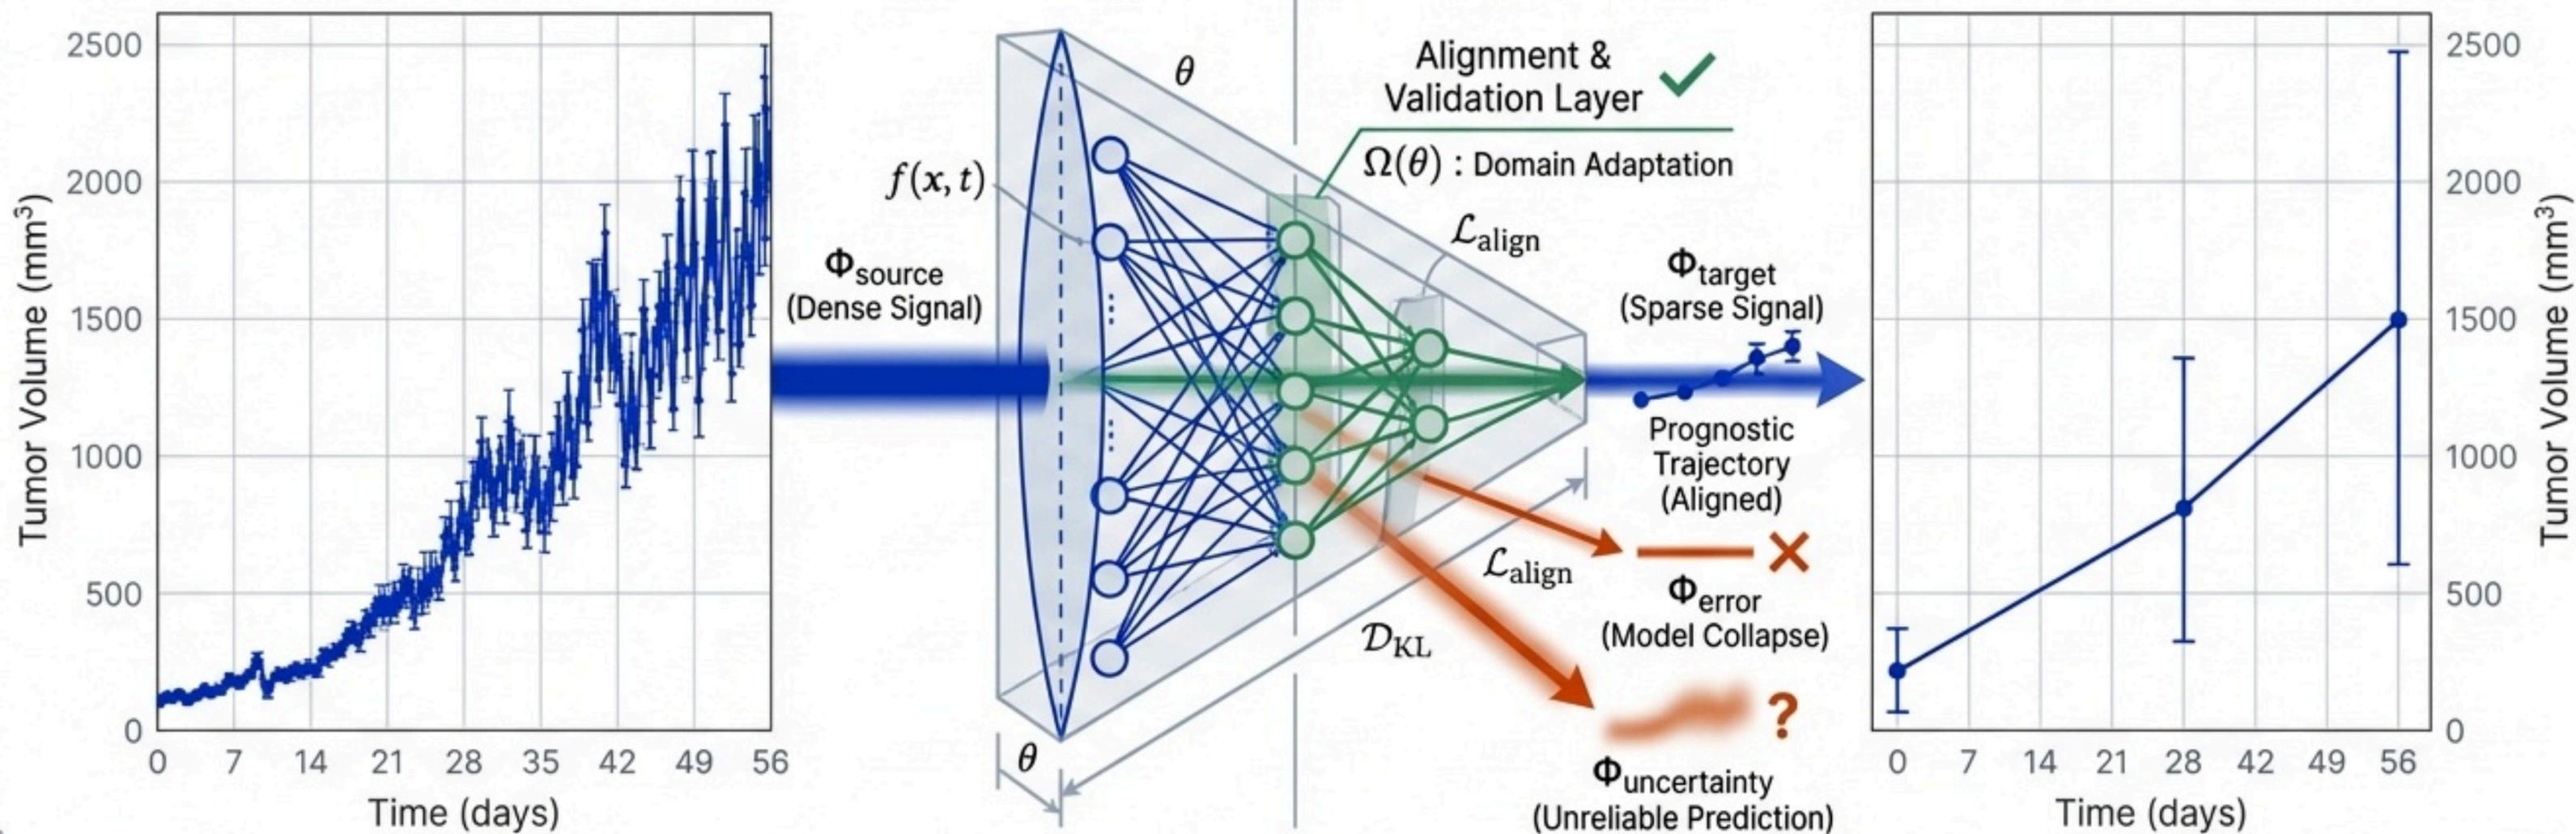


Physics-Constrained Sim-to-Real Transfer Learning in Computational Oncology.

An Integrated Architecture for Cross-Species Tumor Trajectory Simulation

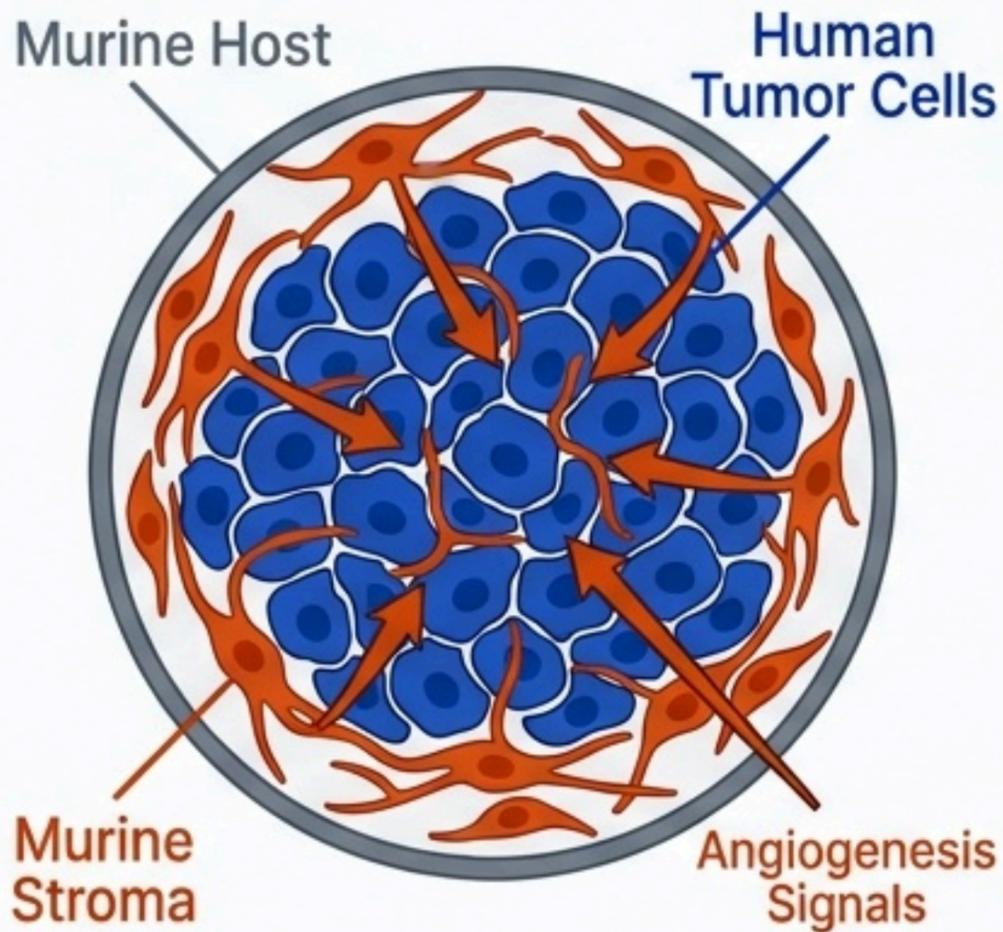
SOURCE DOMAIN: PRECLINICAL

TARGET DOMAIN: CLINICAL



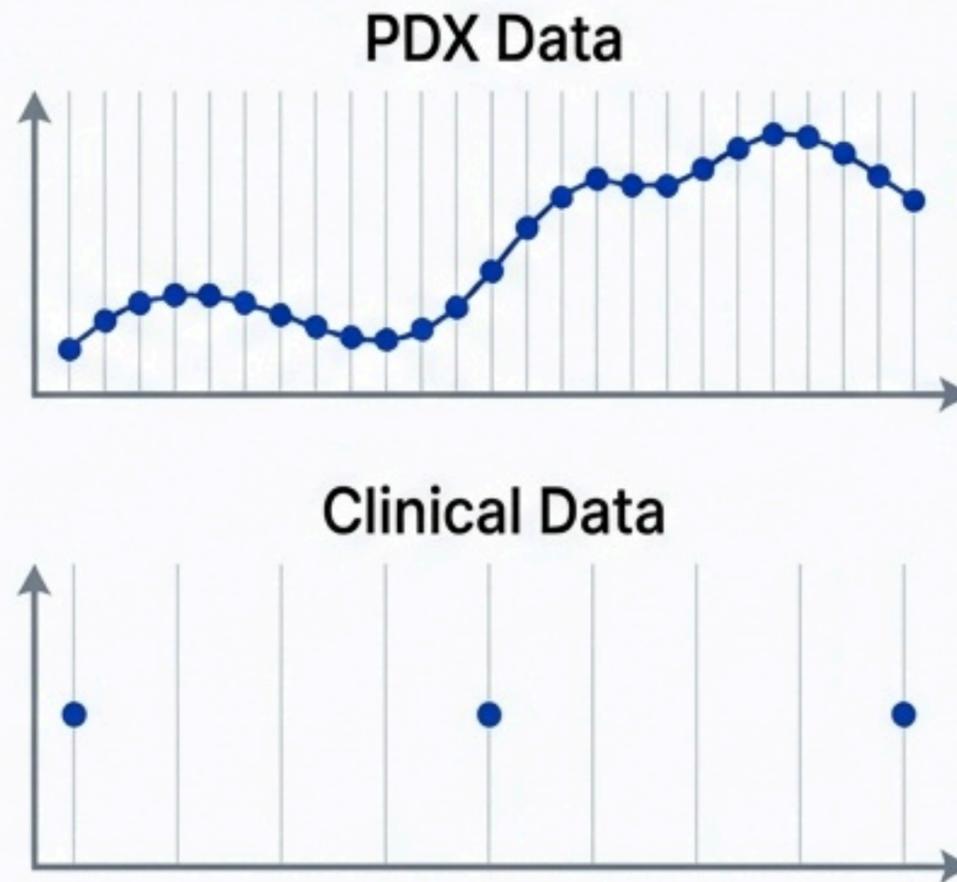
The Translational Gap: Why Preclinical Models Fail in the Clinic

1. Stroma Replacement



Models inadvertently learn murine angiogenesis drivers. Direct transfer leads to Negative Transfer due to microenvironment shift.

