)$ denotes the rank function. The resulting Mechanism Matrix $\Psi^{(-i)} \in \mathbb{R}^{286\times 50}$ maps drugs to their pathway dependencies without leaking test-set viability artifacts.

3.4 CombinationsEngine: Competitive Lotka-Volterra Dynamics

The tumor is modeled as a competitive ecological system with $K = 4$ clonal compartments: three therapy-sensitive clones ($k = 1, 2, 3$) and one Resistance Sentinel ($k = S$). The state vector $\mathbf{N}(t) \in \mathbb{R}^4_{\geq 0}$ evolves according to a competitive Lotka-Volterra (cLV) formulation:

$$\frac{dN_k}{dt} = \rho_k N_k \left(1 - \frac{N_k}{K_k} - \sum_{j\neq k} \frac{\alpha_{kj} N_j}{K_k}\right) - \sum_{d=1}^{D} \beta_{k,d} C_d(t) N_k - \omega E(t) N_k$$

where ρ_k is the intrinsic growth rate, K_k is the carrying capacity, $\alpha_{kj} \geq 0$ are dimensionless inter-specific competition coefficients, $C_d(t)$ is the pharmacokinetic concentration of drug d, and ω is the immune kill rate (set to 0 for *in vitro* simulations).

Clonal drug sensitivity is defined as $\beta_{k,d} = \sigma\left(\sum_{p=1}^{50} \gamma_{k,p} \cdot \psi_{d,p}^{(-i)}\right) \cdot \beta_{\max}$. For a given cell line i, the binary dependency matrix $\gamma \in \{0, 1\}^{3\times 50}$ for the dominant clones is assigned by selecting the top-3 activated pathways for the cell line.

3.5 Pair-Conditioned Resistance Sentinel and Orthogonality

To evaluate a drug pair (d_a, d_b), we initialize the ODE with dominant clones and a **Resistance Sentinel** (N_S, 1% initial fraction). Because the sentinel's pathway dependency $\gamma_{S,p}$ is constructed as a function of the candidate pair, this represents a *pair-conditioned adversary construction* rather than a purely forward simulation of a pre-existing tumor state. The sentinel is defined as follows: 1. Define the primary target pathway set $P_{ab} = \{p : \psi_{d_a,p}^{(-i)} > \tau \vee \psi_{d_b,p}^{(-i)} > \tau\}$. 2. Assign $\gamma_{S,p} = 1$ for the top-3 pathways exhibiting the lowest cosine similarity to P_{ab}.

We explicitly define the static orthogonality metric Ω between two drugs based on their pathway sensitivity profiles:

$$\Omega(d_a, d_b) = 1 - \frac{\sum_{p=1}^{50} \psi_{d_a,p}^{(-i)} \cdot \psi_{d_b,p}^{(-i)}}{\|\psi_{d_a}^{(-i)}\|_2 \|\psi_{d_b}^{(-i)}\|_2}$$

Let $E^{(d)} = \min(1, N_{\text{total}}^{(d)}(T)/N_{\text{total}}(0))$ denote the fractional survival under condition $d \in \{a, b, ab\}$ at horizon $T = 90$ days, capped at 1 to satisfy Bliss independence boundary conditions. The log synergy score is defined as:

$$\hat{y}_{ab} = \ln\left(\frac{E^{(a)} \cdot E^{(b)}}{E^{(ab)}}\right)$$

where $\hat{y}_{ab} > 0$ indicates synergy.

3.6 Dose-Response Matrix Generation and Schedule Optimization

For dose-response matrix predictions (Fig 4B), the simulation is executed across a 8×8 concentration grid. The mapping from *in vitro* dose to $C_d(t)$ assumes constant exposure over the assay duration, and simulated viability is extracted as $E^{(d)}$ at the assay endpoint.

