Pratik Gajane | pratik.gajane@gmail.com | +31-633-313-415 | pratikgajane.github.io

Research Statement

My overarching goal is to create machine learning-driven healthcare solutions that excel in utility while also addressing broader socio-technical goals such as ensuring equity across protected attributes like gender and race, giving precedence to patient privacy, and improving explainability. Below, I explain some of the salient threads of my research.

Personalized and Fair Care for Chronic Pain

Chronic pain significantly diminishes the quality of life for millions worldwide. While psychoeducation and therapy can improve pain outcomes, many individuals experiencing pain lack access to evidence-based treatments or fail to complete the necessary number of sessions to achieve benefit. Reinforcement learning (RL) shows potential in tailoring personalized pain management interventions according to patients' individual needs while ensuring the efficient use of scarce clinical resources. However, clinicians, patients, and healthcare decision-makers are concerned that RL solutions could exacerbate disparities associated with patient characteristics like race or gender. In my recent work [1], conducted in collaboration with Prof. John D. Piette, we studied *gender fairness* in personalized pain care recommendations using a real-world application of reinforcement learning [28]. Here, adhering to gender fairness translates to minimal or no disparity in the utility received by subpopulations as defined by gender. We investigated whether the selection of relevant patient information (referred to as features) used to assist decision-making affects gender fairness. Our experiments, conducted using real-world data [29], indicated that included features can impact gender fairness. Moreover, we proposed an RL solution that demonstrates the ability : i) to adaptively learn to select the features that optimize for utility and fairness, and ii) to accelerate feature selection and in turn, improve pain care recommendations from early on, by leveraging clinicians' domain expertise.

Ongoing and Future Work : Continuing this line of work further, we plan to investigate the clinical significance of the difference in the utility received by subpopulations. We will also investigate disparities related to social determinants of health such as patients' racial identification, educational attainment, geographic access to care, or comorbidities. I intend to explore a refined approach to chronic pain care, framing it as a Markov decision process to better reflect real-world dynamics. Additionally, I aim to expand our existing problem formulation to accommodate dynamic or non-stationary environments, where treatment effectiveness for patients may evolve over time. My previous work [20] in detecting such non-stationarities efficiently leading to proactive decision-making would be useful in this direction. More information about my general research objectives on the topic of personalized, fair, and adaptive pain care using reinforcement learning can be found in the following proposal.

Explainable Machine Learning in Healthcare

Many modern machine learning algorithms are applied in a black-box manner, wherein the method through which the algorithm produces output from input is not explainable. This is unsuitable for high-impact domains like healthcare where the lack of explanations might be unacceptable for patients, healthcare professionals, and policymakers. Accordingly, there has been a surge in research that aims to explain machine learning solutions in the healthcare domain. In [10], we highlighted the inadequacies and inconsistencies of prevailing explanation techniques. To address this gap, we proposed a new data-centric explanation technique. In our experiments conducted on real-world UCI datasets on Diabetes and breast cancer, our proposed solution outperformed the state-of-the-art methods by over 50% considering the performance metric of root mean squared error (RMSE) measured between the ML model to be explained and the provided explanation.

Ongoing and Future Work : In [10] (and related prior research), explanations are generated using synthetic data that disregards the distribution of the original data, namely, the patient population in this context. However, the relevance and effectiveness of an explanation could be contingent upon the distribution of the patient population. To mitigate this concern, I intend to focus on generating explanations using synthetic data that respects the distribution of the patient population.

Fairness-aware Machine Learning in Healthcare

In [21], we presented theoretical as well as empirical critiques of the fairness formalizations used in the machine learning literature. We explained how these critiques limit the suitability of fairness formalizations to certain domains like healthcare. For example, we showed that counterfactual fairness formalizations may not be suitable for healthcare due to the presence of hindsight bias and outcome bias.

Ongoing and Future Work : In an ongoing collaboration with a social scientist from the University of Wisconsin, Madison, we are formalizing an egalitarian fairness notion called *equality of resources* [30]. This principle, previously studied in the social sciences literature, holds potential applications in healthcare resource allocation [30].

References

References [1] to [27] can be found in my CV.

[28] John D. Piette, Sean Newman, Sarah L. Krein, Nicolle Marinec, Jenny Chen, David A. Williams, Sara N. Edmond, Mary Driscoll, Kathryn M. LaChappelle, Robert D. Kerns, Marianna Maly, H. Myra Kim, Karen B. Farris, Diana M. Higgins, Eugenia Buta, and Alicia A. Heapy. Patient-centered pain care using artificial intelligence and mobile health tools: A randomized comparative effectiveness trial. *JAMA Internal Medicine*, 182(9), September 2022. ISSN 2168-6106. doi: 10.1001/jamainternmed.2022.3178. URL `http://dx.doi.org/10.1001/jamainternmed.2022.3178`.

[29] John Piette. Dataset #2 for piette et al, data in brief: Data for a reinforcement learning intervention to treat chronic pain. Mendeley Data, 2022. URL `https://data.mendeley.com/datasets/33mkbm32dz/1`.

[30] Ronald Dworkin. *Sovereign Virtue: The Theory and Practice of Equality*. Harvard University Press, 2000. ISBN 9780674002197. URL `http://www.jstor.org/stable/j.ctv1c3pd0r`.