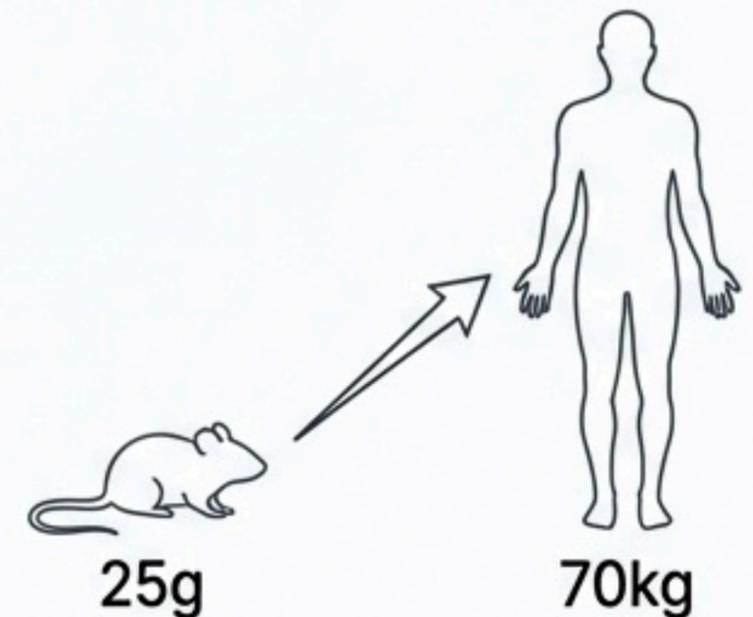
2. Data Asymmetry



Dense Source (Daily/Weekly) vs. Sparse Target (6-12 Weeks). Models overfit to high-frequency noise absent in clinical endpoints.

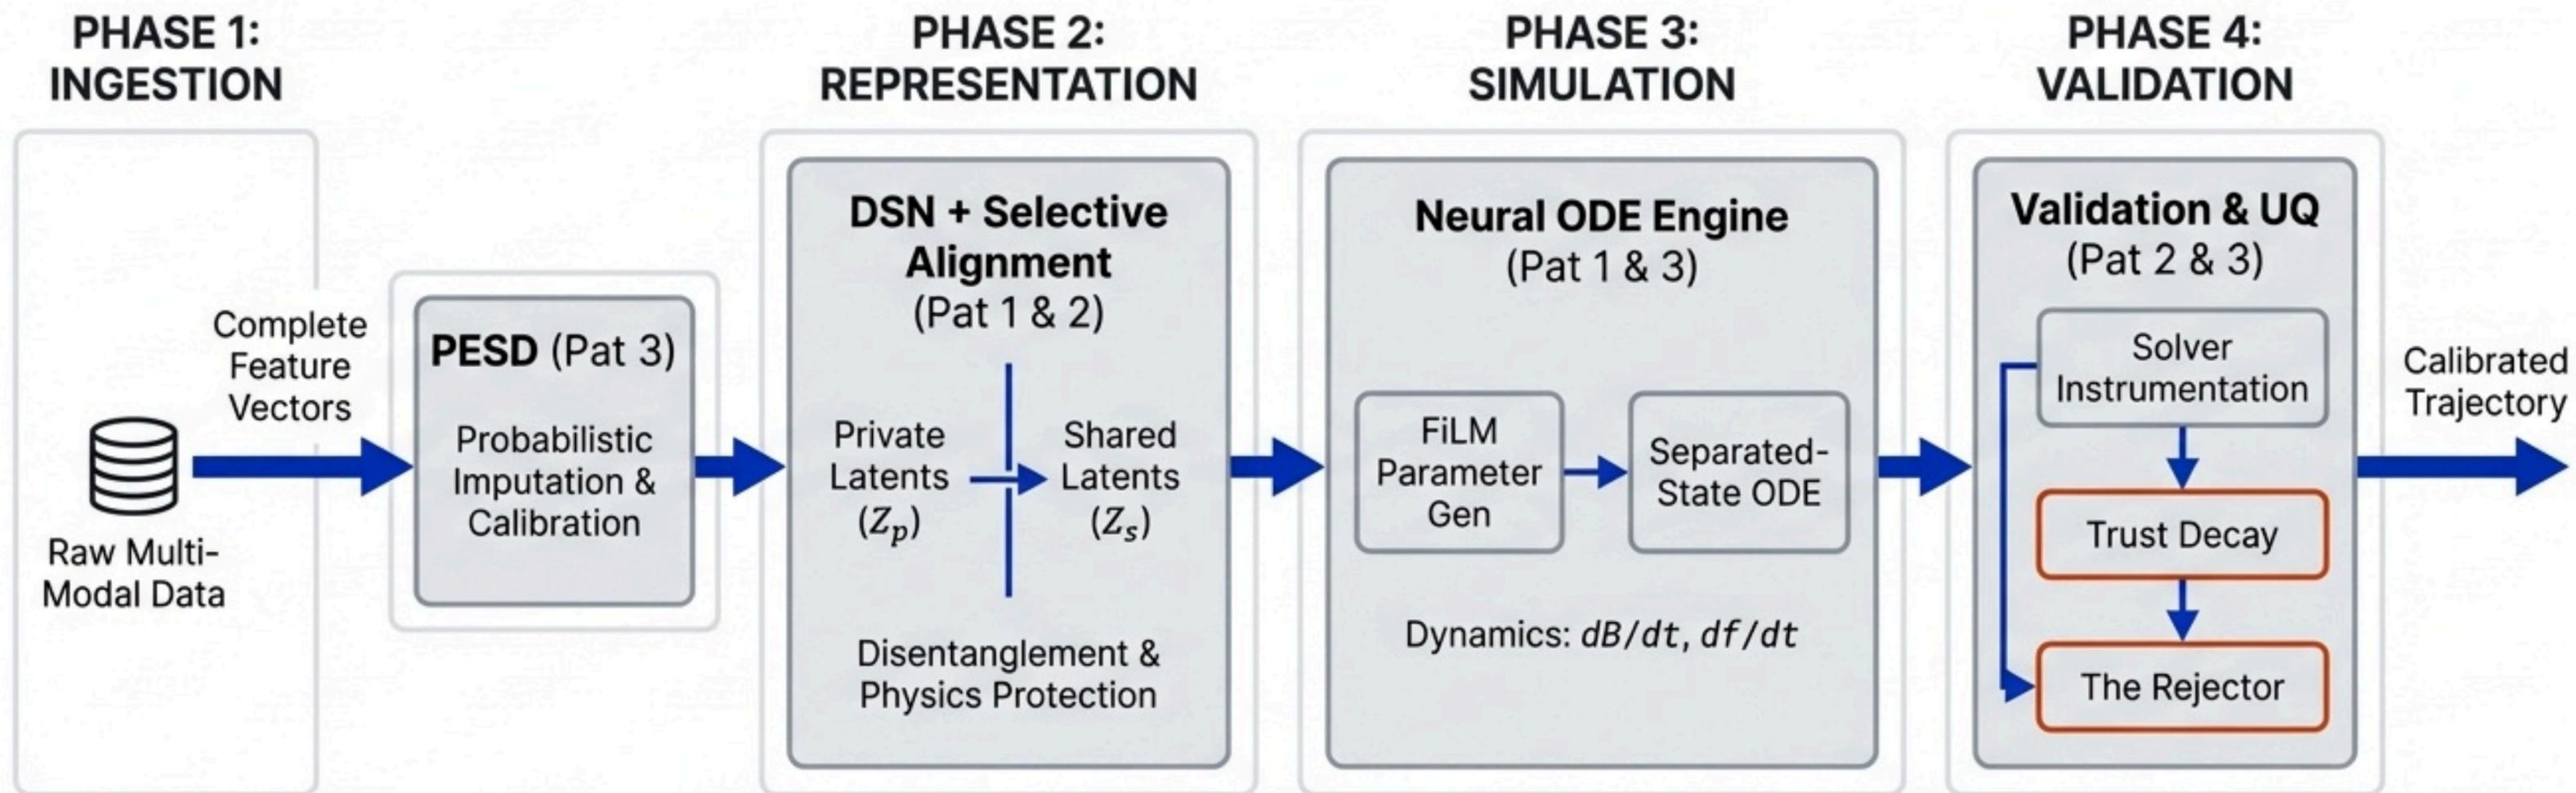
3. Allometric Mismatch

$$r \propto M^{-1/4}$$



Metabolic rates scale with mass. Murine proliferation signals are metabolically accelerated, violating Kleiber's Law if transferred directly.