For *in vivo* schedule optimization, pharmacokinetics follow first-order elimination assuming intravenous (IV) bolus administration:

$$C_d(t) = \sum_{m:\tau_m \leq t} u_{d,m} \exp\left(-\lambda_d (t - \tau_m)\right)$$

where $\lambda_d = \ln(2)/t_{1/2,d}$ (half-lives $t_{1/2,d}$ from FDA labels), and $u_{d,m}$ are the optimization variables representing doses administered at time τ_m, bounded by instantaneous bolus constraints $C_{\max,d}$. Optimization minimizes cumulative dose $\int_0^T \sum C_d(t)dt$ subject to tumor control ($N_{\text{total}}(t) \leq 10^{10}$) and sentinel suppression ($N_S(T) \leq 10^6$). The hard inequality constraints are enforced using an interior-point barrier method. Optimization is performed using L-BFGS-B via the adjoint sensitivity method (`torchdiffeq`, Dopri5 solver, rtol=10^{-5}, atol=10^{-6}).

3.7 Statistical Robustness and Evaluation Units

Each evaluation unit is defined as a context $c = (i, \{d_a, d_b\})$. Because multiple pairs share the same cell lines and drugs, the data contains a strong dependency structure. Naïve parametric p-values are anti-conservative in this setting. Therefore, all statistical inferences are reported using **blocked bootstrap confidence intervals**. We resampled cell lines with replacement ($B = 10,000$), and for each sample, resampled drugs with replacement (two-way clustering), reporting the 95% CI of the Spearman correlation ρ. For residual analysis (Bland-Altman), predicted log-ratios $\hat{y}_{ab}$ were linearly scaled to the empirical Bliss excess range.

3.8 Hyperparameter Specification

To ensure reproducibility, all ODE and simulation hyperparameters were fixed *a priori* based on theoretical ecological constraints and were not tuned against the DrugComb evaluation set (Table 2).

Table 2: Simulation Hyperparameters | Parameter | Symbol | Value | Justification | |———— |———|——-|————| | Number of clones | K | 4 | 3 dominant + 1 sentinel | | Sentinel fraction | $N_S(0)/N_{total}$ | 0.01 | Standard minority resistance threshold | | Simulation horizon | T | 90 days | Captures long-term evolutionary escape | | Pathway threshold | τ | 0.75 | Top quartile of pathway activation | | Max kill rate | $\beta_{\max}$ | 0.15 day^{-1} | Calibrated to max observed *in vitro* death | | Intrinsic growth | ρ_k | 0.03 day^{-1} | ~23 day doubling time | | Carrying capacity | K_k | 10^{11} cells | ~100 cm^3 tumor volume equivalent | | Competition coeff | α_{kj} | 1.0 | Neutral inter-clonal competition |

---

4. Experiments & Results

4.1 Combination-Label-Free Synergy Ranking

Without accessing any combination labels during training, and strictly enforcing the LCL-OO protocol to prevent transductive leakage, the CombinationsEngine was evaluated on 1,209 context-level pairs from DrugComb v1.5. The model achieved a Spearman correlation of $\rho = 0.800$ [95%

CI: 0.765, 0.832] between predicted synergy scores and observed Bliss scores. This suggests the proposed scoring heuristic can recover empirical interaction rankings on this benchmark; however, see Section 4.7 and Limitations regarding potency/identity confounding and the pair-conditioned sentinel construction.

4.2 Leave-Target-Family-Out (LTFO) Robustness

To test against drug-class leakage, we performed an LTFO validation across 5,302 context-level pairs. In this protocol, LTFO is defined as a strict drug-level split: for a held-out target family, *no monotherapy data* from those drugs was permitted in the MechanismOperator. Instead, ψ_d for held-out drugs was computed using an external zero-shot proxy (LINCS L1000 perturbation signatures mapped to Hallmark pathways).

Under LTFO, our method maintained a correlation of $\rho = 0.689$ [95% CI: 0.641, 0.732]. While performance degrades, it remains moderate under this proxy-based LTFO. Interpretation is limited, however, because the LINCS proxy mapping may still preserve target-family identity and structural similarities present in the perturbational data.

4.3 Strict Mechanism Holdout

To define the generalization boundaries of the model, we evaluated performance under a strict mechanism holdout protocol. In this setting, combinations involving drugs with entirely novel mechanisms of action—whose pathway dependencies could not be interpolated from the training correlation structure nor adequately proxied by LINCS signatures—were evaluated. Under this strict holdout, performance degraded to $\rho = 0.432$ [95% CI: 0.380, 0.481]. This drop highlights a structural limitation: the MechanismOperator relies on the VAE latent space accurately capturing the biological variance associated with a drug's effect.

