Plug-and-Play ADMM and Forward-Backward Splitting

Ernest K. Ryu[1] Jialin Liu[1] Sicheng Wang[2] Xiaohan Chen[2]
Zhangyang Wang[2] Wotao Yin[1]

June 12, 2019
International Conference on Machine Learning, Long Beach, CA

[1]UCLA Mathematics
[2]Texas A&M Computer Science and Engineering

Image processing via optimization

Classical variational methods in image processing solve

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad f(x) + \gamma g(x),$$

with methods like ADMM

$$\begin{aligned} x^{k+1} &= \mathrm{Prox}_{\alpha\gamma g}(y^k - u^k) \\ y^{k+1} &= \mathrm{Prox}_{\alpha f}(x^{k+1} + u^k) \\ u^{k+1} &= u^k + x^{k+1} - y^{k+1}. \end{aligned}$$

$$\mathrm{Prox}_{\alpha h}(z) = \mathrm{argmin}_{x \in \mathbb{R}^d} \left\{ \alpha h(x) + (1/2)\|x - z\|^2 \right\}.$$

Plug-and-play image reconstruction

Plug-and-play (PnP) ADMM is a recent non-convex image reconstruction technique for regularizing inverse problems by using advanced denoisers within an iterative algorithm:

$$
\begin{aligned}
x^{k+1} &= H(y^k - u^k) \\
y^{k+1} &= \mathrm{Prox}_f(x^{k+1} + u^k) \\
u^{k+1} &= u^k + x^{k+1} - y^{k+1}.
\end{aligned}
$$

f measures data fidelity and H is a denoiser

$$H : \text{noisy image} \mapsto \text{less noisy image}.$$

Empirically, PnP produces very accurate (clean) reconstructions when it converges. However, there were no theoretical convergence guarantees.

Plug-and-play image reconstruction

We provide the first general convergence analysis of PnP-ADMM.

Theorem

Assume the denoiser H satisfies

$$\|(H-I)(x)-(H-I)(y)\| \leq \varepsilon\|x-y\|, \qquad \forall x,y \tag{A}$$

for some $\varepsilon \geq 0$. Assume f is μ-strongly convex and differentiable. Then PnP-ADMM is a contractive fixed-point iteration and thereby converges in the sense that x^k converges to a fixed point $x^\star$.

(A) means $(H-I)$, the noise estimator, is Lipschitz continuous in the image. We can practically enforce this assumption.

Deep learning denoiser

State-of-the-art denoisers like DnCNN[3] are trained neural networks.

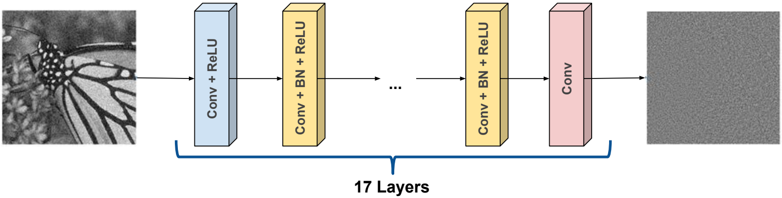

Given a noisy observation $y = x + e$, where x is the clean image and e is noise, the residual mapping R outputs the noise. Learning the residual mapping is a common approach in deep learning-based image restoration.

[3] Zhang, Zuo, Chen, Meng, and Zhang, Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising, IEEE TIP, 2017.

Real Spectral normalization

Enforcing

$$\|(I-H)(x)-(I-H)(y)\| \leq \varepsilon\|x-y\| \tag{A}$$

is equivalent to constraining the Lipschitz constant of R.

For this, we propose Real Spectral Normalization (realSN), a variation of Spectral Normalization of Miyato et al. [4]

RealSN is an approximate projected gradient method enforcing the Lipschitz continuity constraint through a power iteration.

[4] Miyato, Kataoka, Koyama, and Yoshida, Spectral Normalization for Generative Adversarial Networks, ICLR, 2018.

Conclusion

Previously, PnP would produce accurate image reconstructions when it converges, but it would not always converge. Our theory explains when and why PnP converges.

By training the denoiser with realSN, we make PnP converge reliably and thereby make its image reconstruction more reliable.

Longer version of this talk (21.5 minutes) is available on YouTube.

https://youtu.be/V3mbNG5WHPc

Or search in Google:
"Plug-and-Play methods provably converge YouTube"