Master Thesis Proposal for the Master on Computer Vision (2026/27)

Learning Affective Representations for Emotionally Controllable Landscape Generation

Supervisors: Albert Clapés and Sergio Escalera

This thesis aims to advance the emerging field of affective generative AI by investigating how modern generative models, particularly diffusion-based architectures, can be guided to synthesize emotionally expressive landscape images. The research will focus on learning rich visual-affective representations from a novel large-scale dataset that explicitly links natural landscape imagery with human-perceived emotional responses, providing a unique foundation for controllable emotion-driven image generation.

The project will explore state-of-the-art Diffusion Transformer (DiT) architectures and advanced conditioning strategies, including Adaptive Layer/Group Normalization (AdaLN/AdaGN), parameter-efficient adaptation through LoRA, and ControlNet-inspired structural guidance. These techniques will be studied in both text-to-image and image-to-image generation settings, with the objective of enabling precise control over the emotional content while preserving visual realism and semantic consistency.

Beyond image generation, the thesis will investigate the relationship between individual landscape constituents—such as water, vegetation, terrain, sky, lighting, and atmospheric conditions—and the emotional responses they evoke. By disentangling and modeling these contributions, the research seeks to uncover interpretable visual factors that shape human affective perception, providing valuable insights into the connection between natural environments and emotional experience.

Ultimately, the thesis will contribute a novel, explainable, and highly controllable framework for affective landscape generation, bridging advances in conditional generative modeling with computational models of human emotion. The project is expected to push the state of the art in emotionally guided image generation and image-to-image translation through the development of improved conditional diffusion models. In addition, the resulting dataset, trained models, and implementation will be released as open research resources, with the goal of producing a high-impact scientific publication.