KE LI, Simon Fraser University

Rethinking Regression: Insights from Machine Learning

Regression problems arise every time one would like to predict a continuous-valued variable, be it the colour of a pixel, a 3D position, a system configuration or a feature vector. It is well known that regression with square loss yields the conditional mean as the prediction. This is undesirable when there could be many predictions that are all correct, since the conditional mean would effectively average over these predictions and could be far from any of them. As an example, when the prediction takes the form of an image, the conditional mean tends to be blurry and desaturated. On the other hand, in classification problems, ambiguity in labels does not cause an issue because classifiers produce a distribution over class labels as output. Is it possible to get the best of both worlds? In this talk, I will show how to do so using a simple technique, known as conditional Implicit Maximum Likelihood Estimation.