

Developing a Machine Learning Tool for Dynamic Cancer Treatment Strategies

Jiaming Zeng

Department of Management Science and Engineering, Stanford University, Stanford, CA
jiaming@stanford.edu

Abstract

With the rising number and complexity of cancer therapies, it is increasingly difficult for clinicians to identify an optimal combination of treatments for a patient. Our research aims to provide a decision support tool to optimize and supplant cancer treatment decisions. Leveraging machine learning, causal inference, and decision analysis, we will utilize electronic medical records to develop dynamic cancer treatment strategies that advise clinicians and patients based on patient characteristics, medical history, and etc. The research hopes to bridge the understanding between causal inference and decision analysis and ultimately develops an artificial intelligence tool that improves clinical outcomes over current practices.

Summary

My research applies ML to develop a tool for clinicians to optimize cancer treatment decisions. In the past decade, there has been a reversal of 146 standard medical practices (Prasad et al. 2013), and many questions remain about the best treatment strategies. This is especially true for cancer treatments; with the rising number and complexity of cancer therapies, clinicians find it increasingly difficult to identify the **best combination of treatments and when to switch treatments**. Thus, there is an urgent need for better longitudinal strategies that manage treatment combination and timing while considering patient characteristics, medical history, and personal preferences. Ideally, an optimal treatment strategy will prolong life, decrease hospitalization, and improve patient quality of life.

For my Ph.D. research, **I will leverage machine learning (ML) to develop a decision support tool for dynamic cancer treatment strategies that will advise clinicians and patients based on patient characteristics, medical history, response to therapy, and side effects**. Specifically, I will begin by focusing on initial treatment planning for localized prostate, esophageal, and oropharynx cancer. The research is highly interdisciplinary, combining causal inference, decision analysis, and machine learning for medical application and advancement. Causal inference measures the effect of a treatment on an outcome. Decision analysis tools provide the framework that allows an optimal decision under uncertainty. With ML, the two fields can be combined to

help clinicians choose optimal treatments. Finally, the temporal nature of medical data makes it is an ideal subject for research on longitudinal methods for causal inference.

Aims and Contribution

Using the rich and growing information available in electronic medical records (EMR), I aim to develop a tool that can 1) build a longitudinal causal inference model for the effect of a therapy on patient outcomes, 2) design a cancer treatment strategy that can adapt to changing circumstances in a patient's disease progression, and 3) help plan systemic (e.g., chemotherapy, immunotherapy), radiation, and surgery treatments for various types of cancer. This research thus has the potential to help clinicians better apply ML tools and provide decision support systems that improve patient outcomes. The use of high-dimensional EMR data offers a key advantage over past studies, which used population-based databases that lack detailed information about many important confounders, i.e., variables that affect both treatment and outcomes (Williams et al. 2017). Furthermore, my research will help build generalizable knowledge on performing causal inference from observational data, develop methods for identifying personalized treatment strategy based on an individual's EMR data, and ultimately help improve clinical outcomes over current practices.

Research Plan

I will use the Stanford Cancer Institute Research Database (SCIRDB), a rich set of data that integrates resources from EPIC, STARR, databases in surgical pathology and radiation oncology, and the Stanford Cancer Registry. It contains data on over 100,000 patients collected since 2008 and provides rich clinical information from oncology, inpatient stays, radiology, nursing, and labwork. Internal Review Board approval has already been acquired for the research.

Step 1: Creating computer-understandable format of clinical data A challenge with EMR data is that much of it is unstructured (i.e., in the form of clinical notes). To create computable representations of this data, I will build on the work in Banerjee et al. (2018), which developed a hybrid pipeline that combines semantic data mining with neural embedding to extract a vector-based dataset from the multiple types of free-text clinical notes. I will combine the learned vector-based datasets in Banerjee et al. (2018)

with the structured EMR data to create a comprehensive dataset with patient characteristics (e.g., age, gender), medical history, tumor characteristics, hospitalization history, and treatment history. I am currently developing the natural language processing (NLP) language model for extracting the necessary information from the free-text clinical notes. The groundwork of this research is submitted as a student abstract for the AAAI-20 Conference and I expect to have completed this portion by AAAI-20 Conference.

Step 2: Decision-making tool I will build on the existing methods in Chen et al. (2017) and Schulam and Saria (2017). Chen et al. (2017) adapted a Markov Decision Process (MDP) into a sequential decision-making tool for therapy choice in a class of metastatic breast cancer patients. This model maintained a running belief about the effectiveness of each therapy treatment as observations were made and used it to make therapy recommendations. Schulam and Saria (2017) introduced the Counterfactual Gaussian Process (CGP), an algorithm based on the potential-outcome model (Rubin 1974), Gaussian processes (Rasmussen 2003), and marked point processes (Daley and Vere-Jones 2007), and designed to support decision-making in the longitudinal setting. When applied to individualized treatment planning, this algorithm showed promising results.

