

In the universal simulation setting, a training sequence from an unknown information source of a certain parametric family (e.g., the family of memoryless sources or Markov sources), and a stream of purely random bits, are provided. The objective is to generate another random sequence (of the same length or shorter) whose conditional entropy given the training sequence is maximal, subject to a certain distance constraint between the probability laws governing the output sequence and the input, training sequence. This problem has been recently studied in a series of papers in which the distance constraint requires the two sources to be identical. In this talk, we will first review the known results (including the delay-limited case) and their connections to universal data compression. We will then discuss the case in which the Markov model order is unknown, which naturally leads to relaxing the requirement of identical distributions. Finally, we will consider the case in which the distance constraint allows for the two distributions to differ (universal simulation with a fidelity criterion).