

Learning Dynamic Generative Models via Causal Optimal Transport

Beatrice Acciaio¹, Michael Munn², Tianlin Xu³

^{1,3} London School of Economics and Political Science, ² Google, NY

In this talk I will present a new method to train generative models, based on non-anticipative (causal) optimal transport in conjunction with Sinkhorn divergences. The generator is trained to produce real-looking evolutions of processes, given long time-series of data. To improve its training, a discriminator is set against it, with the task of evaluating the distance between real and fake data. We employ recurrent neural network architectures to account for the dynamic nature of the problem. A conditional modification of our model leads to prediction of sequential data. Several applications will be discussed.