

Can Requirements Make Machine Learning Safer? From constrained prediction to structured generation

Eleonora Giunchiglia¹

¹*Imperial College London, United Kingdom*

Abstract

As machine learning systems move from narrow prediction tasks to generative and decision-making settings, their ability to satisfy domain requirements is becoming increasingly important. Yet deep neural networks can still produce outputs that violate even simple logical, structural, or regulatory constraints. This talk argues that requirements should be treated as first-class citizens in machine learning: not only as criteria for evaluating a model after deployment, but as objects that can shape what models learn, generate, and optimize for from the outset. I will show how logical requirements can define safe output spaces and be compiled into neural architectures, enabling predictions and generations that are compliant by design. Through examples in autonomous driving, tabular data generation, and neurosymbolic reasoning, the talk will illustrate how requirements can make learning systems safer, more data-efficient, and better aligned with the structure of the problems they are meant to solve.