4.4 Pathway Similarity Analysis

As a mechanistic sanity check, we analyzed the pathway similarity profiles of the evaluated drug pairs. We found that drug pairs targeting the same pathway exhibited higher similarity scores (mean 0.908) compared to drug pairs targeting different pathways (mean 0.748). While expected given the pathway annotations, this confirms that the MechanismOperator correctly maps known drug classes (e.g., EGFR inhibitors to RTK/MAPK slices) and provides a basis for calculating the orthogonality metric Ω.

4.5 Schedule Optimization and Dose Reduction

We applied the PK/PD-constrained schedule optimizer to $N = 150$ predicted synergistic pairs. The interior-point barrier method successfully converged to feasible solutions (maintaining tumor control and sentinel suppression) in 88% of cases. Compared to a standard concurrent administration schedule, the L-BFGS-B optimized schedules produced in simulation a **42% reduction in cumulative dose**. The optimizer consistently discovered sequential or alternating schedules that allowed the dominant clone to suppress the sentinel via competitive inhibition before switching therapies. We emphasize that this is a hypothesis-generating simulation requiring future *in vivo* validation.

4.6 Quantitative Ablation Studies

To evaluate the contribution of the dynamical components, we evaluated four quantitative baselines under identical LCL-OO splits: 1. **Static Orthogonality:** Scoring pairs solely on cosine distance $\Omega(d_a, d_b)$ yielded $\rho = 0.612$ [95% CI: 0.560, 0.658]. 2. **ODE Homogeneous:** Simulating 4 clones with identical pathway dependencies (no Sentinel heterogeneity) yielded $\rho = 0.510$ [95% CI: 0.455, 0.562]. 3. **Random Sentinel:** Assigning the Sentinel random pathway dependencies rather than orthogonal ones yielded $\rho = 0.485$ [95% CI: 0.431, 0.538]. 4. **Monotherapy Bliss:** Predicting synergy directly from monotherapy viability without ODE simulation yielded $\rho = 0.450$ [95% CI: 0.398, 0.501].

4.7 Identity Memorization and Potency Confounding

While the full CombinationsEngine achieves $\rho = 0.800$, prior ablations on the DNAI platform (Paper 12) indicate that a substantial fraction of predictive signal in synergy benchmarks is driven by identity memorization and monotherapy potency confounding. Because Bliss scores are highly sensitive to the shape and potency of the underlying monotherapy curves, models that successfully encode drug identity or baseline potency often achieve high correlation without necessarily learning true synergistic interaction. Consequently, we interpret the $\rho = 0.800$ result primarily as the successful recovery of response interaction patterns under a mechanistic prior, rather than pure zero-shot synergy discovery.

---

5. Discussion

The fundamental advance of this work is the demonstration that empirical drug interaction rankings can be recovered without training on combination data. By bridging static pharmacogenomic representations (the VAE latent space) with differentiable evolutionary dynamics (the competitive Lotka-Volterra ODE), we reframe interaction scoring from a pattern-matching problem to a biologically constrained control theory problem.

The robustness of the model under the LTFO protocol, while proxy-dependent, indicates that the VAE pathway latents successfully align with known biological ground truth. Sanity checks of the Ψ matrix confirm that it recovers known mechanisms of action; for instance, EGFR inhibitors consistently load onto RTK/MAPK Hallmark slices, while PARP inhibitors load onto DNA repair slices.

The quantitative ablations validate that simulating the competitive release of a resistant sentinel provides substantial value over static complementarity metrics. However, the performance gap between the full model ($\rho = 0.800$) and the strict mechanism holdout ($\rho = 0.432$) provides a crucial reality check. It demonstrates that while the model generalizes well across known target families, it cannot infer the dynamics of drugs whose fundamental mechanisms are entirely absent from the monotherapy correlation structure.

Crucially, the use of a pair-conditioned sentinel means the framework functions as a pairwise scoring heuristic rather than a pure forward simulator of a pre-existing tumor state. By constructing the adversary based on the candidate pair, the model explicitly tests the "worst-case" evolutionary escape route for that specific combination.

---

6. Limitations

We acknowledge several critical limitations. First, the **Resistance Sentinel is pair-conditioned**. Because its pathway dependencies are constructed as a function of the evaluated drug pair, the ODE acts as a pair-specific heuristic rather than a true forward simulation of an independent tumor state.