System Architecture: The Data Lifecycle

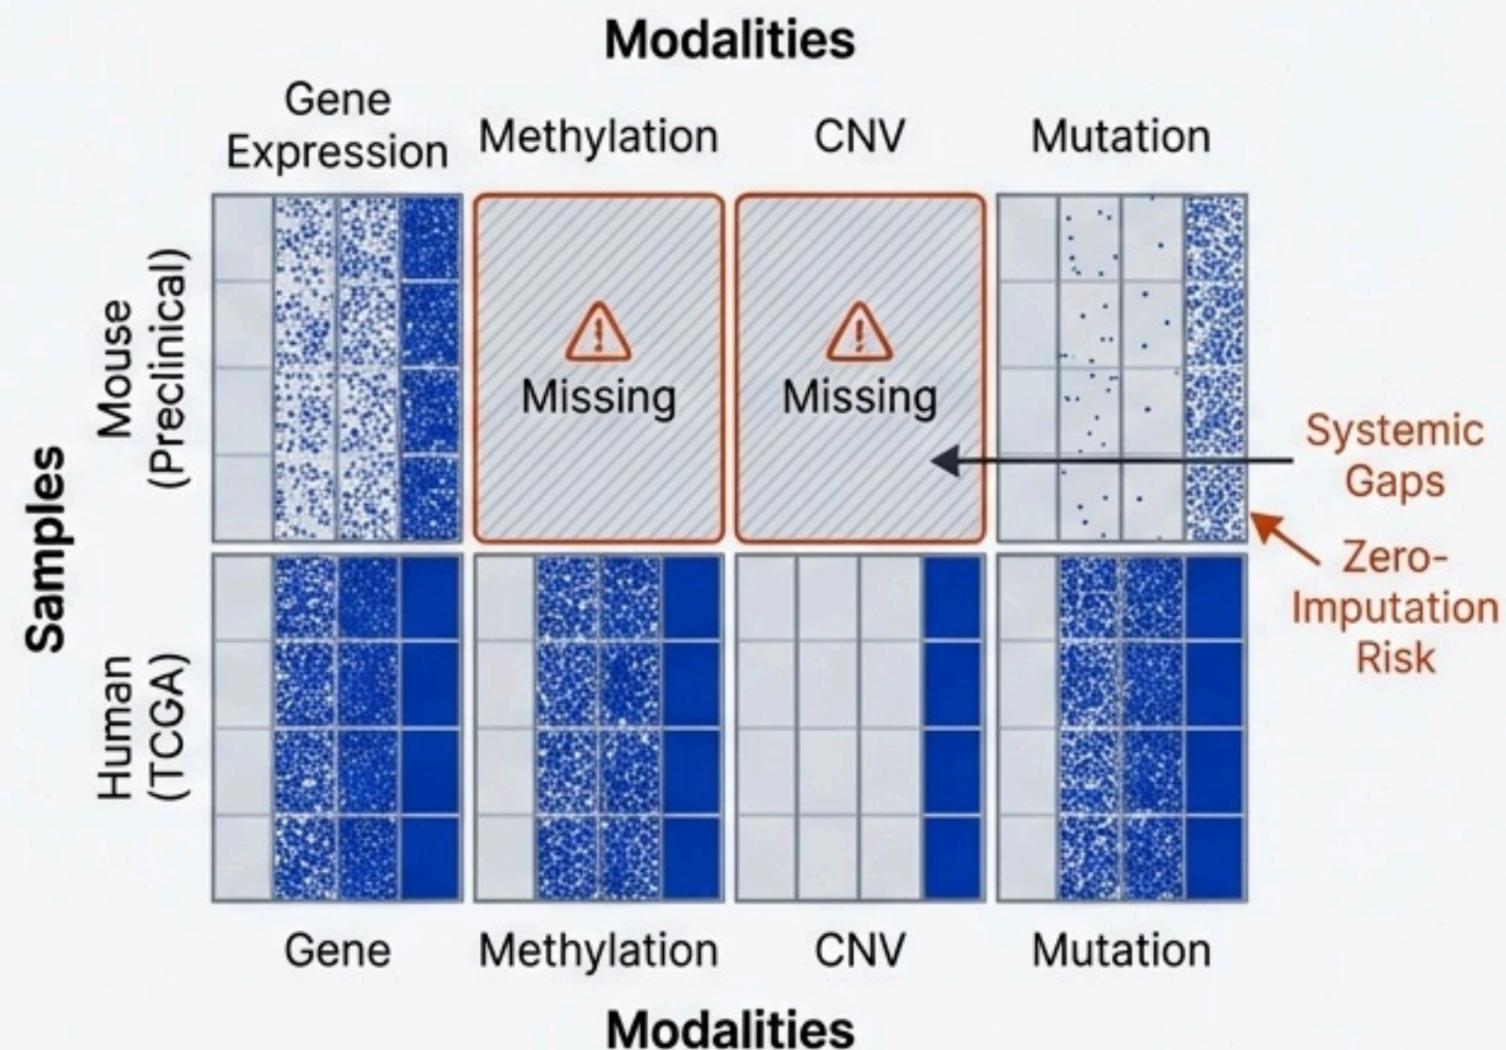


A unified pipeline converting noisy, single-modality preclinical inputs into validated, uncertainty-calibrated human trajectory predictions.

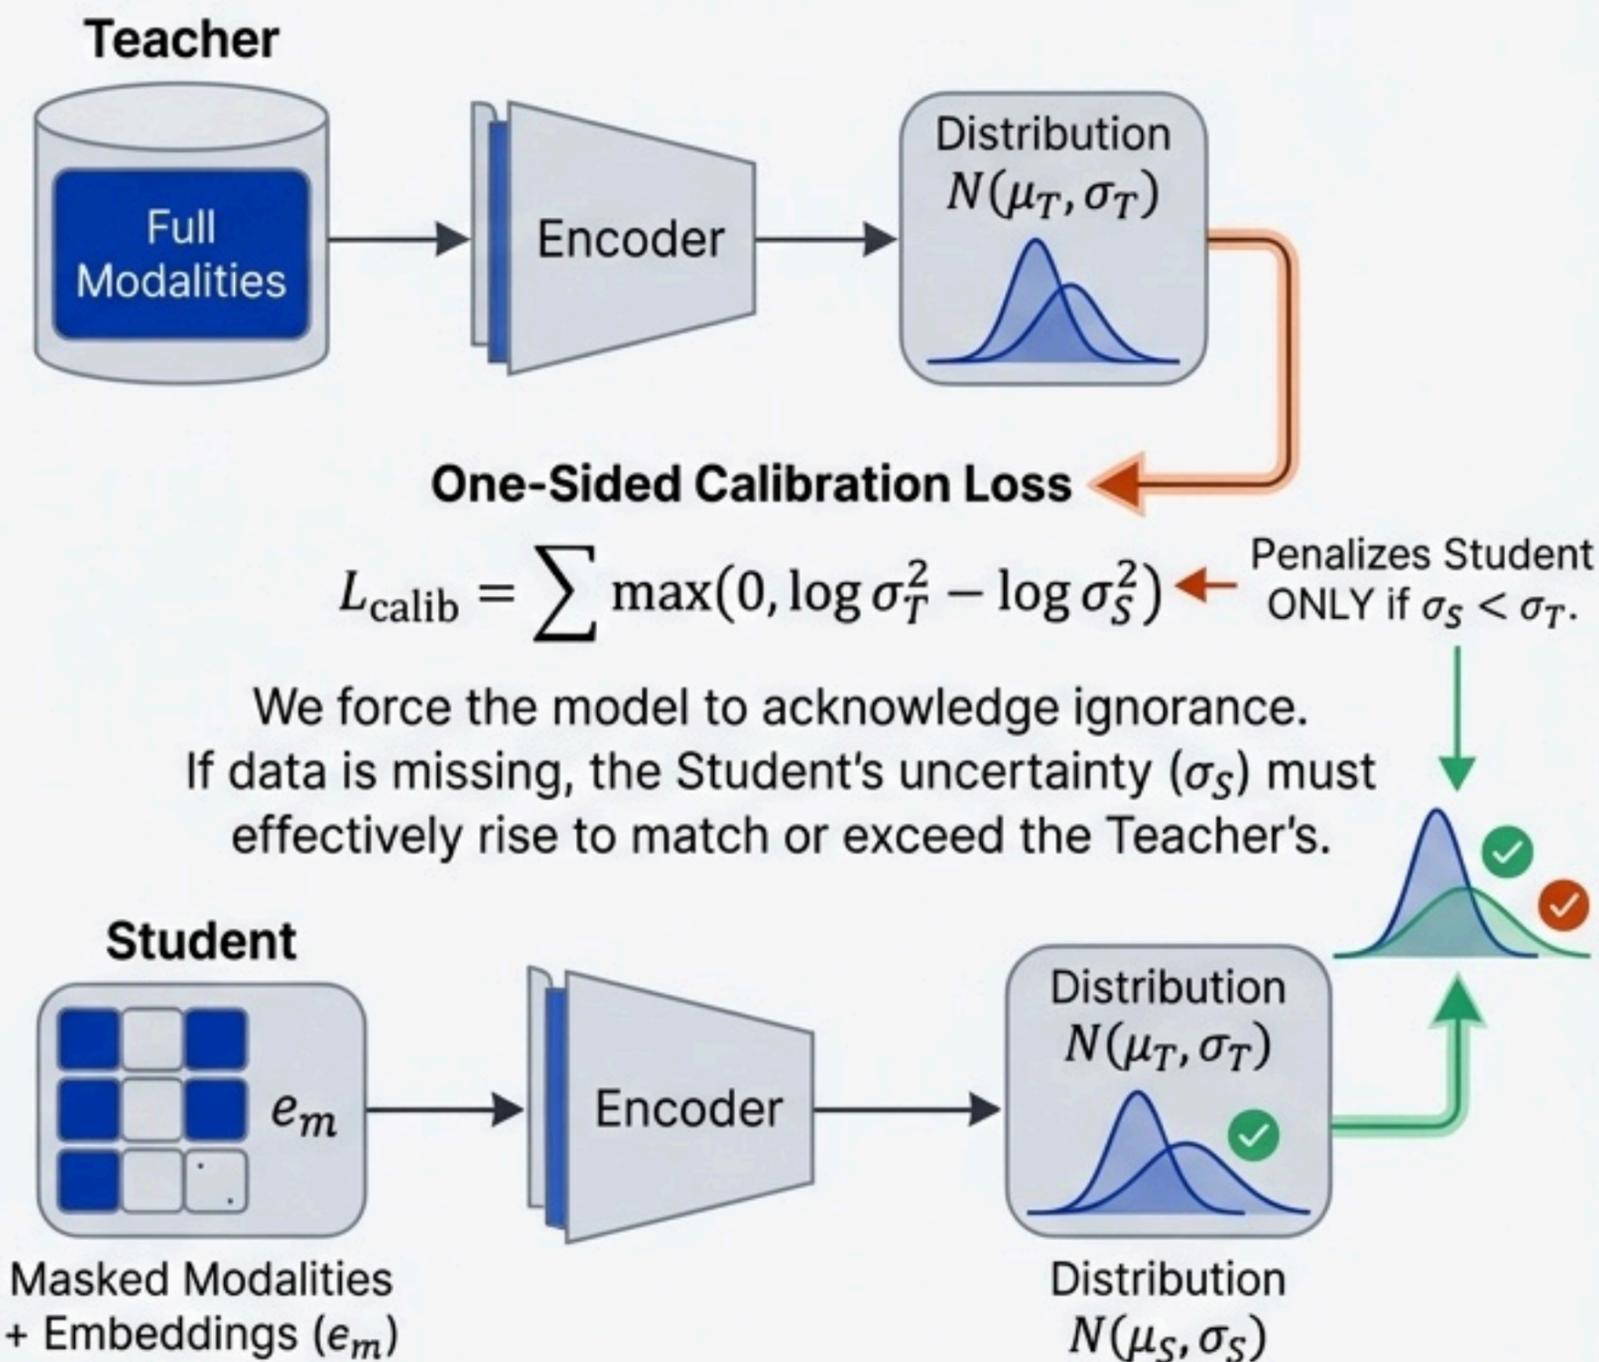
Solving Data Gaps: Probabilistic Encoder Self-Distillation (PESD)

