

Building Contextual Clinical AI: Learning What Matters and When to Speak Up

Brett Beaulieu-Jones

Assistant Professor, Department of Medicine, University of Chicago

September 24, 2026

1:30PM-2:50PM

Li Ka Shing Center, Room:120

Abstract:

Electronic health records contain rich information about patients but do so through the lens of clinical care. As a result, predictive models may achieve strong predictive performance while learning artifacts of care delivery rather than a patient's underlying health. In this talk I will examine the challenge of separating clinical decision making and care processes from patient health to make progress towards contextual AI. Such AI must understand what is happening, why it is happening, when it matters, and how it fits within the broader trajectory of care. It must also know when to speak up and when to remain silent, because even a correct alert or recommendation can be distracting, disruptive, or harmful when delivered in the wrong context.

Reading List:

1. Beaulieu-Jones, Brett K., William Yuan, Gabriel A. Brat, Andrew L. Beam, Griffin Weber, Marshall Ruffin, and Isaac S. Kohane. "Machine learning for patient risk stratification: standing on, or looking over, the shoulders of clinicians?." *NPJ digital medicine* 4, no. 1 (2021): 62.
2. Wilson, F. Perry, Melissa Martin, Yu Yamamoto, Caitlin Partridge, Erica Moreira, Tanimu Arora, Aditya Biswas et al. "Electronic health record alerts for acute kidney injury: multicenter, randomized clinical trial." *bmj* 372 (2021).
3. Burkhart, Michael C., Bashar Ramadan, Luke Solo, William F. Parker, and Brett K. Beaulieu-Jones. "Quantifying surprise in clinical care: Detecting highly informative events in electronic health records with foundation models." In *Biocomputing 2026: Proceedings of the Pacific Symposium*, pp. 173-188. 2025.