Toward fast, modular and explainable physicsinformed generative AI for image restoration: a diffusion deep unfolding, distillation and Langevin sampling approach

Marcelo Pereyra

This talk introduces a novel mathematical and computational framework for Bayesian image and video restoration that leverages distilled generative denoising diffusion models as powerful and flexible image/video priors. Our approach is built on two key innovations. First, we develop a new plug-and-play paradigm for embedding distilled generative models within Langevin diffusion processes, enabling the formulation of generative restoration techniques with clear, interpretable components that separately model statistical image priors and data sensing. This structure allows instrumental and degradation models to be specified at inference time and seamlessly integrated with pre-trained generative priors such as modern Stable Diffusion models. The second innovation combines Langevin sampling with deep unfolding to design new neural network architectures tailored for Bayesian imaging. When coupled with modern adversarial distillation techniques, this leads to highly efficient algorithms that achieve state-of-the-art restoration performance with as few as 4 to 8 neural function evaluations. We validate our framework through comprehensive experiments, demonstrating that it not only matches or exceeds the accuracy of current generative restoration methods but also offers significant improvements in computational efficiency and adaptability across diverse sensing scenarios. The material for this talk is drawn from our recent work on generative AI for image restoration, namely https://arxiv.org/abs/2503.12615 that utilises a distilled latent Stable Diffusion XL image-text model in a zero-shot manner (ICCV 2025) and its extension to video https://arxiv.org/abs/2510.01339, as well as https://arxiv.org/pdf/2507.02686 which integrates these ideas with deep unfolding and diffusion distillation techniques.