

Stochastic Video Generation with a Learned Prior

Rob Fergus, NYU / Facebook

Generating video frames that accurately predict future world states is challenging. Existing approaches either fail to capture the full distribution of outcomes, or yield blurry generations, or both. In this paper we introduce an unsupervised video generation model that learns a prior model of uncertainty in a given environment. Video frames are generated by drawing samples from this prior and combining them with a deterministic estimate of the future frame. The approach is simple and easily trained end-to-end on a variety of datasets. Sample generations are both varied and sharp, even many frames into the future, and compare favorably to those from existing approaches. Joint work with Emily Denton.