The Challenge

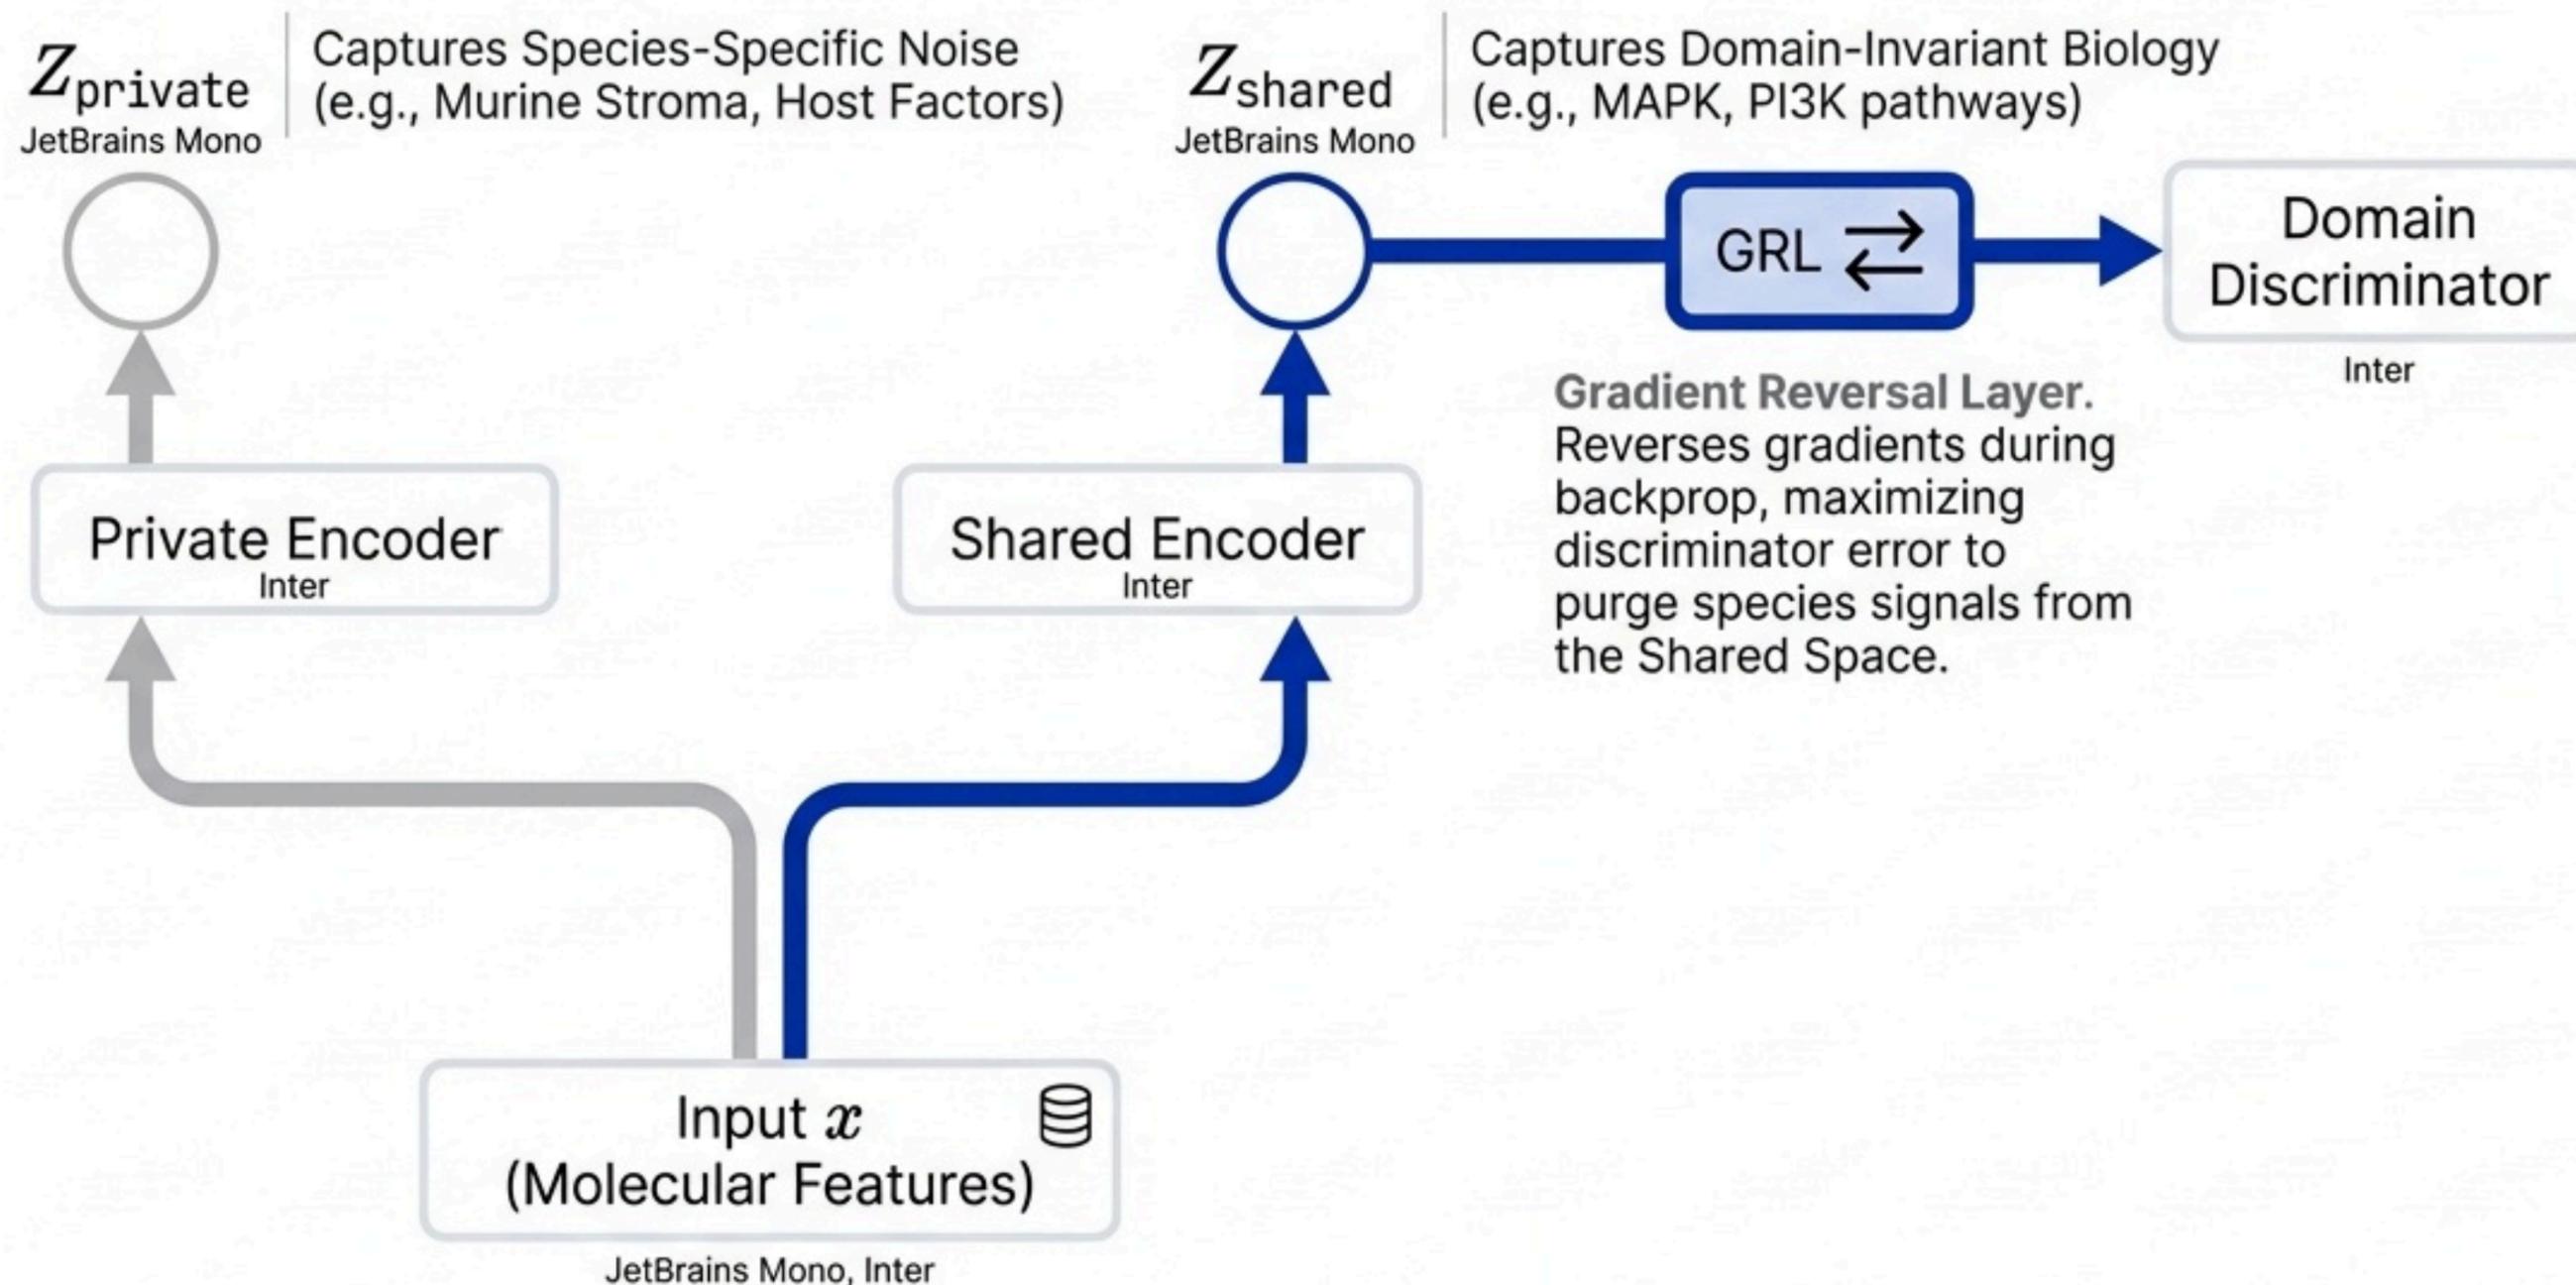
Preclinical datasets often lack modalities (e.g., Methylation, CNV) available in TCGA human datasets. Zero-imputation leads to overconfident errors.



The Solution: One-Sided Calibration



Disentangling Species from Biology: The Domain Separation Network



The Failure Mode: Metabolic Scaling-Induced Representational Collapse

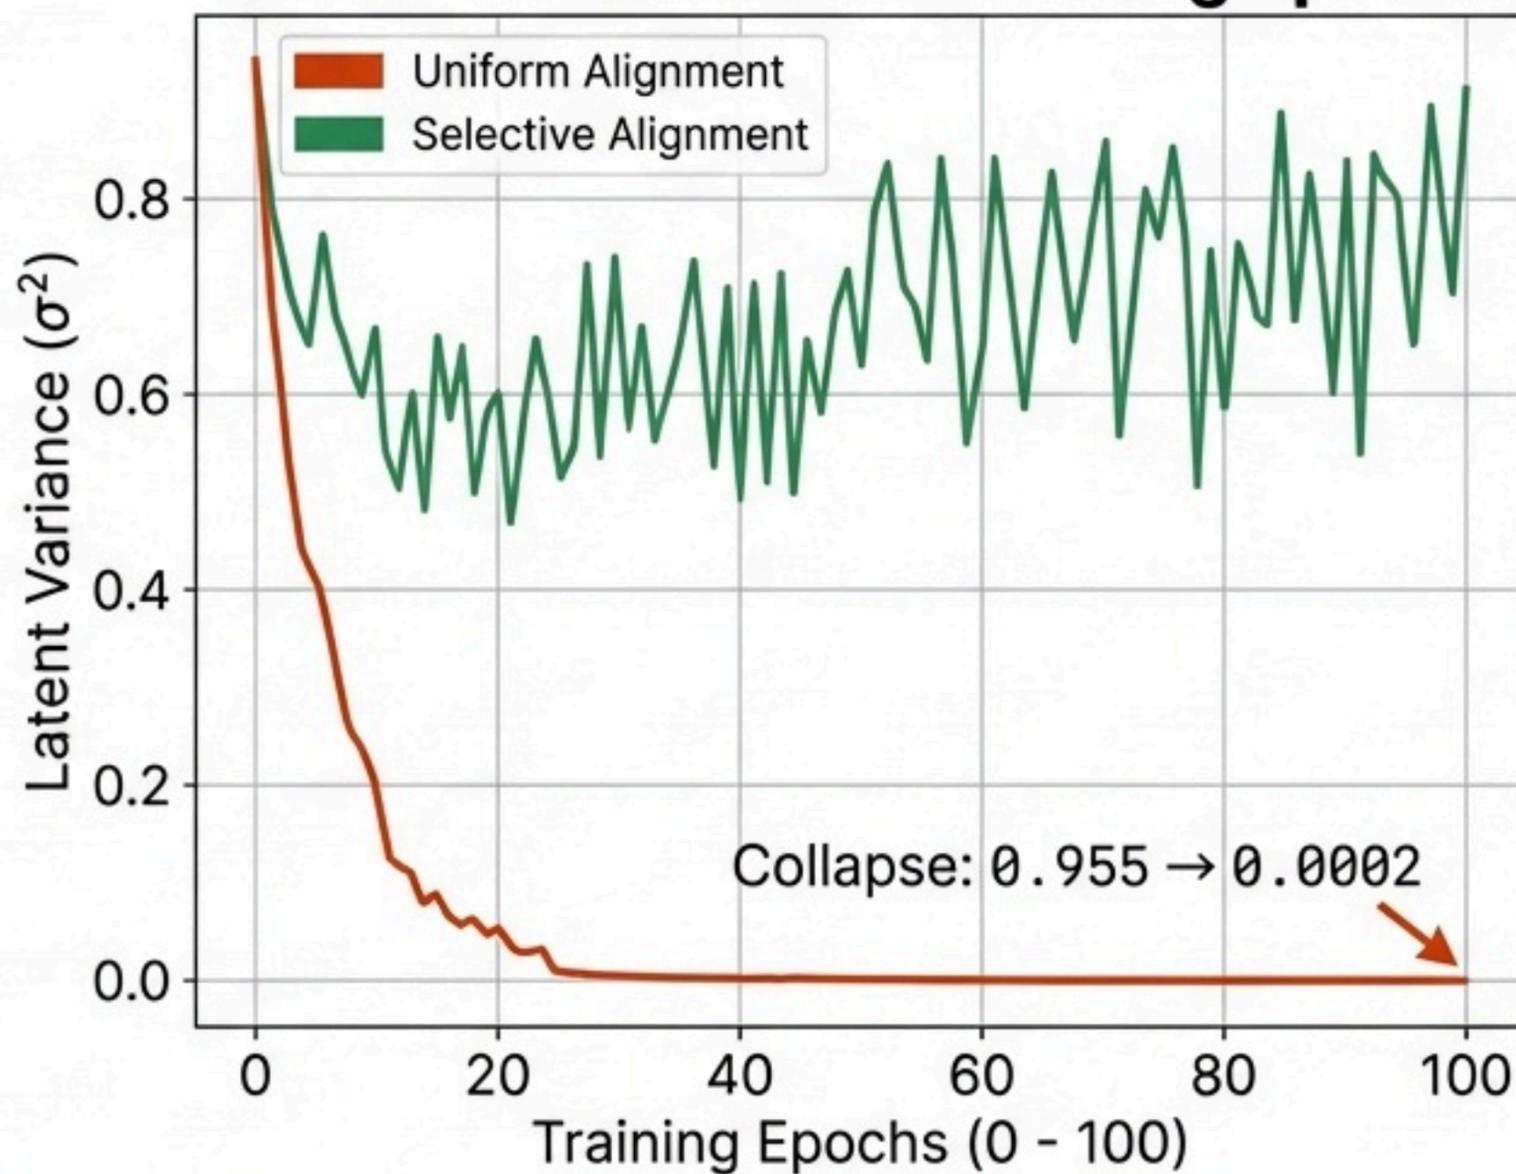
The Conflict

Standard domain adaptation forces all features to align statistically. However, proliferation signals scale allometrically ($r \propto M^{-1/4}$). Forcing alignment on these biology-physics conflicts causes the encoder to compress the latent space to zero variance to “hide” the difference.

