TRAINING-FREE GUIDANCE WITH APPLICATIONS TO PROTEIN ENGINEERING

Lewis Cornwall, Joshua Meyers, James Day, Lilly Wollman, Neil Dalchau, Aaron Sim
Synteny
{lewis,josh,james,lilly,neil,aaron}@synteny.ai

ABSTRACT

Diffusion models facilitate powerful control over the generative process. Here we introduce training-free guidance, a method for sampling from a broad class of conditional distributions that can be considered generalisations of inpainting. The method is grounded in annealed Langevin dynamics which ensures convergence to the exact conditional distribution, unlike popular methods for inpainting which rely on heuristics. We demonstrate training-free guidance using pretrained unconditional models for protein structure, and protein sequence generation and improve upon state-of-the-art approaches. We show the versatility of training-free guidance by addressing a wider range of tasks, including multi-motif scaffolding and amino acid mutagenesis of T cell receptors, with applications to biologics design.