

How Do Tumor Cytogenetics Inform Cancer Treatments? Dynamic Risk Stratification and Precision Medicine Using Multi-armed Bandits

Zhijin Zhou, Yingfei Wang, Hamed Mamani

Foster School of Business, University of Washington, Seattle, Washington 98195,
{zjzhou@uw.edu, yingfei@uw.edu, hmamani@uw.edu}

David G. Coffey*

Program in Immunology, Fred Hutchinson Cancer Research Center, Seattle, WA 98109,
Department of Medicine, School of Medicine, University of Washington, Seattle, WA 98195, dcoffey@fredhutch.org

Multiple myeloma is an incurable cancer of bone marrow plasma cells with a median overall survival of 5 years. With newly approved drugs to treat this disease over the last decade, physicians are afforded more opportunities to tailor treatment to individual patients and thereby improve survival outcomes and quality of life. However, since the optimal sequence of therapy is unknown, selecting a treatment that will result in the most effective outcome for each individual patient is challenging. To understand patients' treatment responses, we develop an econometric model – the Hidden Markov model, to systematically identify changes in patients' risk levels. Based on a fine-grained clinical dataset from Seattle Cancer Care Alliance (Seattle, WA) that includes patient-level cytogenetic information, we find that, other than the manifestation of cytogenetic features, previous exposure to certain drugs also affect patients' underlying risk levels. The effectiveness of different treatments varies significantly among patients, which calls for personalized treatment recommendations.

We then formulate the treatment recommendation problem as a Bayesian contextual bandit, which sequentially selects treatments based on contextual information about patients and therapies, with the goal of maximizing overall survival outcomes. Facing the difficulty of evaluating the performance of the policy without field experiments in medical practice, we integrate the structural econometric model into bandit optimization and generate counterfactuals to support the theoretical exploration/exploitation framework with empirical evidence. Compared with clinical practices and benchmark strategies, our method suggests a rise in overall survival outcomes, with higher improvement for aging or high-risk patients with more complications.

Key words: Multiple Myeloma, Precision Medicine, Multi-Armed Bandit, Thompson Sampling, Hidden Markov Model (HMM)

* Clinical data abstraction was funded by financial support from the Seattle Translational Tumor Research group at the Fred Hutchinson Cancer Research Center, Seattle, WA. David Coffey received financial support from NIH grant 2K12CA076930-16A1.