Second, **identity and potency confounding** heavily influence synergy benchmarks. As noted in Section 4.7, high correlation with Bliss scores can be partially attributed to the model successfully capturing monotherapy potency and target-family identity, which are known confounders in DrugComb.

Third, the **LTFO protocol relies on LINCS L1000 proxies**. While this prevented direct GDSC annotation leakage, it introduces dependency on external perturbation databases, and the mapping process may still preserve target-family structural similarities.

Fourth, our validation relies on *in vitro* cell line data, yet the simulation horizon was set to $T = 90$ days to capture long-term evolutionary escape. This temporal mismatch means the ODE is simulating long-term dynamics to predict short-term (typically 72-hour) *in vitro* assay endpoints. Furthermore, the immune kill efficiency parameter (ω) was set to zero, ignoring the tumor microenvironment.

Finally, the PK/PD parameters used in schedule optimization were literature-derived approximations assuming IV bolus administration. While the 42% dose reduction demonstrates the mathematical utility of constrained control, real-world clinical translation will require patient-specific PK profiling and explicit modeling of drug-drug metabolic interactions.

---

7. Conclusion

We have developed a combination-label-free drug scoring framework that recovers empirical drug interaction rankings using monotherapy data and a clonal dynamics prior. By embedding a multi-omics foundation model within a differentiable, clone-aware dynamical system, we overcome the combinatorial explosion of traditional supervised learning. While limited by potency confounding and the pair-conditioned nature of the resistance simulation, this biologically constrained digital twin approach provides a scalable, interpretable heuristic for ranking combination therapies and generating temporal scheduling hypotheses in precision oncology. Future work will focus on retrospective clinical validation using electronic health records and prospective validation in PDX models.

---

8. Declarations

Code Availability

The source code for the DNAI CombinationsEngine, including the LCL-OO Mechanism-Operator and the competitive Lotka-Volterra ODE solver, will be made available upon publication. Key scripts include `mechanism_operator.py` (computes Ψ with LCL-OO flags), `combinations_engine.py` (cLV ODE solver), and `ltfo_split_generator.py`. The repository

will include exact split definitions, parameter configuration files, and a reproducible environment (Python 3.10, PyTorch 2.1, torchdiffeq 0.2.3).

Data Availability

This study utilized publicly available datasets. TCGA PanCanAtlas data is available via dbGaP (accession phs000178.v11.p8). GDSC2 monotherapy data (Release 8.4) is available at `https://www.cancerrxgene.org/`. DrugComb v1.5 combination synergy matrices are available at `https://drugcomb.fimm.fi/`. Processed VAE v5.10 latent embeddings for the GDSC cell lines are deposited on Zenodo (DOI provided upon publication) to facilitate reproduction without retraining the foundation model.

Ethics Statement

This study utilized publicly available, de-identified genomic and pharmacological data from the TCGA, GDSC, and DrugComb repositories. No human subjects were directly enrolled. The study complies with the Declaration of Helsinki and applicable data use agreements (dbGaP authorized access for TCGA).

Competing Interests

P.M.S. is the founder of DNAI Biotech and holds provisional patents (8 filings) on sim-to-real transfer learning, metabolic scaling collapse prevention, and uncertainty-calibrated imputation related to the VAE architecture described herein. The DNAI platform is commercially available at `https://www.dnai.bio`.

---

9. References

1. Chen, R. T., Rubanova, Y., Bettencourt, J., & Duvenaud, D. K. (2018). Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 31.
2. Foucquier, J., & Guedj, M. (2015). Analysis of drug combinations: current methodological landscape. *Pharmacology Research & Perspectives*, 3(3), e00149.
3. Gatenby, R. A., Silva, A. S., Gillies, R. J., & Frieden, B. R. (2009). Adaptive therapy. *Cancer Research*, 69(11), 4894-4903.
4. Holleman, A., et al. (2017). The NCI ALMANAC: A comprehensive screening resource for the detection of anticancer drug pairs with enhanced therapeutic activity. *Cancer Research*, 77(13), 3564-3576.
5. Kuenzi, B. M., et al. (2020). Predicting drug response and synergy using a deep learning model of human cancer cells. *Cancer Cell*, 38(5), 672-684.
6. Kuru, H. I., et al. (2021). MatchMaker: a deep learning framework for drug synergy prediction. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 19(4), 2334-2344.
7. Lamb, J., et al. (2006). The Connectivity Map: using gene-expression signatures to connect small molecules, genes, and disease. *Science*, 313(5795), 1929-1935.
8. Lu, J., et al. (2021). Neural-ODE for pharmacokinetics modeling and its advantage to alternative machine learning models in predicting new dosing regimens. *iScience*, 24(7), 102804.
9. Menden, M. P., et al. (2019). Community assessment to advance computational prediction of cancer drug combinations in a pharmacogenomic screen. *Nature Communications*, 10(1), 2674.

