CancerDyn

Satya Thavanesh Yalla

1 · PROBLEM

what and why

Most machine-learning approaches to cancer treatment response treat the problem as a static prediction task: given a tumor's molecular profile and a drug, predict whether the tumor will respond. However, tumors are dynamic systems whose molecular states change over time and in response to treatment. This project proposes CancerDyn, a framework that models treatment response as a continuous dynamical process rather than a single static prediction. The central hypothesis is that biological interaction structure should not be treated as equally influen- tial in every tumor or treatment context. CancerDyn therefore begins with a fixed, biologically grounded interaction graph and learns a context-dependent effective interaction structure whose edge weights vary with the tumor's molecular state and treatment. This structure drives a treatment-conditioned dynamical model that predicts how the tumor state evolves over time. The framework will additionally use an interpretable surrogate model to approximate the computationally expensive dynamics, allowing large treatment spaces to be searched efficiently.

2 · APPROACH

model / method, and why this one

CancerDyn consists of four primary components. 1. Biological state encoder. Multi-omic measurements such as gene expression, mutations, copy-number alterations, and other available molecular features will be compressed into a patient/tumor state representa- tion. Multiple encoders will be evaluated, including a conventional MLP representation and a graph-based representation. 2. Context-dependent effective interaction graph. A fixed biological interaction graph will provide the structural prior. For each known interaction $i \rightarrow j$, CancerDyn learns: ( ) $w_{ij}(t, u) = g_\theta(h_i(t), h_j(t), e_{ij}, u(t))$ where h_i, h_j are the current molecular states, e_{ij} represents the known biological interaction, and $u(t)$ rep- resents treatment. The model therefore does not claim that treatment creates or destroys physical molecular interactions. Instead, it estimates how strongly an established interaction contributes to the modeled tumor dynamics in a particular context. Edges absent from the biological prior cannot be created by the model, providing a biological constraint on the learned structure. 3. Treatment-conditioned tumor dynamics. The resulting effective graph is used to evolve the tumor state: $dz/dt = f_\theta(z(t), G(z, u), u(t))$ The primary implementation will use a Neural ODE, allowing the model to represent continuous tumor- state trajectories under different treatment conditions. However, the framework will not assume that Neural ODEs are inherently superior. A GRU-based discrete-time dynamical model and static models will be evaluated as predefined alternatives. This allows us to determine whether continuous-time dynamics actually provide value rather than simply assuming that they do. 4. Surrogate treatment optimizer. Neural ODE simulation can be computationally expensive when evaluating many treatment schedules. CancerDyn will therefore generate simulated treatment trajectories and train an XGBoost surrogate to approximate the expensive dynamical model. The surrogate can then rapidly evaluate large numbers of candidate drugs, doses, treatment durations, and treatment sequences. The highest-performing candidates will subsequently be re-evaluated using the original dynamical model rather than trusting the surrogate blindly. Interpretability. SHAP will be incorporated at two levels. Graph-level attribution identifies which molecular features and treatment variables contribute to context-dependent changes in effective interaction weights. Prediction-level attribution identifies which features contribute to the final predicted treatment response. These explanations will be treated as model-level attribution rather than evidence of biological causality. The thing to emphasize most in this proposal is not "we use a Neural ODE." It is: CancerDyn treats treatment response as movement through a context-dependent biological dy- namical system. A fixed biological interaction prior constrains the model's interaction space, while tumor state and

treatment determine the effective influence of those interactions over time. An expensive dynamical simulator is subsequently distilled into a fast surrogate for treatment-space exploration. That is the part that makes this a framework rather than a model-comparison project. It is also deliberately distinct from the companion BioTopo work: BioTopo → characterize biological topol- CancerDyn → model how tumor state ogy. evolves through that topology under intervention. That distinction should stay absolutely central as the project is developed.

3 · DATASET

source, size, preprocessing

The project will use publicly available cancer datasets containing molecular profiles and perturbational treatment-response information. The initial dataset audit will prioritize resources containing combinations of gene expression, mutations, copy-number alterations, cancer-cell identity, drug perturbations, drug sen- sitivity/response, and time-dependent or repeated measurements where available. Potential sources include DepMap, GDSC, TCGA/CPTAC, and compatible perturbational datasets. The exact dataset combination will be finalized only after verifying that the available measurements contain sufficient temporal or perturbational information to justify a dynamical model. Preprocessing will include gene identifier harmonization, molecular feature normalization, missing-value handling, drug identity/dose normalization, construction of the biological interaction prior, alignment of molecular and treatment-response measurements, and removal of duplicated or incompatible measure- ments. A major requirement will be leakage-resistant splitting. Evaluation will include held-out cancer models, treatments, and/or cancer types where the available data permit.

