

---

**TORYN QWYLLYN KLASSEN**, University of Toronto

*Remembering to Be Fair: Non-Markovian Fairness in Sequential Decision Making*

Fair decision making has largely been studied with respect to a single decision. Here we investigate the notion of fairness in the context of sequential decision making where multiple stakeholders can be affected by the outcomes of decisions. We observe that fairness often depends on the history of the sequential decision-making process, and in this sense that it is inherently non-Markovian. We further observe that fairness often needs to be assessed at time points within the process, not just at the end of the process. To advance our understanding of this class of fairness problems, we explore the notion of non-Markovian fairness in the context of sequential decision making. We identify properties of non-Markovian fairness, including notions of long-term, anytime, periodic, and bounded fairness. We explore the interplay between non-Markovian fairness and memory and how memory can support construction of fair policies. Finally, we introduce the FairQCM algorithm, which can automatically augment its training data to improve sample efficiency in the synthesis of fair policies via reinforcement learning.

This is joint work with Parand A. Alamdari, Elliot Creager, and Sheila A. McIlraith.