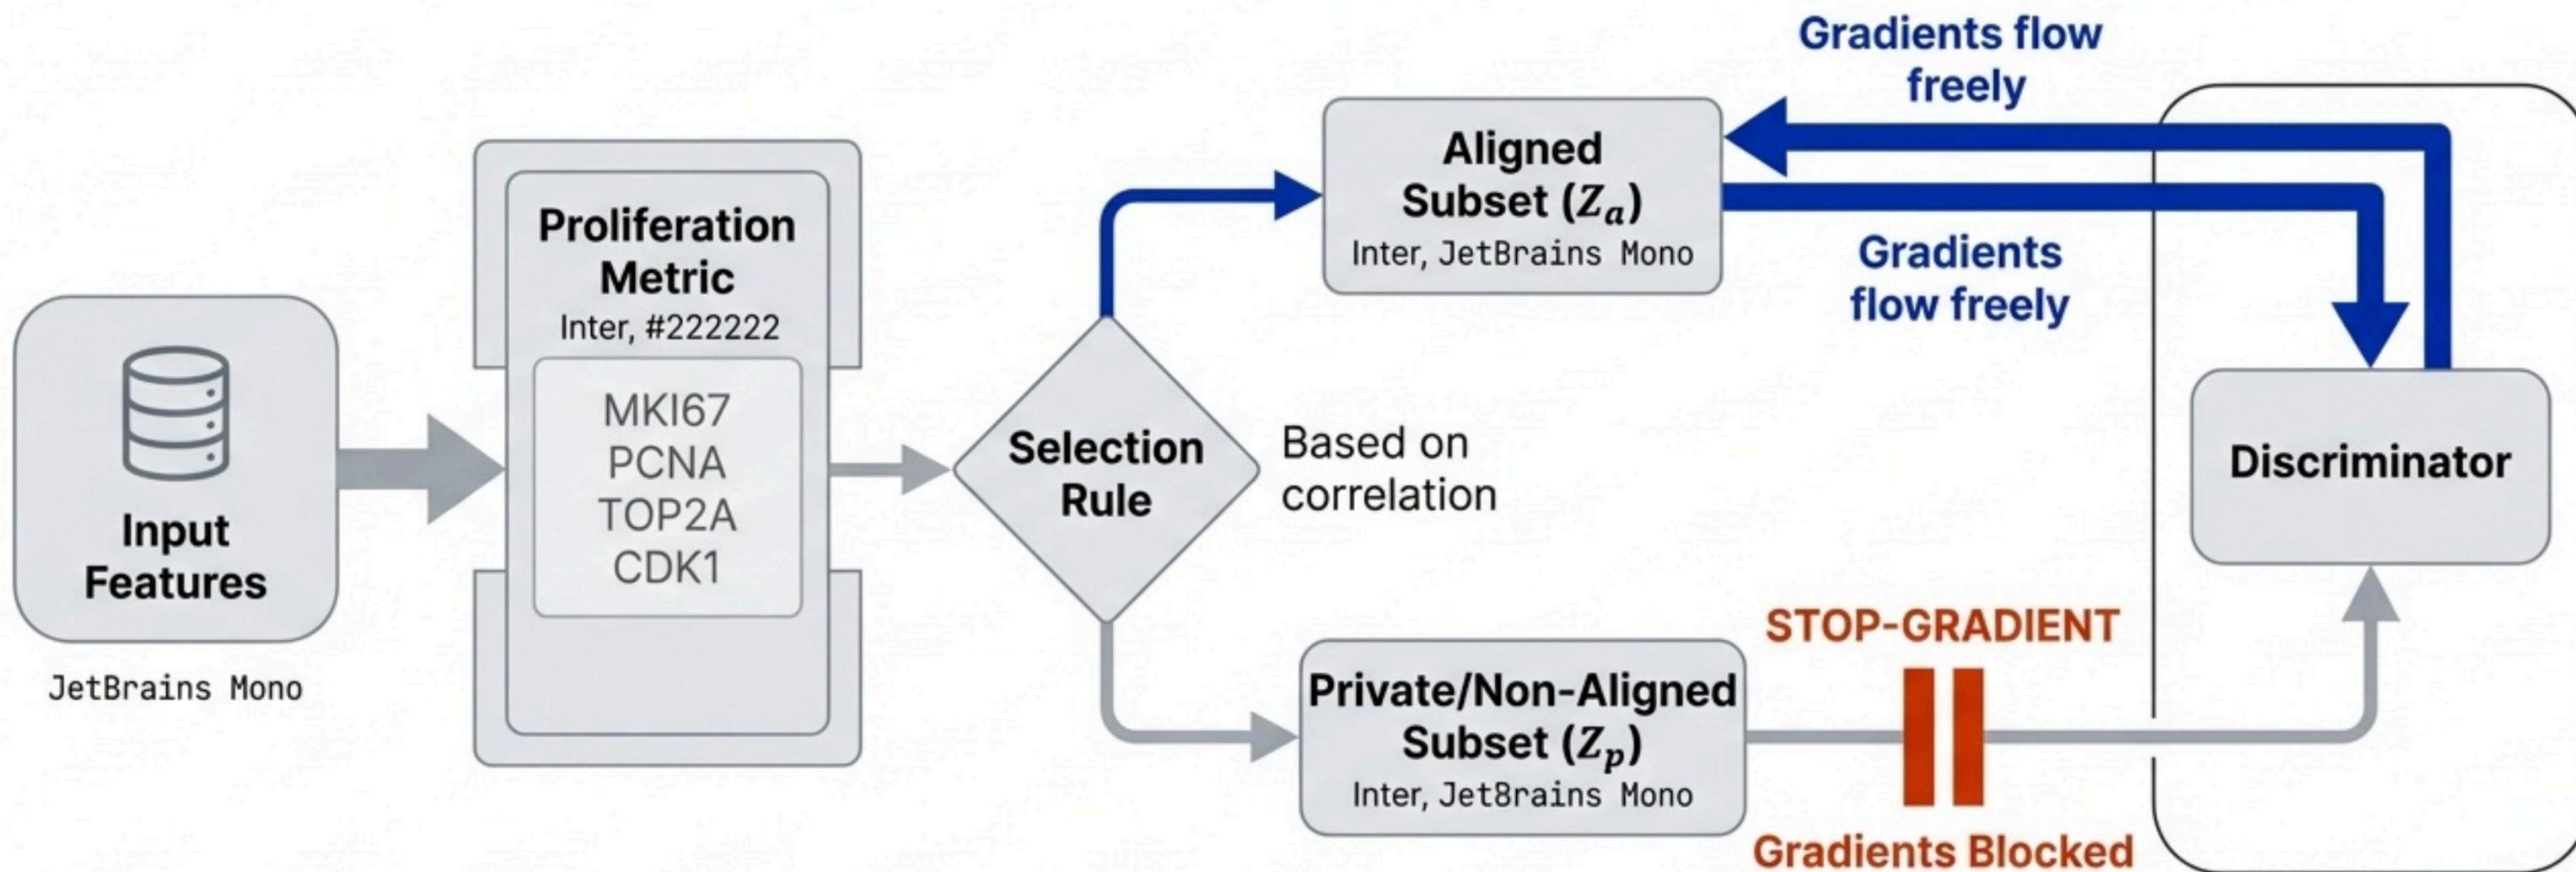
The Consequence

Latent Information Capacity drops to near-zero. Downstream ODE solvers fail due to lack of signal.

