Title: An Immune-Oncology Multi-Modal Foundation Model for Adaptive Oncology Drug Development

Abstract:
Cancer drug development increasingly fails not because targets are unknown, but because patient-level biology is heterogeneous and dynamic, therapeutic benefit emerges from evolving interactions among tumor cells, the immune system, and the host microenvironment. The complexity of cancer and therapeutic response necessitates a holistic approach that moves beyond single-modality biomarkers. Traditional methods struggle with the fragmented, incomplete, and heterogeneous nature of real-world biomedical data, limiting generalizability and clinical actionability.

A novel, immune-centric omnimodal foundation model is described to unify diverse data layers—including genomics, transcriptomics, pathology imaging (H&E and multiplexed/spatial imaging), proteomics/immune profiling, and clinical metadata, into a single, co-trained AI framework that captures system-level immune–tumor dynamics.

At the core of the model is a conditional diffusion engine that leverages a shared multimodal conditioning vector to (i) impute missing modalities via cross-modality reconstruction (e.g., inferring immune programs or gene-expression states from histology, or estimating tumor burden dynamics from plasma signals) and (ii) simulate system-level interventions through perturbation-based modeling.

This whole-system approach is designed to address core translational questions at the level of the individual patient: likelihood of response, mechanisms and timing of resistance, risk of immune-related toxicity, and rational selection of combinations and schedules grounded in patient state. The model co-trains across complementary modalities reflecting both tumor and systemic immunity, including tumor and blood RNA sequencing, tumor/normal DNA sequencing, T-cell receptor (TCR) repertoire features, plasma-based sequencing (ctDNA/cfDNA), multimodal imaging, and clinical context, learning shared representations under heterogeneity and partial observability to enable robust inference when some modalities are unavailable.

These capabilities enable Digital Twins, patient-specific latent profiles updated longitudinally to support blood-first monitoring, early response prediction, and detection of trajectory divergence that can precede progression or severe toxicity.

At BostonGene, these oncology foundation models are already being deployed in large, real-world clinical programs with major industry partners (including **Takeda, AstraZeneca, and Unilever**), with a primary focus on solving practical pharma and precision-medicine problems end-to-end rather than optimizing isolated molecular interactions.