

# Masked diffusions: how to choose an optimal schedule and how much is gained?

Prof. Dr. Hugo Lavenant  
(Università Bocconi Milan, Italy)

June 26, 2026

We look at the computational error made in the generation phase of Masked Diffusion Models, which were recently proposed for generative modelling of discrete data. These models make approximations of conditional dependence to speed up generation of new samples at the price of some bias in the sampling distribution. We explain how these models can be simply described once you do not try to make a parallel with continuous diffusion models. From this elementary description, we study the resulting computation-vs-accuracy trade-off, providing general error bounds in relative entropy, and we investigate the gains obtained by using non-constant schedule sizes.

This is joint work with Giacomo Zanella.