Latent Variance over Training Epochs



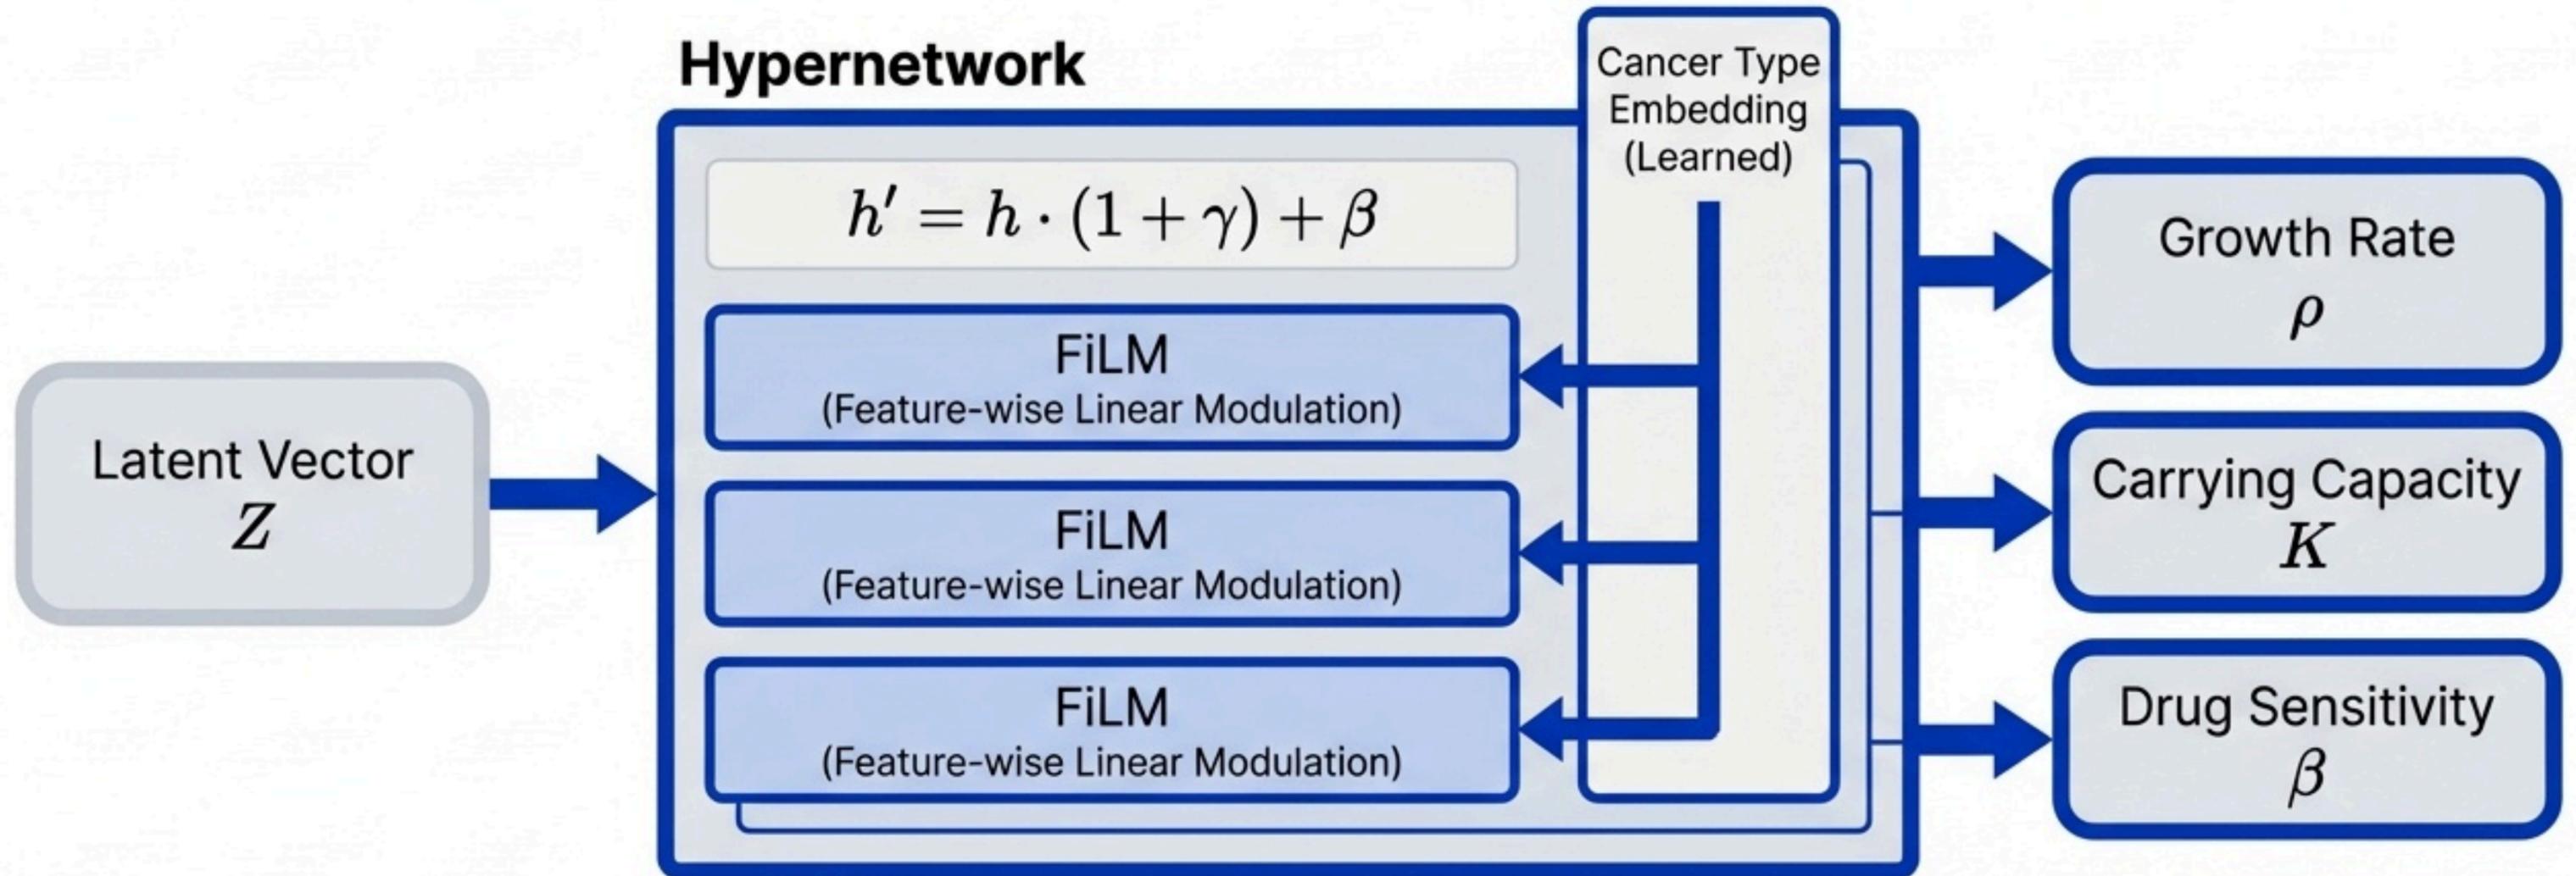
Proliferation-Aware Selective Alignment



Adversarial updates are physically blocked from cell-cycle genes. We preserve the biological fidelity of growth signals (Z_p) while aligning the remaining pathway biology (Z_a).

Cancer-Aware Parameter Generation

Transforming Latents into ODE Parameters via FiLM



Identity Initialization

Parameters γ and β are initialized to zero. The model begins as a universal identity transform and learns specific histological deviations (e.g., Lung vs. Breast kinetics) during training.

Physics-Informed Neural ODEs: Separated-State Dynamics

Traditional Approach

$$\frac{dN_k}{dt} = \dots$$

Tracks absolute populations directly.
Suffers from identifiability issues
(Scale vs. Composition confounded).

Separated-State Formulation (Patent 3)

Decomposes state into **Size** (Burden) and **Composition** (Fractions).

$$\frac{dB}{dt} = \bar{\phi}(t) \cdot B(t)$$

Driven by Mean
Fitness $\bar{\phi}$

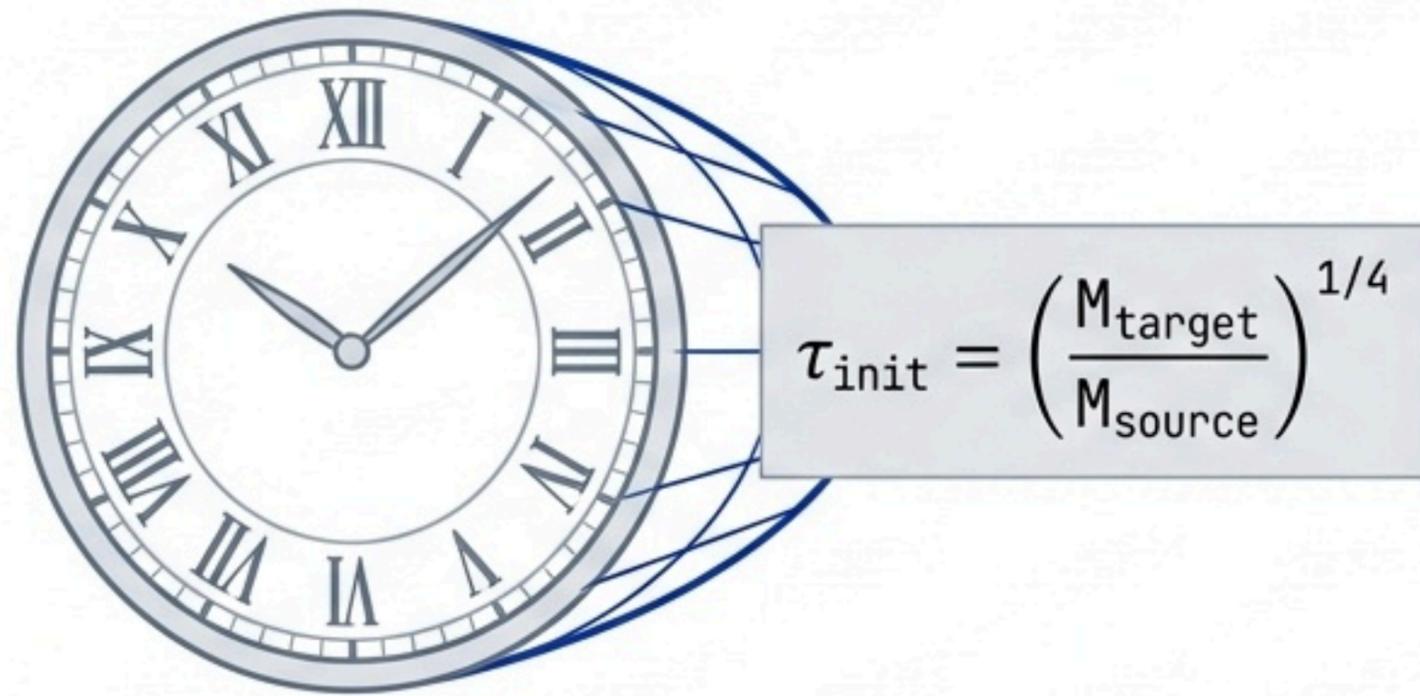
$$\frac{df_k}{dt} = f_k(t)(\phi_k(t) - \bar{\phi}(t))$$

Guarantees $\sum f_k = 1$
and $f_k \in [0,1]$ by
construction.

Removes need for external projection penalties,
eliminating solver stiffness.

Correcting Time & Immune Interaction

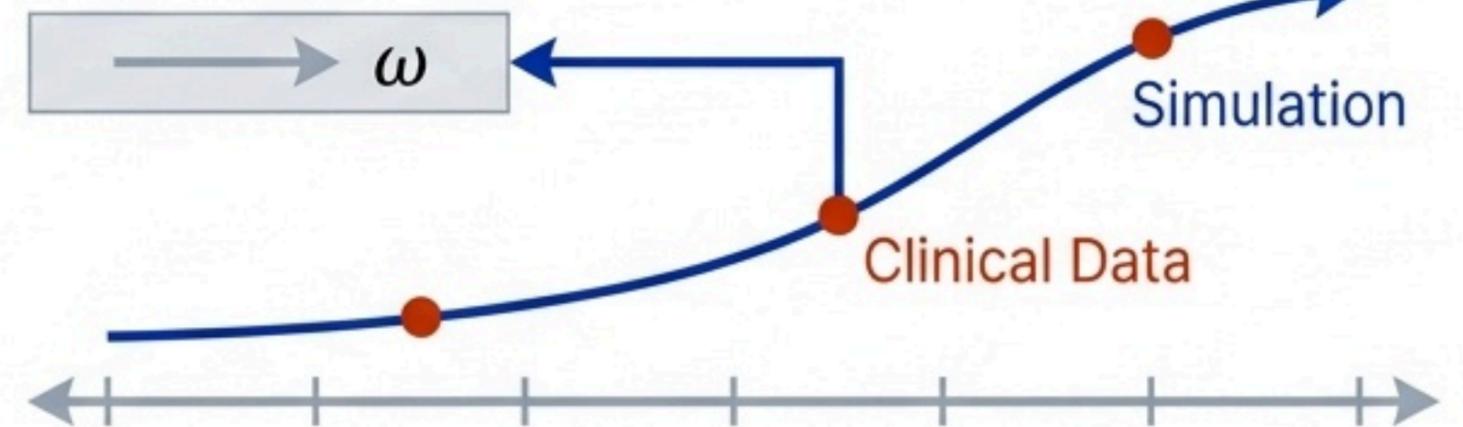
Allometric Time Dilation



The simulation clock is “dilated” by parameter τ to match human metabolic rates. This adapts the fast murine growth physics to the slower human timescale.

