ANASTASIS KRATSIOS, McMaster University
Incremental Generation is Necessity and Sufficient for Universality in Flow-Based Modelling

Incremental flow-based denoising models have reshaped generative modelling, but their empirical advantage still lacks a rigorous approximation-theoretic foundation. We show that incremental generation is necessary and sufficient for universal flow-based generation on the largest natural class of self-maps of $[0,1]^d$ compatible with denoising pipelines, namely the orientation-preserving homeomorphisms of $[0,1]^d$. All our guarantees are uniform on the underlying maps and hence imply approximation both samplewise and in distribution.

Using a new topological-dynamical argument, we first prove an impossibility theorem: the class of all single-step autonomous flows, independently of the architecture, width, depth, or Lipschitz activation of the underlying neural network, is meagre and therefore not universal in the space of orientation-preserving homeomorphisms of $[0,1]^d$. By exploiting algebraic properties of autonomous flows, we conversely show that every orientation-preserving Lipschitz homeomorphism on $[0,1]^d$ can be approximated at rate $\mathcal{O}(n^{-1/d})$ by a composition of at most K_d such flows, where K_d depends only on the dimension. Under additional smoothness assumptions, the approximation rate can be made dimension-free, and K_d can be chosen uniformly over the class being approximated. Finally, by linearly lifting the domain into one higher dimension, we obtain structured universal approximation results for continuous functions and for probability measures on $[0,1]^d$, the latter realized as pushforwards of empirical measures with vanishing 1-Wasserstein error.

Joint work: Hossein Rouhvarzi