4 · METRICS

how you'll know it worked

CancerDyn will be evaluated at several levels rather than using a single prediction metric. A. Static prediction. Compare XGBoost, MLP, and GNN using RMSE, MAE, R2 , and Spearman corre- lation. B. Dynamic prediction. Compare a static model, GNN + GRU, and GNN + Neural ODE using trajectory MAE, trajectory RMSE, endpoint prediction error, and correlation of predicted vs. observed trajectories. The primary question is: does explicitly modeling tumor dynamics improve prediction of future treatment response? C. Context-dependent graph. Compare Gfixed against G(z, u) to determine whether treatment- and state- dependent effective interactions improve held-out prediction. Additional analysis will measure stabil- ity of learned edge weights, sparsity, reproducibility across seeds, and enrichment of important learned edges/pathways against independent biological evidence. D. Surrogate performance. The XGBoost surrogate will be compared directly with the Neural ODE using MAE, RMSE, R2 , treatment-ranking correlation, and top-k treatment agreement. The surrogate's computational advantage will also be measured: Neural ODE inference time Speedup = Surrogate inference time E. Treatment optimization. The final experiment will compare random search vs. surrogate-guided search vs. direct dynamical-model search where computationally feasible. The primary outcome will be whether the surrogate can identify high-performing treatment strategies while substantially reducing computational cost.

5 · COMPUTE

training time, memory, dependencies

The project is intentionally designed to be substantially more computationally demanding than a conven- tional tabular cancer ML project. Expected software: Python 3.10+, PyTorch, PyTorch Geometric, SciPy, XGBoost, scikit-learn, SHAP, pandas / NumPy, Matplotlib, and experiment tracking through MLflow or an equivalent lightweight system. The most computationally expensive components will be graph representation learning, Neural ODE train- ing, repeated ODE trajectory simulation, multiple random seeds, architecture ablations, treatment-scenario generation, and surrogate training. The project will ideally use GPU compute for the main experiments, while preprocessing, XGBoost, and many ablation analyses can remain CPU-based. Every experiment will record random seed, dataset version, model configuration, hyperparameters, train- ing time, inference time, hardware, and evaluation metrics, following the proposal

format’s emphasis on recording reproducible configurations and computational requirements.

6 · DELIVERABLES

model, paper, demo

1. CancerDyn framework. A reproducible implementation containing biological graph construction, context-dependent graph transformation, tumor-state encoder, treatment-conditioned dynamics, Neural ODE implementation, alternative GRU dynamics, XGBoost surrogate, treatment-search module, and SHAP explainability layer. 2. Experimental model suite. A complete ablation hierarchy establishing which components actually contribute to the framework: XGBoost ↓ MLP ↓ GNN ↓ GNN + GRU ↓ GNN + Neural ODE ↓ GNN + Context-Dependent Neural ODE ↓ CancerDyn + XGBoost surrogate 3. Context-aware interaction analysis. Visualizations showing the baseline biological graph, the context- dependent effective graph, treatment-specific changes, important molecular drivers, and SHAP explana- tions. 4. Treatment optimization demonstration. A computational demonstration in which CancerDyn searches a large treatment space and verifies the highest-ranked candidates using the original dynamical model. 5. Research paper. The final paper will document framework design, biological motivation, datasets, model architecture, ablation experiments, dynamic prediction results, context-dependent graph analysis, surrogate validation, treatment optimization, explainability, and failure cases and limitations.

7 · RISK

what breaks first, and plan B

what breaks first, and plan B Risk Plan B Neural ODE does not outperform Determine whether the available data contain enough static models temporal information; retain the best validated archi- tecture rather than forcing the ODE GRU outperforms Neural ODE Retain GRU as the dynamical component and inves- tigate whether continuous-time modeling provides value only under irregular sampling GNN does not outperform MLP Test alternative biologically grounded graphs; if graph structure remains unhelpful, remove the graph compo- nent Context-dependent graph does not Compare fixed vs. context-aware structures and deter- improve prediction mine whether context-dependent weighting provides measurable value Learned graph becomes unstable Increase sparsity/regularization and constrain the model more strongly to the biological prior Surrogate poorly approximates Compare XGBoost with MLP/Random Forest; if ap- Neural ODE proximation remains inadequate, omit surrogate opti- mization Treatment optimization produces Use the Neural ODE as the final verification model unstable recommendations and evaluate robustness across perturbations Longitudinal data are insufficient Reframe the dynamical component around perturba- tional response rather than claiming continuous tumor evolution Results do not outperform existing Analyze which components fail and report the condi- methods tions under which the framework succeeds rather than tuning until it wins This failure-first design is deliberate. Negative results are treated as scientifically useful rather than as- suming the proposed method must win.