Sebastian Vollmer.
University of Warwick.

23rd October 2020, 1pm UK time

Part 1: Risk prediction and Risk prediction (work in progress)
Pre-empting and avoiding emergency admission (EA) is potentially advantageous both to individual health and the overall efficiency of the healthcare system. We present a model which predicts the risk of an emergency admission for each person in Scotland with predictors derived from their electronic health records (EHRs). This builds on an existing simpler model which has been in clinical use since 2012. Our data comprise acute hospital records (emergency admissions, elective admissions, day cases, outpatient attendances, A&E attendances, long term conditions) and community prescribing activity with NHS Scotland over the pre-prediction period (approximately 80% of the Scottish population), with outcome defined as an emergency admission to an inpatient unit in the year following prediction. Motivated by this problem we also study the abstract problem of updating of a predictive score for a binary outcome when an existing predictive score forms part of the standard workflow, driving interventions. In this setting, the existing score induces an additional causative pathway which leads to miscalibration when the original score is replaced.

Part 2: Machine Learning in Julia and benchmarking results on predictive fairness.
MLJ (Machine Learning in Julia) (Blaom, 2019) is a toolbox written in Julia that provides a common interface and meta-algorithms for selecting, tuning, evaluating, composing and comparing machine learning model implementations written in Julia and other languages. More broadly, the MLJ project hopes to bring cohesion and focus to a number of emerging and existing, but previously disconnected, machine learning algorithms and tools of high quality, written in Julia. A welcome corollary of this activity will be increased cohesion and synergy within the talent-rich communities developing these tools. As a use case we benchmark evaluate the performance of pre-processing and post-processing debiasers for improving fairness in random forest classifiers trained on a suite of data sets. Specifically, we study how these debiasers generalize with respect to both out-of-sample test error for computing fairness–performance and fairness–fairness trade-offs, and on the change in other fairness metrics that were not explicitly optimised. Our results demonstrate that out-of-sample performance on fairness and performance can vary substantially and unexpectedly. Moreover, the variance in estimation arises from class imbalances with respect to both the outcome and the protected classes. Our results highlight the importance of evaluating out-of-sample performance in practical usage. https://arxiv.org/pdf/2007.12285.pdf