Inverse-Solving Immunity

Model Parameters

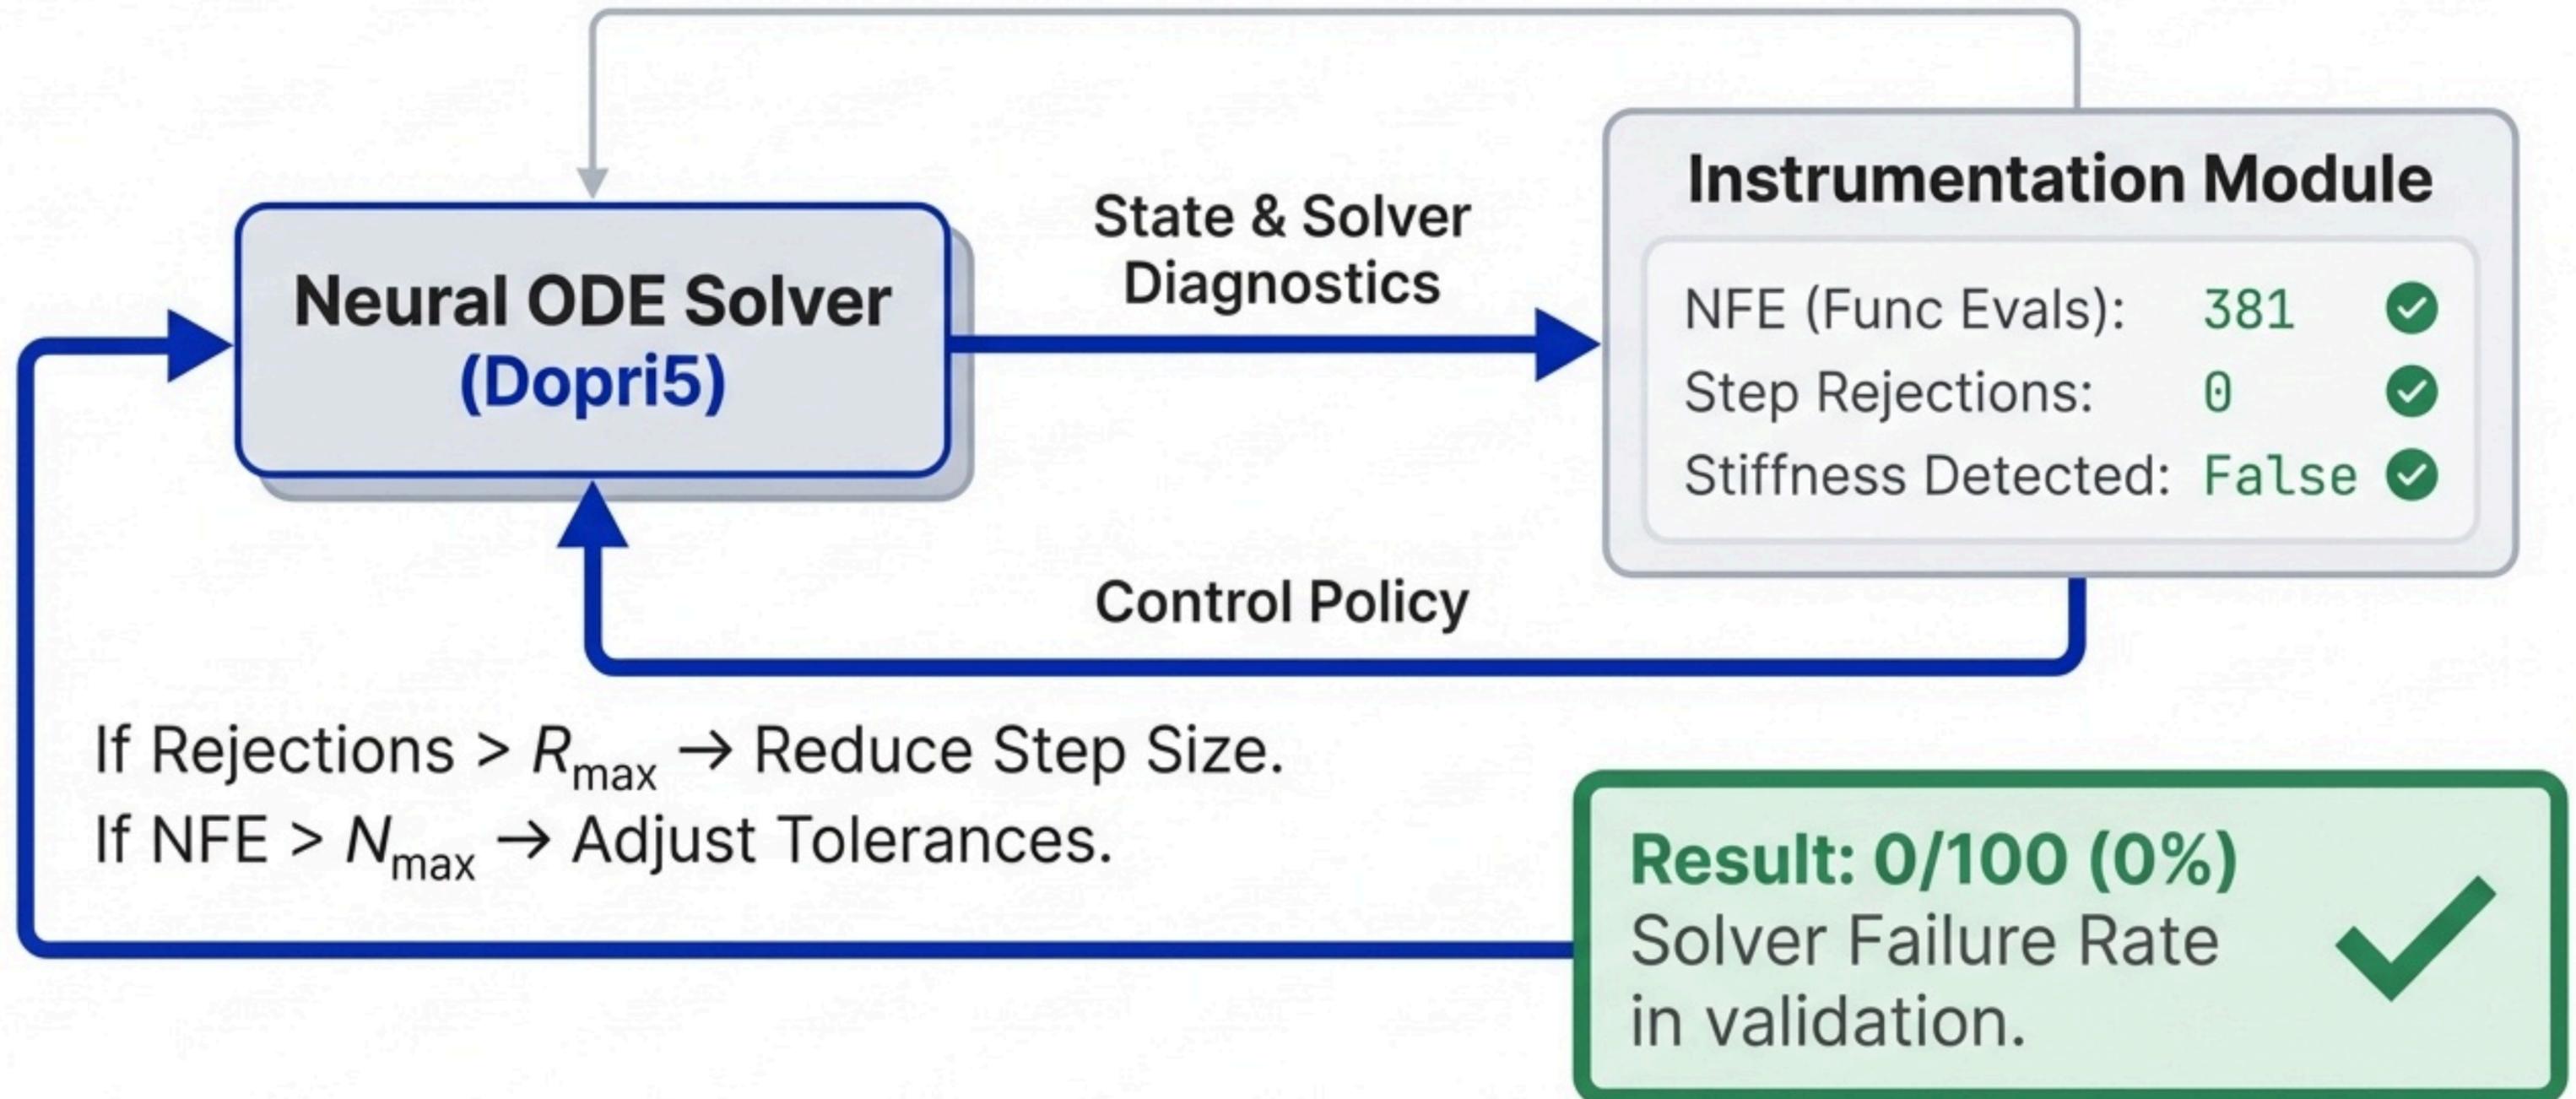


Inverse Problem: Solving for ω (Immune Clearance).

$\omega \in (0, \infty)$ via softplus constraint.

Since human data is too sparse to train a full immune model, we mathematically invert the ODE to find the single clearance parameter that explains the deviation from “pure” growth.

Ensuring Solver Stability: Instrumentation & Feedback



Three-Layer Uncertainty Quantification

Oscillation Damping

Post-processing "Severity-Graded Damping" (Lowpass/EMA) to remove numerical artifacts.

Temporal Trust

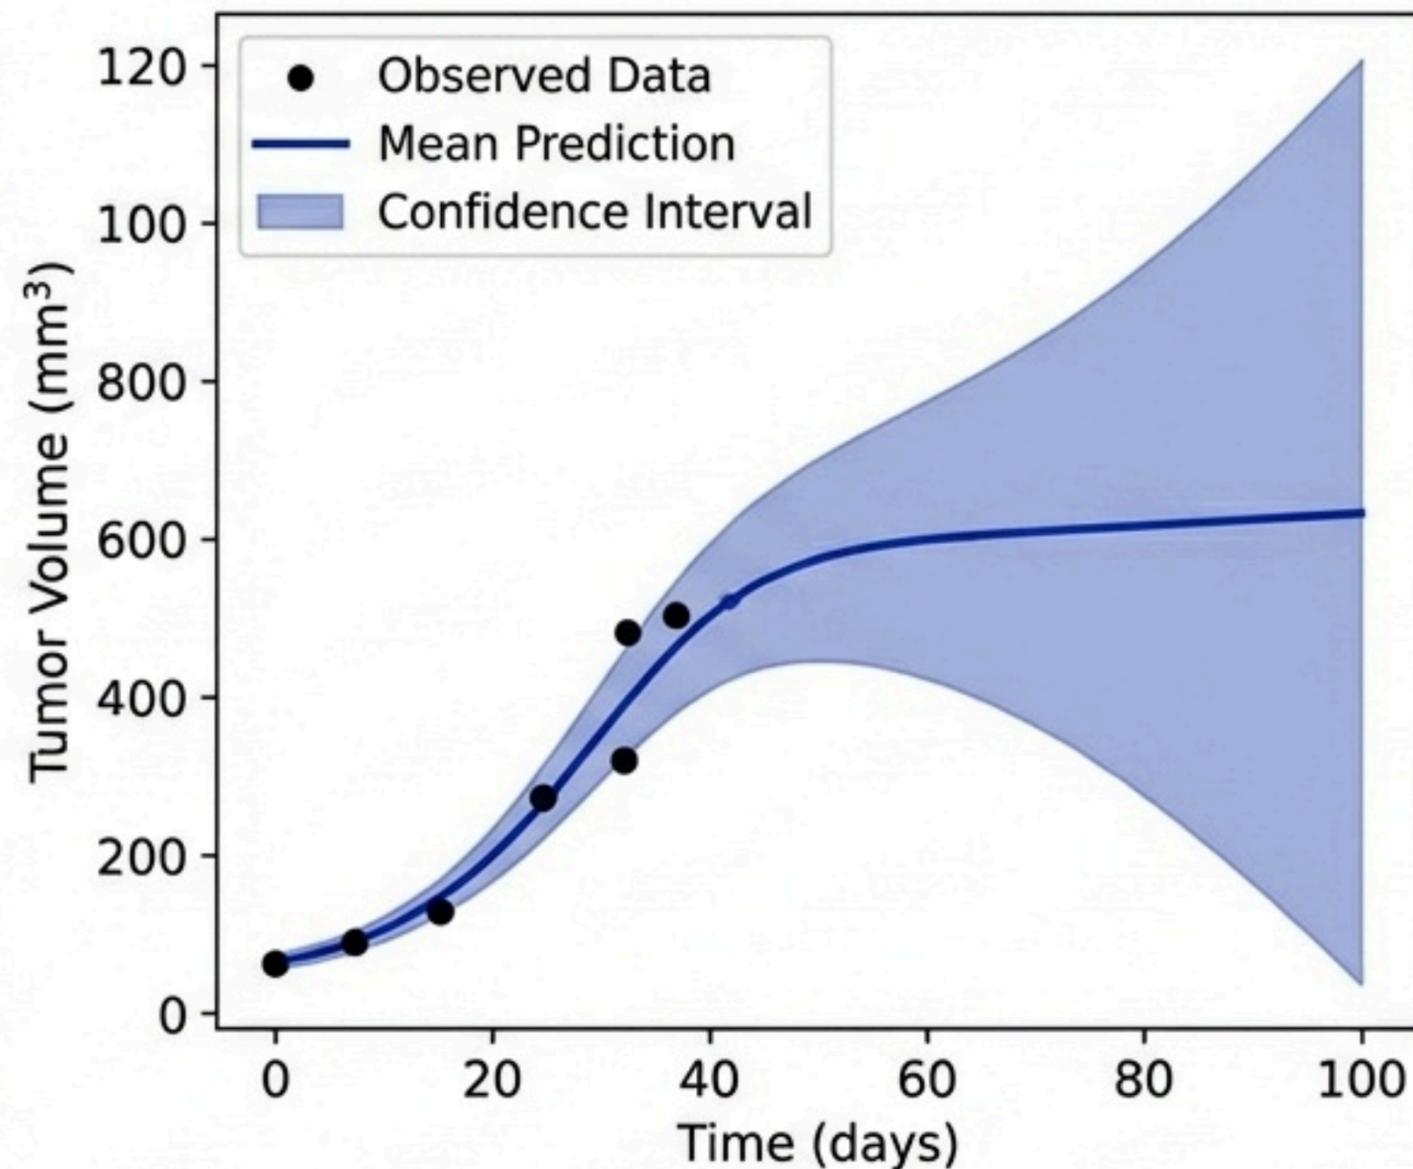
Trust decays as distance from data increases.

