Developing an Inverse Reinforcement Learning Methodology to Predict the Progression of Colorectal Cancer

SILBA S. DOWELL, The University of Wyoming, USA
DANIEL HINTZ, The University of Wyoming, USA
TYSON LIMATO, The University of Wyoming, USA
SHAD SELLERS, The University of Wyoming, USA
MILANA WOLFF, The University of Wyoming, USA
DR. NICHOLAS CHIA, Argonne National Lab, USA
DR. LIUDMILA S. MAINZER, The University of Wyoming, USA

In cancer biology, large amounts of high dimensional data (genomic, transcriptomic, proteomic, phenotypic, etc.) are required for any computationally relevant work. The problem is further complicated by the sheer size of the human genome, roughly three billion base pairs long. Therefore, computation is time-consuming and data-intensive. To solve this problem for human colorectal cancer, we are implementing a machine learning engine that is based on inverse reinforcement learning, and includes several different kinds of neural networks to perform data preparation, training, and prediction. Our work aims to reconstruct the progression of tumor development in a sample, and predict the next steps of its evolution, to aid in diagnosis and treatment. This poster will be presented as a work in progress methodology.

CCS Concepts: • **Inverse Reinforcement Learning**; • **Long Short Term Memory**; • **Bidirectional Encoder-Decoder network**; • **Tokenizers**; • **Transformers**;

Additional Key Words and Phrases: Cancer biology, Inverse Reinforcement Learning, Siamese encoder-decoder network, bidirectional PhyloWGS, Augur Pipeline, expert demonstrations, GISAID, Long Short Term Memory, embeddings, reward function, DNABERT-2, HyenaDNA, GenSLM, Markov Decision Process

Complex systems consist of many interacting agents, resulting in conglomerations of higher-level behaviors known as collective effects. Collective effects manifest as emergent properties that originate from the nature of the underlying dynamics of the system without any apparent central driving force [2] [5] [6]. The underlying principles of collective effects create order in living systems across all scales, from biomolecules to people. The same principles dictate the behavior of diseases ranging from viral infections to cancer. For instance, cellular population-level dynamics lead to cancer evolution in patients over time. We propose a model capable of effectively predicting cancer progression by using biological feature representations for clinical use.

Traditional efforts to understand the molecular biology of disease focus on observing simple correlations between a protein function and a disease, followed by decades of validation in different animal models. These static classical approaches usually address disease emergence or a particular stage in its progression. However, disease progression implies a temporal process marked by evolutionary changes in the underlying mechanisms of the disease. Dynamic

Manuscript submitted to ACM

approaches, such as Bayesian methods, Markov chain modeling, and various machine learning techniques, focus on characterizing a small number of well-defined stages of disease via critical mutations in a few genes and predicting the overall clinical outcome (e.g. survival, likelihood of recurrence) or binary classification (e.g. drug responders vs. non-responders) [3] [8]. These approaches are not designed to unravel the complexity of the entire evolutionary process that gradually drives disease progression via individual mutations, as is the case in cancer. Furthermore, these approaches offer limited insight into specific molecular mechanisms of disease.

Machine learning approaches are uniquely suited to the biomedical field due to the biological complexity of disease phenotypes [4] [9]. In the case of cancer, genomic variants, rearrangements, and clonal behavior affect tumor phenotypes. The interplay of these characteristics inhibit mutation identification efforts. Consequently, cancer research operates in a vast parameter space, rendering pure statistical approaches insufficient. Artificial intelligence offers a promising tool to reduce dimensionality and noise, simplifying research efforts [1] [10].

Specialized inverse reinforcement learning (IRL) algorithms parallel the step-by-step accumulation of somatic alterations in cancer evolution, distinguishing this approach from other machine learning methods. IRL attempts to estimate the reward function of a Markov decision process (MDP) from examples provided by an "expert demonstration" of optimal behavior, such as effective cancer cell replication. For cancer, this reward function represents the possible sequences of mutations through which cancer evolves, or traverses through its phylogenetic tree, assigning transition probabilities for every mutation. The time-ordered paradigm of IRL enables identification of potential causal relationships in carcinogenesis. Specifically, the reward function can be used to approximate the most influential causal mechanisms within clonal populations of cancer tumor cells. In contrast to previously used methods relying on association alone, which can only hypothesize co-occurrence, we propose the application of IRL to reconstructing the evolutionary history of a tumor. This paradigm-shifting AI approach can significantly improve our understanding of the molecular mechanisms of disease, allow us to accurately predict its course of progression, and enable more effective drug target identification for treatment. Our prior implementation of IRL demonstrates success in colorectal cancer progression prediction, but requires impractical amounts of time and computational resources [7].

