Deconvolution using ADMM with Diffusion Denoising Prior

Zakaria Patel, Akash Haridas, Kartikaeya Kumar
MSc in Applied Computing (MScAC), University of Toronto

Motivation

- Deconvolution is an inverse problem wherein the goal is to recover a clean image from a blurry one, with applications in medical imaging, astronomy, microscopy, etc.
- Alternating Direction Method of Multipliers (ADMM) [1] is a general algorithm for solving such inverse problems which can be guided by our understanding of what the solution should look like by using a prior.
- The presence of high noise in the image makes this problem challenging, necessitating an effective denoising prior in ADMM.

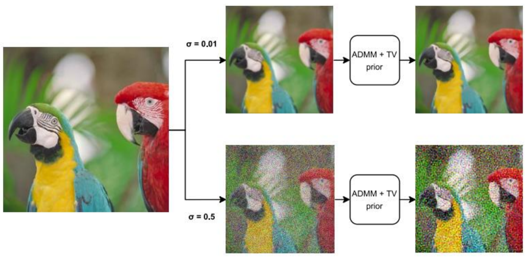

- Diffusion models have recently been shown to produce high quality images from pure noise through iterative denoising [3].

Background

- The optimal solution to a deconvolution problem can be formulated as,

$$\underset{\{x\}}{\text{minimize}} \; \underbrace{\frac{1}{2}\|\mathbf{Ax}-\mathbf{b}\|_2^2}_{f(\mathbf{x})} + \underbrace{\lambda\Psi(\mathbf{z})}_{g(\mathbf{z})} \qquad \text{subject to } \mathbf{Dx}-\mathbf{z}=0$$

- Using Lagrangian optimization, this simplifies to the iterative ADMM algorithm (flowchart below), where any general denoiser can be plugged into the z-update to guide the ADMM process [2].

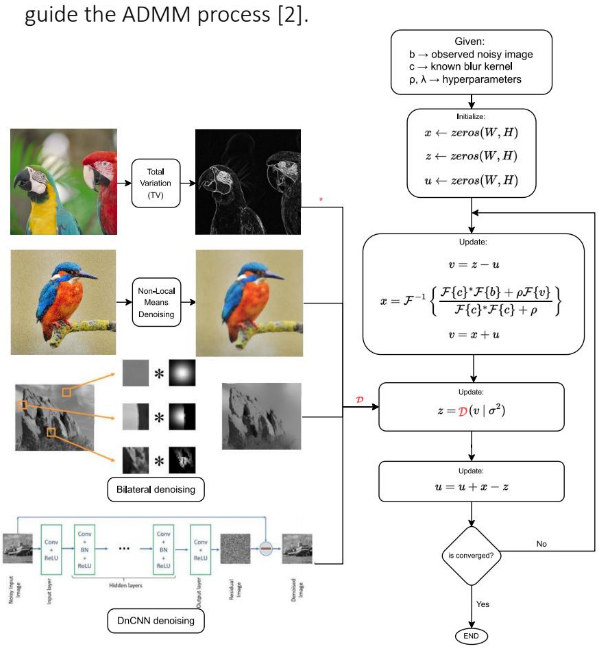

New Technique

- The proposed algorithm uses a diffusion denoiser for the $\mathbf{z}$-update in ADMM. A diffusion model (DM) progressively adds noise to an image in a forward Markov chain defined by:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}) \qquad (1)$$

and then recovers a plausible sample from the data distribution through a reverse denoising process parametrized by a neural network μ_θ.

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) := \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)) \qquad (2)$$

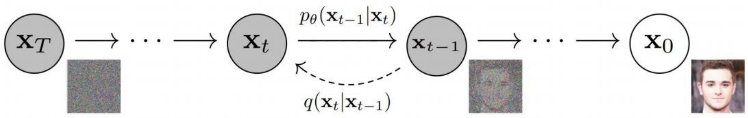

- From equation (1), the noisy image $\boldsymbol{x_t}$ at timestep $\boldsymbol{t}$ in the forward diffusion process can be written in terms of the noise-free image $\boldsymbol{x_0}$ as:

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1-\bar{\alpha}_t)\mathbf{I})$$

$$\text{where} \quad \bar{\alpha}_t = \prod_{i=0}^{t} 1 - \beta_t$$

- **For a noisy image with known noise variance β^*, we inject it into the reverse diffusion process at timestep t^* such that $\beta^* \approx 1 - \bar{\alpha}_{t^*}$**
- **The resulting noise-free image is then used in the z-update of ADMM as a denoising prior.**

Experimental Results

- DM acts as an excellent denoising prior when the image belongs to the class it was trained on. On a different class, only works with low noise.
- PSNR is not reflective of visual quality, since the DM hallucinates details

References

[1] Boyd, Stephen, et al. "Distributed optimization and statistical learning via the alternating direction method of multipliers." Foundations and Trends® in Machine learning 3.1 (2011): 1-122

[2] Venkatakrishnan, Singanallur V., Charles A. Bouman, and Brendt Wohlberg. "Plug-and-play priors for model based reconstruction." 2013 IEEE Global Conference on Signal and Information Processing. IEEE, 2013

[3] Ho, Jonathan, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models." Advances in Neural Information Processing Systems 33 (2020): 6840-6851