

Diffusion Models for Visual Content Generation

Niloy Mitra¹ 

¹University College London, United Kingdom

Abstract

Diffusion models are now the state-of-the-art for producing images. These models have been trained on vast datasets and are increasingly repurposed for various image processing and conditional image generation tasks. We expect these models to be widely used in Computer Graphics and related research areas. Image generation has evolved into a rich promise of new possibilities, and in this tutorial, we will guide you through the intricacies of understanding and using diffusion models. This tutorial is targeted towards graphics researchers with an interest in image/video synthesis and manipulation. Attending the tutorial will enable participants to build a working knowledge of the core formulation, understand how to get started in this area, and study practical use cases to explore this new tool. Our goal is to get more researchers with expertise in computer graphics to start exploring the open challenges in this topic and explore innovative use cases in CG contexts in image synthesis and other media formats. From understanding the underlying principles to hands-on implementation, you'll gain practical skills that bridge theory and application. Throughout the tutorial, we'll explore techniques for generating lifelike textures, manipulating details, and achieving remarkable visual effects. By the end, you'll have a solid foundation in utilizing diffusion models for image generation, ready to embark on your creative projects. Join us as we navigate the fascinating intersection of computer graphics and diffusion models, where pixels become canvases and algorithms transform into brushes.
