PRODUCT OF GAUSSIAN MIXTURE DIFFUSION MODELS

MARTIN ZACH

Diffusion model have demonstrated remarkable abilities in image generation. However, interpreting the classical deep networks with respect to the diffusion partial differential equation (DPDE) remains challenging. We lay a foundation for tackling this problem by considering one-layer networks with (log-) Gaussian mixture activations. For models acting on filter-, wavelet-, and shearlet-responses, we show how to adapt the activations such that the model fulfills the DPDE exactly. The models can be trained over the entire diffusion horizon using empirical Bayes. By interpreting our models as time-conditional likelihood densities, they can be used for noise estimation and blind heteroscedastic denoising. Numerical results for image denoising show that our models are competitive while being tractable, interpretable, and having only a small number of learnable parameters.