

## **Pre-read: healthcare (part 1)**

COS 597E/SOC 555, Princeton University, Fall 2020

If we can more accurately diagnose health conditions, prognose disease progression, and predict disease risk, we may be able to intervene earlier and more effectively. This will save lives and trillions of dollars. Do you have a qualitative sense of how widely data-driven methods are used for these purposes today in healthcare? How do you expect this to change in the future?

The three readings for Tuesday form a progression advocating for increasingly complex statistical methods in healthcare domains (for different types of tasks): the first is about simple formulas and decision rules, the second is about interpretable models of intermediate complexity, and the third is about deep learning.

The Dawes et al. paper is a classic that is based on a few early studies comparing the performance of expert predictions against statistical methods. Make sure you understand what the authors mean by “clinical” and “actuarial”. As you read the paper, think about (1) the similarities and differences between the experts-vs-statistical-predictions comparisons described here and those by Tetlock; (2) whether you agree with the overall conclusions and especially the claim that these findings have encountered resistance from clinicians due to ignorance and irrationality.

The Caruana et al. paper is about the dangers of black-box models in healthcare. The main lesson is nicely illustrated by the story that opens the paper. Can you try to succinctly summarize the root cause of the problem in a way that can potentially be applied to other domains?

The third paper is one of a series of papers describing deep learning techniques developed (in part) by Google that made waves in the press for claiming to match or beat the performance of medical experts at detecting health conditions from medical scans or images. Note that this paper is about “predicting the present”. What sort of impact do you expect this type of research to have on healthcare? Will it replace the work of some types of medical practitioners?

Thursday’s readings will be more explicitly about limits: ethical concerns about predictive technologies, pitfalls of performance evaluation, and scientific limits to prediction.