I will use Schulam and Saria (2017)'s CGP model in place of Chen et al. (2017)'s MDP model. Because the CGP builds a causal inference model for longitudinal data and predicts the future progression from sequences of observations and treatments at irregular time intervals, it is ideal for this research. Based on the framework in Chen et al. (2017), I will design a strategy tailored to a patient or class of patients that answers 1) what the most effective current therapy is, and 2) when clinicians should switch to another therapy.

For the potential-outcomes framework (Rubin 1974), the treatments are the various types of existing cancer treatments (i.e., systemic, radiation, and surgery). The outcome will be the life expectancy and quality of life. I will parametrize the CGP inputs with domain knowledge from clinical trials or expert judgment from clinicians. I will also perform sensitivity analysis on various parameters – e.g., chemotherapy toxicity – to identify when the causal assumptions conflict with the dataset, thereby making these methods more acceptable in practice. Patient preferences can also play a major role in treatment decisions and I will explore how they can be incorporated into the model as additional parameters.

An important part of the research will be identifying the most significant potential confounders and finding ways to control for them. As variables, confounders can obscure the real effect of the treatment. In the longitudinal setting, we have the additional problem that both the type and timing of treatments may depend on the outcome. Although the CGP explicitly addresses the confounding problem, it still requires that relevant confounders be observed. By AAAI-20 Conference, I expect to be in the process of developing the causal inference models and have developed some method for identifying confounding from the EMR dataset.

Step 3: Validate causal inference model against randomized studies As described in Schulam and Saria (2017),

the validity of the causal inference models cannot be quantitatively evaluated without prospective experimental data. In addition to measuring fit on observational data, I will compare the treatment recommendations against the final results from comprehensive randomized clinical trials. Specifically, I will validate the causal inference model by examining how well it can reproduce the results from these established clinical studies for prostate (Hamdy et al. 2016), esophageal (van Hagen et al. 2012), and oropharynx cancer (Nichols et al. 2019). I will then benchmark the performance of the CGP model against the MDP model in Chen et al. (2017) and compare the quality of the decisions recommended by the two approaches.

References

- Banerjee, I.; Gensheimer, M. F.; Wood, D. J.; Henry, S.; Aggarwal, S.; Chang, D. T.; and Rubin, D. L. 2018. Probabilistic prognostic estimates of survival in metastatic cancer patients (ppes-met) utilizing free-text clinical narratives. *Scientific reports* 8(1):10037.
- Chen, X.; Shachter, R. D.; Kurian, A. W.; and Rubin, D. L. 2017. Dynamic strategy for personalized medicine: An application to metastatic breast cancer. *Journal of biomedical informatics* 68:50–57.
- Daley, D. J., and Vere-Jones, D. 2007. *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media.
- Hamdy, F. C.; Donovan, J. L.; Lane, J. A.; Mason, M.; Metcalfe, C.; Holding, P.; Davis, M.; Peters, T. J.; Turner, E. L.; Martin, R. M.; et al. 2016. 10-year outcomes after monitoring, surgery, or radiotherapy for localized prostate cancer. *New England Journal of Medicine* 375(15):1415–1424.
- Nichols, A. C.; Theurer, J.; Prisman, E.; Read, N.; Berthelet, E.; Tran, E.; Fung, K.; de Almeida, J. R.; Bayley, A.; Goldstein, D. P.; et al. 2019. Radiotherapy versus transoral robotic surgery and neck dissection for oropharyngeal squamous cell carcinoma (orator): an open-label, phase 2, randomised trial. *The Lancet Oncology* 20(10):1349–1359.
- Prasad, V.; Vandross, A.; Toomey, C.; Cheung, M.; Rho, J.; Quinn, S.; Chacko, S. J.; Borkar, D.; Gall, V.; Selvaraj, S.; et al. 2013. A decade of reversal: an analysis of 146 contradicted medical practices. In *Mayo Clinic Proceedings*, volume 88, 790–798. Elsevier.
- Rasmussen, C. E. 2003. Gaussian processes in machine learning. In *Summer School on Machine Learning*, 63–71. Springer.
- Rubin, D. B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology* 66(5):688.
- Schulam, P., and Saria, S. 2017. Reliable decision support using counterfactual models. In *Advances in Neural Information Processing Systems*, 1697–1708.
- van Hagen, P.; Hulshof, M.; Van Lanschot, J.; Steyerberg, E.; Hene-gouwen, M. V. B.; Wijnhoven, B.; Richel, D.; Nieuwenhuijzen, G.; Hospers, G.; Bonenkamp, J.; et al. 2012. Preoperative chemoradiotherapy for esophageal or junctional cancer. *New England Journal of Medicine* 366(22):2074–2084.
- Williams, S. B.; Huo, J.; Chamie, K.; Smaldone, M. C.; Kosarek, C. D.; Fang, J. E.; Ynalvez, L. A.; Kim, S. P.; Hoffman, K. E.; Giordano, S. H.; et al. 2017. Discerning the survival advantage among patients with prostate cancer who undergo radical prostatectomy or radiotherapy: the limitations of cancer registry data. *Cancer* 123(9):1617–1624.