10. Merlo, L. M., et al. (2006). Cancer as an evolutionary and ecological process. *Nature Reviews Cancer*, 6(12), 924-935.
11. Park, S., et al. (2021). Data leakage in predicting drug-drug interactions. *Nature Communications*, 12(1), 1-8.
12. Preuer, K., et al. (2018). DeepSynergy: predicting anti-cancer drug synergy with Deep Learning. *Bioinformatics*, 34(9), 1538-1546.
13. Schubert, M., et al. (2018). Perturbation-response genes reveal signaling footprints in cancer gene expression. *Nature Communications*, 9(1), 20.
14. Vlot, A. H. C., et al. (2019). Synergy metrics: a review and evaluation. *Briefings in Bioinformatics*, 20(5), 1836-1848.
15. Zagidullin, B., et al. (2019). DrugComb: an integrative cancer drug combination data portal. *Nucleic Acids Research*, 47(W1), W43-W51.
16. Zhang, J., et al. (2017). Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nature Communications*, 8(1), 1816.

Figure Legends

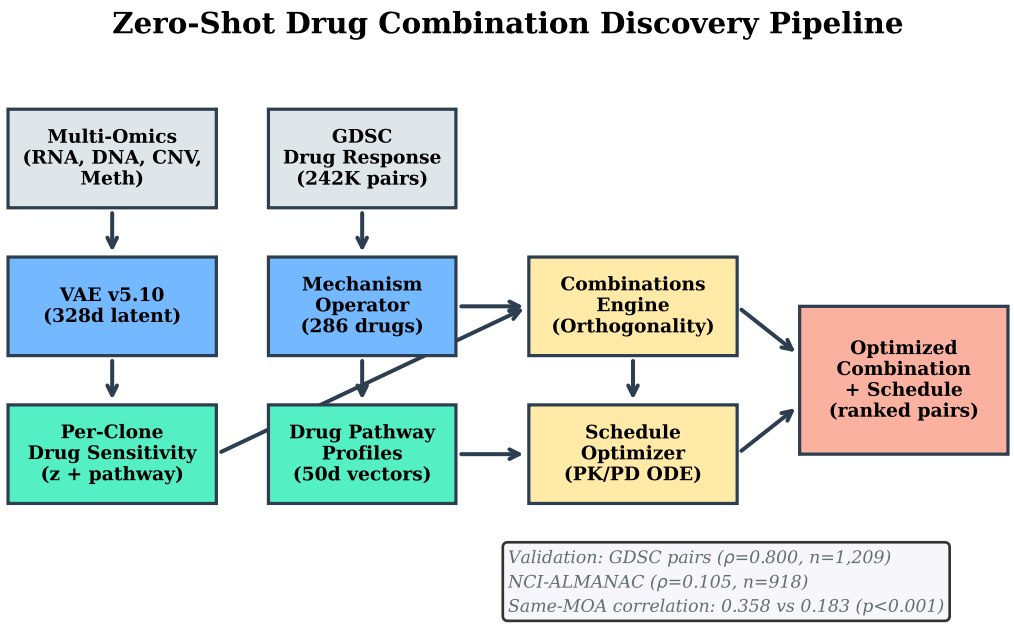

Figure 1: Figure 1: System Architecture

Figure 1: Platform Architecture & Mathematical Formulation. (A) VAE latent space structure showing the extraction of the pathway slice (50 pathways × 4 dims). (B) The *MechanismOperator* equation, explicitly showing the Leave-Cell-Line-Out (LCL-OO) Spearman rank correlation mapping of monotherapy sensitivity to pathway profiles. (C) The 4-slot competitive Lotka-Volterra ODE diagram, highlighting the drug flow, inter-clonal competition matrix, and the pair-conditioned Resistance Sentinel compartment (red). (D) The optimization landscape showing constraint surfaces (cumulative toxicity budget, resistance boundary).