We introduce efficient DNA encoding methods based on bidirectional encoder representations from transformers, convert existing code from Java and C to Python in order to enable parallel execution on multiple GPUs, and rewrite existing IRL implementations in terms of neural network approximations of reward functions. Furthermore, we leverage Oracle Cloud allocations to conduct an extensive parameter search to optimize neural network hyperparameters. Overall, these changes reduce the computational requirements, allowing use of larger amounts of data to obtain representation of features across numerous types of cancer.

Our proposed algorithm affords oncology researchers three major advantages over extant approaches. First, IRL identifies the molecular time-ordered events of cancer progression at unprecedented single-mutation resolution, rather than reducing complex mechanisms to univariate approximations. Second, IRL is generative. While a traditional classifier can only consume data and return a label, our model offers the novel feature of forward-simulating cancer progression, effectively predicting future tumor development. Third, the encoding schema improvements and code parallelization on GPUs, increase algorithm speed, making IRL usable on a clinically relevant time scale. We will use these advantages to understand the causal relationships implied by the learned parameters of the model-free paradigm, closing the distance between model performance and scientific insight.

ACKNOWLEDGMENTS

We would like to thank WY-INBRE, WY-REDD, the Wyoming Public Health Laboratory, University of Wyoming Department of Statistics for their financial support, and finally Oracle Research Cloud for the awarded $100,000 Project Award.

REFERENCES

[1] Nupur Biswas and Saikat Chakrabarti. 2020. Artificial Intelligence (AI)-Based Systems Biology Approaches in Multi-Omics Data Analysis of Cancer. *Frontiers in Oncology* 10 (2020). https://www.frontiersin.org/articles/10.3389/fonc.2020.588221

[2] Nicholas Chia, Carl R. Woese, and Nigel Goldenfeld. 2008. A collective mechanism for phase variation in biofilms. *Proceedings of the National Academy of Sciences* 105, 38 (Sept. 2008), 14597–14602. https://doi.org/10.1073/pnas.0804962105 Publisher: Proceedings of the National Academy of Sciences.

[3] Jiadong Chu, Na Sun, Wei Hu, Xuanli Chen, Nengjun Yi, and Yueping Shen. 2022. The Application of Bayesian Methods in Cancer Prognosis and Prediction. *Cancer Genomics & Proteomics* 19, 1 (Jan. 2022), 1–11. https://doi.org/10.21873/cgp.20298 Publisher: International Institute of Anticancer Research Section: Review.

[4] Marina De Rosa, Daniela Rega, Valeria Costabile, Francesca Duraturo, Antonello Niglio, Paola Izzo, Ugo Pace, and Paolo Delrio. 2016. The biological complexity of colorectal cancer: insights into biomarkers for early detection and personalized care. *Therapeutic Advances in Gastroenterology* 9, 6 (Nov. 2016), 861–886. https://doi.org/10.1177/1756283X16659790 Publisher: SAGE Publications Ltd STM.

[5] Nigel Goldenfeld and Carl Woese. 2007. Biology's next revolution. *Nature* 445, 7126 (Jan. 2007), 369–369. https://doi.org/10.1038/445369a Number: 7126 Publisher: Nature Publishing Group.

[6] Nigel Goldenfeld and Carl Woese. 2011. Life is Physics: Evolution as a Collective Phenomenon Far From Equilibrium. *Annual Review of Condensed Matter Physics* 2, 1 (2011), 375–399. https://doi.org/10.1146/annurev-conmatphys-062910-140509 _eprint: https://doi.org/10.1146/annurev-conmatphys-062910-140509.

[7] John Kalantari, Heidi Nelson, and Nicholas Chia. 2020. The Unreasonable Effectiveness of Inverse Reinforcement Learning in Advancing Cancer Research. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 01 (April 2020), 437–445. https://doi.org/10.1609/aaai.v34i01.5380 Number: 01.

[8] Konstantina Kourou, Themis P. Exarchos, Konstantinos P. Exarchos, Michalis V. Karamouzis, and Dimitrios I. Fotiadis. 2015. Machine learning applications in cancer prognosis and prediction. *Computational and Structural Biotechnology Journal* 13 (Jan. 2015), 8–17. https://doi.org/10.1016/j.csbj.2014.11.005

[9] Evangelia Koutsogiannouli, Athanasios G. Papavassiliou, and Nikolaos A. Papanikolaou. 2013. Complexity in cancer biology: is systems biology the answer? *Cancer Medicine* 2, 2 (2013), 164–177. https://doi.org/10.1002/cam4.62 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cam4.62.

[10] Norman E. Sharpless and Anthony R. Kerlavage. 2021. The potential of AI in cancer care and research. *Biochimica et Biophysica Acta (BBA) - Reviews on Cancer* 1876, 1 (2021), 188573. https://doi.org/10.1016/j.bbcan.2021.188573

Received 12 August 2023