

Arshak Navruzyan

Startup.ML

Avoiding bad machine learning predictions in critical decision domains

As machine learning gets used for more and more critical decisions, what is the human cost of bad predictions? So far ML has been used for guiding advertising spend decisions, information retrieval, consumer-grade image processing and we've already seen examples of ML models making blatantly bad predictions. What will happen when ML is making medical decision, navigating a freight truck or in a criminal justice case?

We focus on ways to understand how bad predictions happen and how to mitigate their effect.