$$\tau(t) = \exp(-\lambda \cdot \min |t - t_{\text{obs}}|)$$

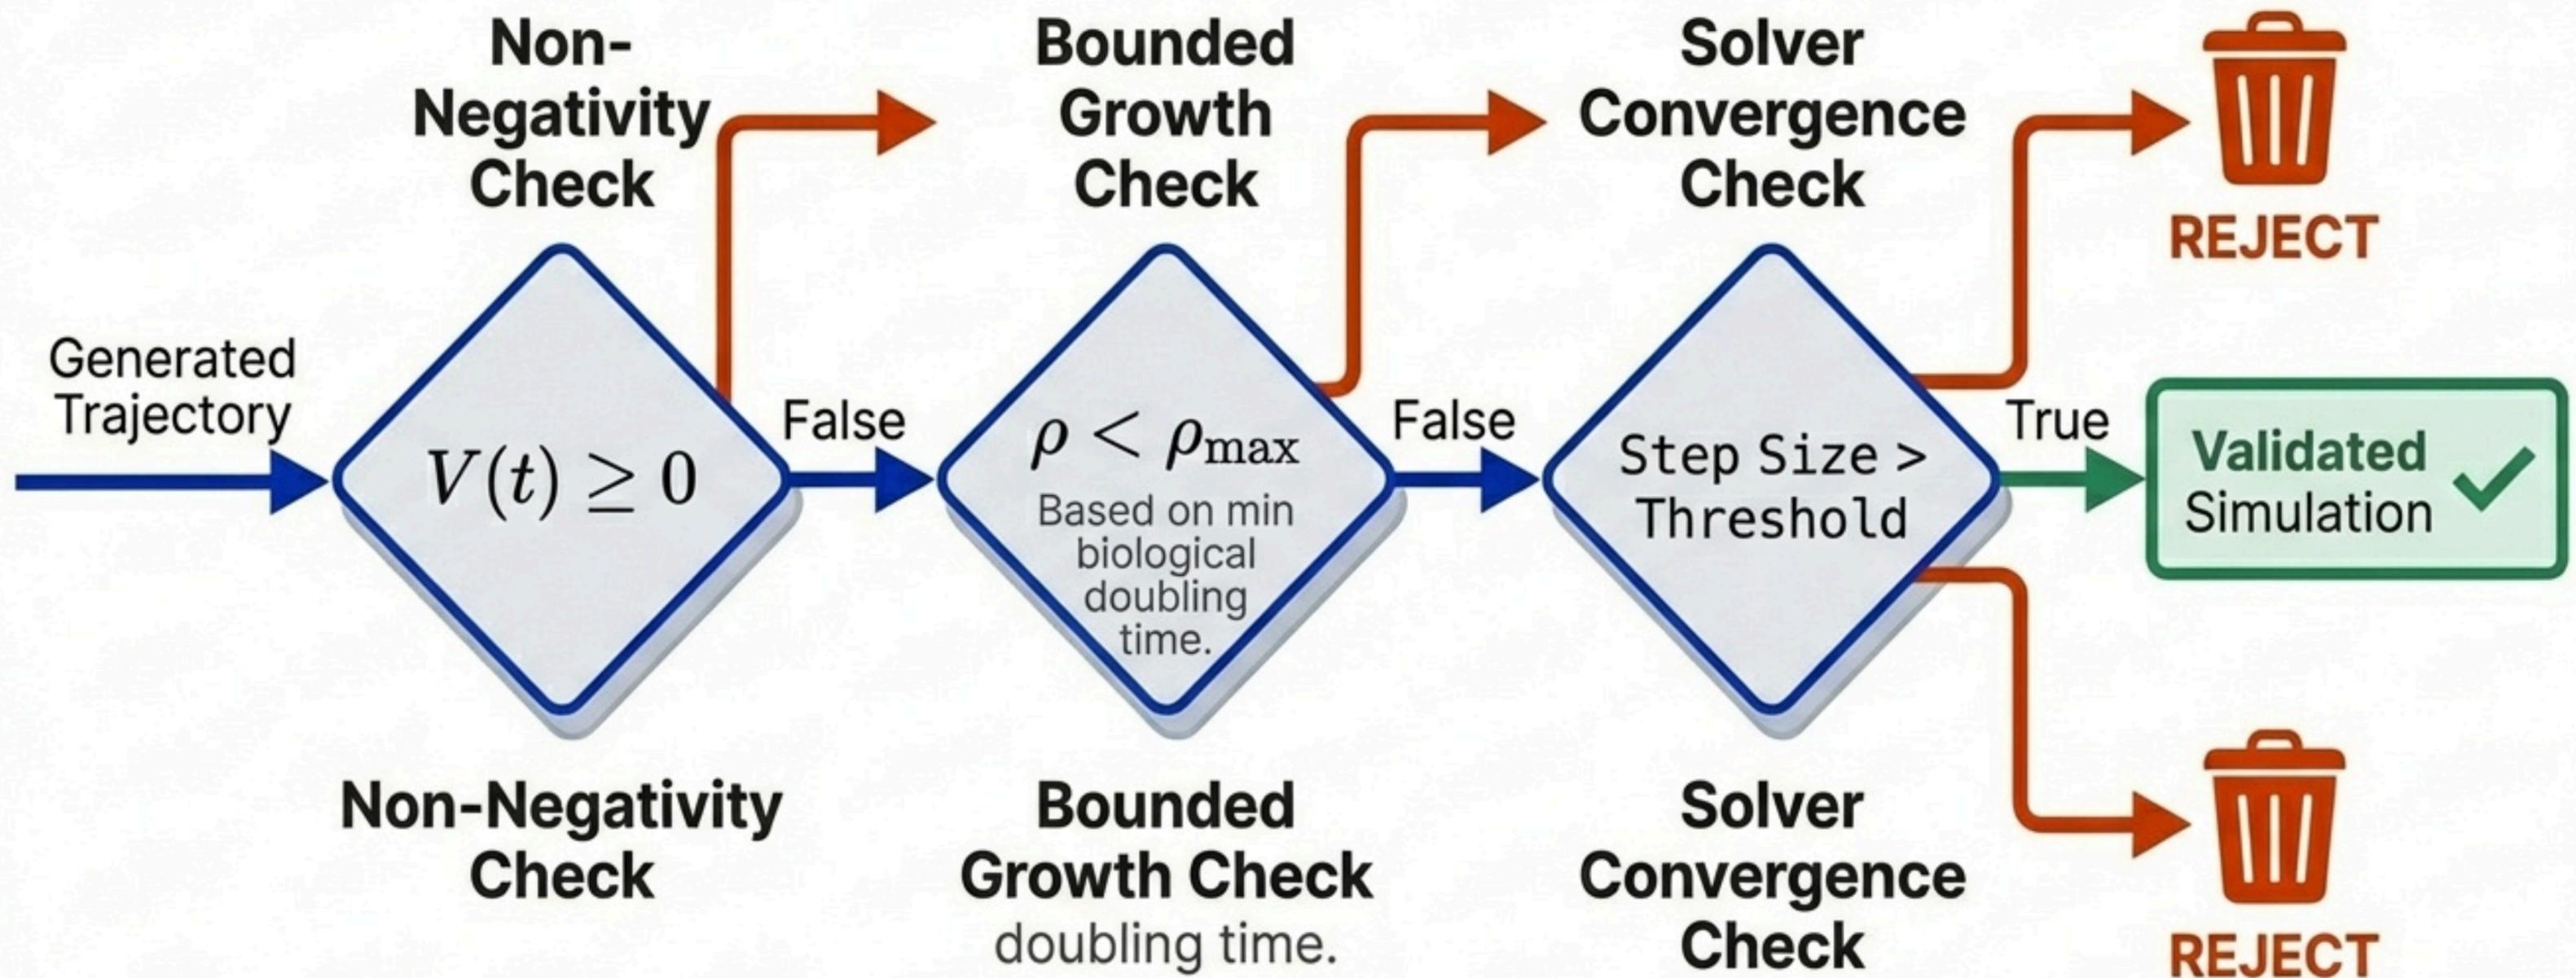
Parameter Ensembles

Blending model predictions with Population Priors. N ensemble members with random perturbations.

Tumor Trajectory Plot (Output)



Deterministic Validation Logic: The Rejector



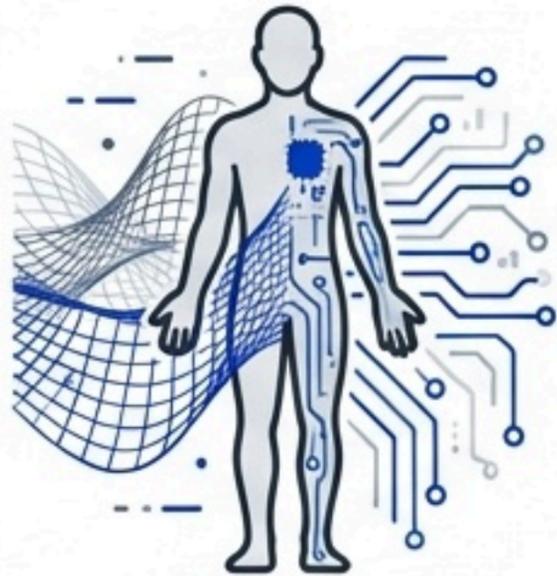
Physics constraints are hard boundaries. We do not output “hallucinated” physics.

Technical Innovation Matrix

TRANSLATIONAL CHALLENGE	TECHNICAL INNOVATION	PATENT SOURCE
Data Sparsity (Missing Modalities)	PESD with One-Sided Calibration	Patent 3
Stroma Replacement (Negative Transfer)	Domain Separation Network (DSN)	Patent 1
Metabolic Scaling (Representational Collapse)	Proliferation-Aware Selective Alignment	Patent 2
Solver Instability (Stiffness)	Separated-State Dynamics & Instrumentation	Patent 2 & 3
Uncertainty Quantification	Ensembles + Trust Decay + Damping	Patent 3

The Vision: Rigorous In Silico Clinical Trials

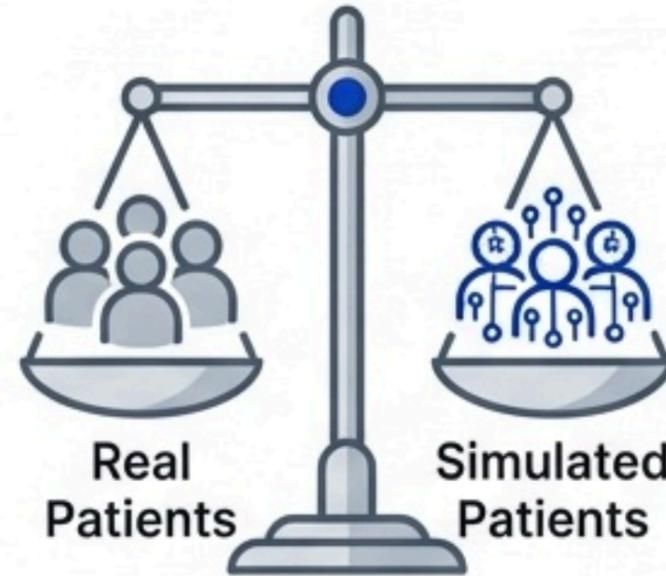
Digital Twin



Physics-Constrained
Patient Simulation.

Inter

Hybrid Control Arms



Augmenting Clinical Trials
with Synthetic Data.

Inter

Decision Support



De-risking Decisions with
Calibrated Uncertainty.

Inter

Mastery over design. Fidelity to physics. Trust in prediction.