

## **The generalization of flow matching and its temporal phases: why and when does it work?**

Mathurin Massias

A growing body of research aims to understand why diffusion and flow matching generalize so effectively, with the puzzling observation that, trained perfectly, they should only reproduce their training data. In this talk, we study reasons why generative models fail to learn the exact minimizer of their training loss. We first rule out the noisy nature of the loss as primary driver of generalization, showing that the stochastic and exact versions of the flow matching loss yield the same performance. We then show that failure to properly learn the exact solution occurs at small and large times, with small times being the most important for generalization. Finally, by adopting a denoising perspective on flow matching, we provide new characterization and insights on the temporal phases of the generative process.