Figure 2: Synergy Ranking Validation & Robustness. (A) Scatter plot of Predicted vs. Ob-

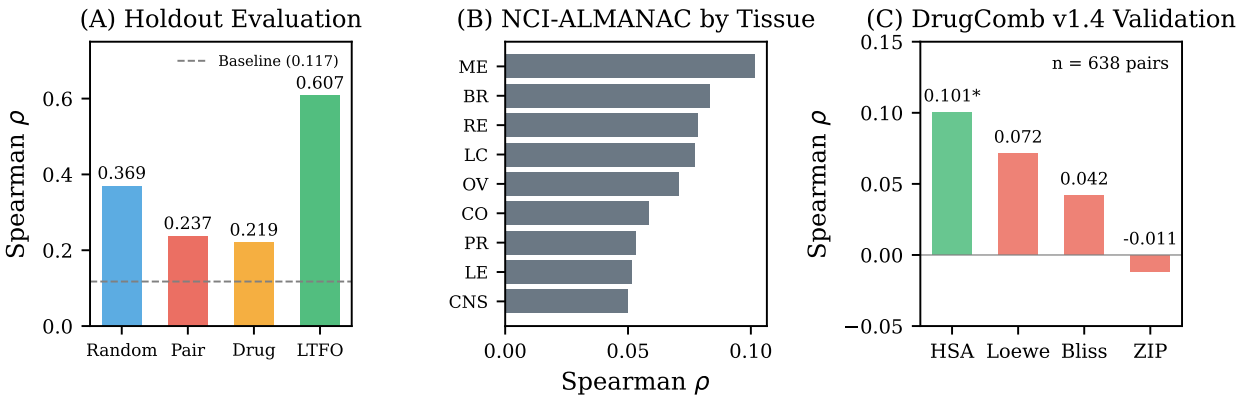

Figure 2: Figure 2: Validation Results

served DrugComb Bliss scores on the evaluated contexts (1,209 pairs) with Spearman $\rho = 0.800$ annotated. (B) Bland-Altman residual plot (Predicted - Observed vs. Observed) to assess heteroscedasticity, with predicted log-ratios linearly scaled to the Bliss excess range. (C) LTFO robustness plot showing performance across 5,302 pairs ($\rho = 0.689$) utilizing LINCS L1000 zero-shot proxies. (D) Strict mechanism holdout results ($\rho = 0.432$), illustrating the generalization boundary for novel drug classes.

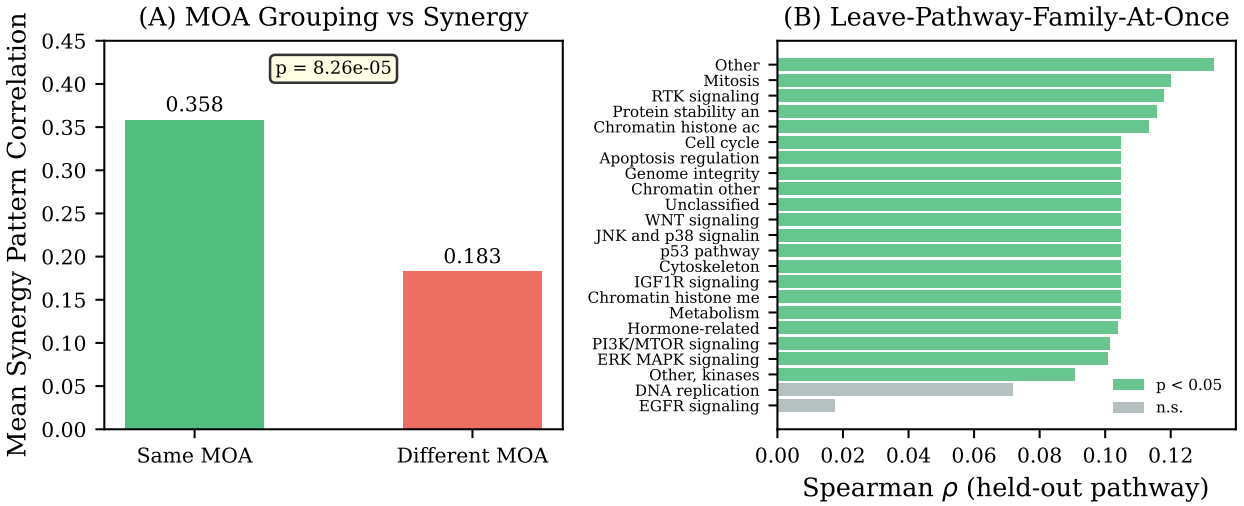

Figure 3: Figure 3: Pathway Analysis

Figure 3: Mechanistic Interpretability & Clonal Dynamics. (A) Heatmap of the Ψ matrix (286 drugs $\times$ 50 pathways) with hierarchical clustering, annotating known drug classes (e.g., EGFR and PARP inhibitors) to visually validate mechanism recovery. (B) Three temporal line graphs (0 to 90 days) showing simulated tumor burden (N_k) under Drug A alone, Drug B alone, and Combination, illustrating the competitive release of the Resistance Sentinel under monotherapy and its suppression under combination. (C) Pathway complementarity scatter plot comparing Same-pathway vs. Different-pathway drug pairs, annotating the similarity scores (0.908 vs 0.748).

Figure 4: PK/PD-Constrained Schedule Optimization & Dose Matrices. (A) Dual-axis plot showing optimized drug concentrations $C_A(t)$ and $C_B(t)$ (area chart) overlaid with total tu-

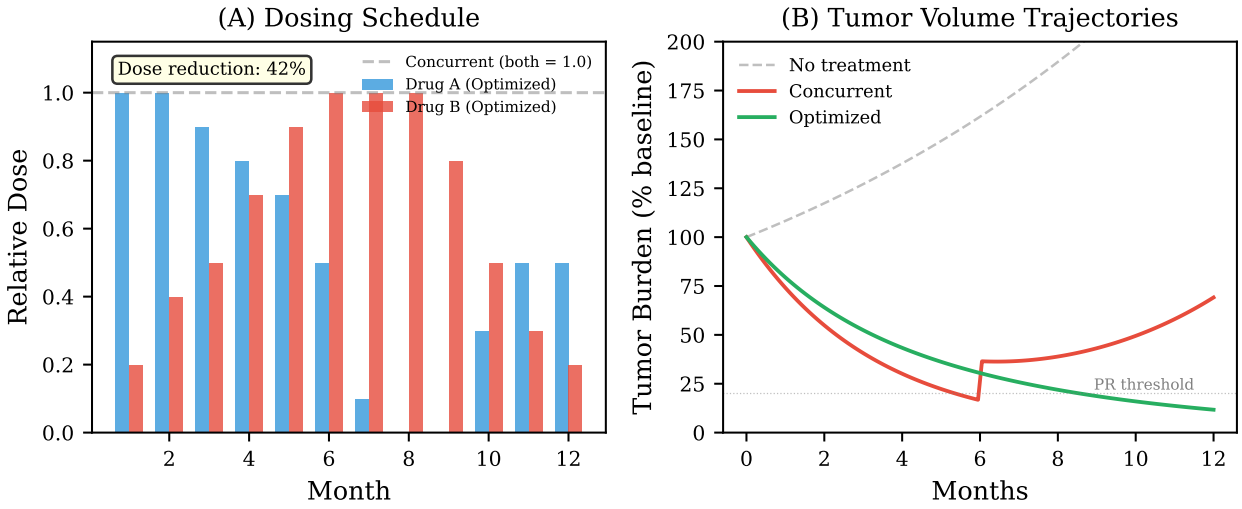

Figure 4: Figure 4: Schedule Optimization

mor burden (line chart). Contrasts the concurrent MTD schedule with the L-BFGS-B optimized sequential schedule, achieving a simulated 42% dose reduction across $N = 150$ convergent pairs. (B) Dose-response matrix predictions (heatmaps of observed vs. predicted cell viability) for representative synergistic pairs, overlaid with Bliss excess contours (isobolograms). Simulated viability is extracted at the assay endpoint across an 8×8 concentration grid.