Uncertainty-Driven Adaptive Sampling via GANs

Thomas Sanchez, Igor Krawczuk, Zhaodong Sun and Volkan Cevher
Laboratory for Information and Inference Systems (LIONS), EPFL, Switzerland

Motivation

Given limited measurements about some unknown ground truth $\boldsymbol{x}$, how can we select a measurement that would give us the most information about it?

Proposed solution | Estimate the posterior with a GAN and use its variance to drive sampling adaptively.

- GANs can model the posterior for inverse problems [1, 2].
- The variance of the posterior provides a natural criterion to guide sampling adaptively.
- The posterior quickly contracts to the correct mode.

Linear inverse problem

Recover a signal $\boldsymbol{x} \in \mathbb{C}^P$ from partial observations $\boldsymbol{y}_\omega$ obtained by subsampling a unitary transform matrix $\boldsymbol{A}$:

$$\boldsymbol{y}_\omega = \boldsymbol{P}_\omega \boldsymbol{A}\boldsymbol{x} + \boldsymbol{\eta} \in \mathbb{C}^N \quad (1)$$

- $\boldsymbol{\eta}$: noise vector;
- $\omega \subseteq [P]$: index set of sampling locations with $|\omega| = N$;
- $\boldsymbol{P}_\omega$: diagonal matrix s.t. $(\boldsymbol{P}_\omega)_{ii} = \mathbb{1}_{i\in\omega}$.

Posterior estimation

Draw samples of the (approximate) posterior distribution:

$$\hat{\mathbf{x}}_\theta = G_\theta(\boldsymbol{y}_\omega, \mathbf{z}), \quad (2)$$

where $\mathbf{z}$ is a random vector drawn from a simple distribution; $G_\theta(\boldsymbol{y}_\omega, \mathbf{z})$ is a deep generative model mapping the input to a sample from $p(\mathbf{x}|\boldsymbol{y}_\omega)$. Such models typically rely on conditional GANs [2] or VAEs [3].

Sampling optimization

Ideal sampling optimization algorithm: tailor the mask to each instance of $\mathbf{x} \sim p(\mathbf{x})$ solving

$$\underset{\omega:|\omega|\le N}{\operatorname{argmin}}\ \ell(\mathbf{x}, \hat{\boldsymbol{x}}_\theta(\mathbf{y}_\omega = \mathbf{P}_\omega \boldsymbol{A}\mathbf{x})). \quad (3)$$

Solving this directly is impossible since this requires using the unknown ground truth signal $\mathbf{x}$ at testing time. Two main approaches circumvent this problem.

Sampling optimization (cont'd)

Fixed (open-loop) sampling. Design a mask at training time and keep it fixed at inference time.

$$\underset{\omega:|\omega|\le N}{\operatorname{argmin}}\ \mathbb{E}_{\mathbf{x}\sim p(\mathbf{x})}\left[\ell(\mathbf{x}, \hat{\boldsymbol{x}}_\theta(\mathbf{y}_\omega = \boldsymbol{P}_\omega \boldsymbol{A}\mathbf{x}))\right], \quad (4)$$

- constrained to a maximal sampling budget N
- want to find a mask that minimizes the loss ℓ
- estimate the expected loss using the empirical estimate with the training samples: $\hat{\boldsymbol{x}}_\theta(\mathbf{y}_\omega = \boldsymbol{P}_\omega \boldsymbol{A}\mathbf{x})$ denotes an estimate of the mean, obtained by averaging on $\mathbf{z}$ for equation 2, i.e., $\hat{\boldsymbol{x}}_\theta(\mathbf{y}_\omega) = \mathbb{E}_{\mathbf{z}}[G_\theta(\boldsymbol{y}_\omega, \mathbf{z})]$.

Adaptive (closed-loop) sampling. Adapt at test time to a fixed, unknown data sample $\boldsymbol{x}$

$$v_t = \underset{v:v\in[P]}{\operatorname{argmin}}\ \ell(\boldsymbol{x}, \hat{\boldsymbol{x}}_\theta(\boldsymbol{y}_{\omega_t})) \text{ s.t } \omega_t = \{\omega_{t-1}, v\} \quad (5)$$

- v_t: individual pixels to be sequentially observed from $\boldsymbol{x} \sim p(\mathbf{x})$.
- Use information of the previously obtained measurements $\boldsymbol{y}_{\omega_{t-1}}$ to determine what should be acquired at time t by conditioning on previous measurements.
- Optimize a heuristic ℓ available at test time, or train a surrogate model to guide sampling [4, 7].

GAS: Generative adaptive sampling

Starting at $t = 0$ we iteratively sample the locations v_t with the largest variance in the domain $\boldsymbol{A}$:

$$v_t = \underset{v:v\in[P]}{\operatorname{argmax}}\ \mathbf{P}_{\{v\}}\mathrm{Var}[\boldsymbol{A}\mathbf{x}|\mathbf{y}_{\omega_t}] \quad (6)$$

- $\mathrm{Var}[\boldsymbol{A}\mathbf{x}|\boldsymbol{y}_{\omega_t}]$: element-wise variance in $\boldsymbol{A}\mathbf{x}$.
- No need for a surrogate model: the variance from the posterior model is sufficient to guide sampling adaptively.

Results

Results using a cascade of residual networks (c-ResNet) [7] for the generator, comparing three sampling approaches:

1. Variable density sampling (VDS) [5]
2. Learning-based Compressive Sensing (LBC) [6]
3. Our generative adaptive sampling methods (GAS)

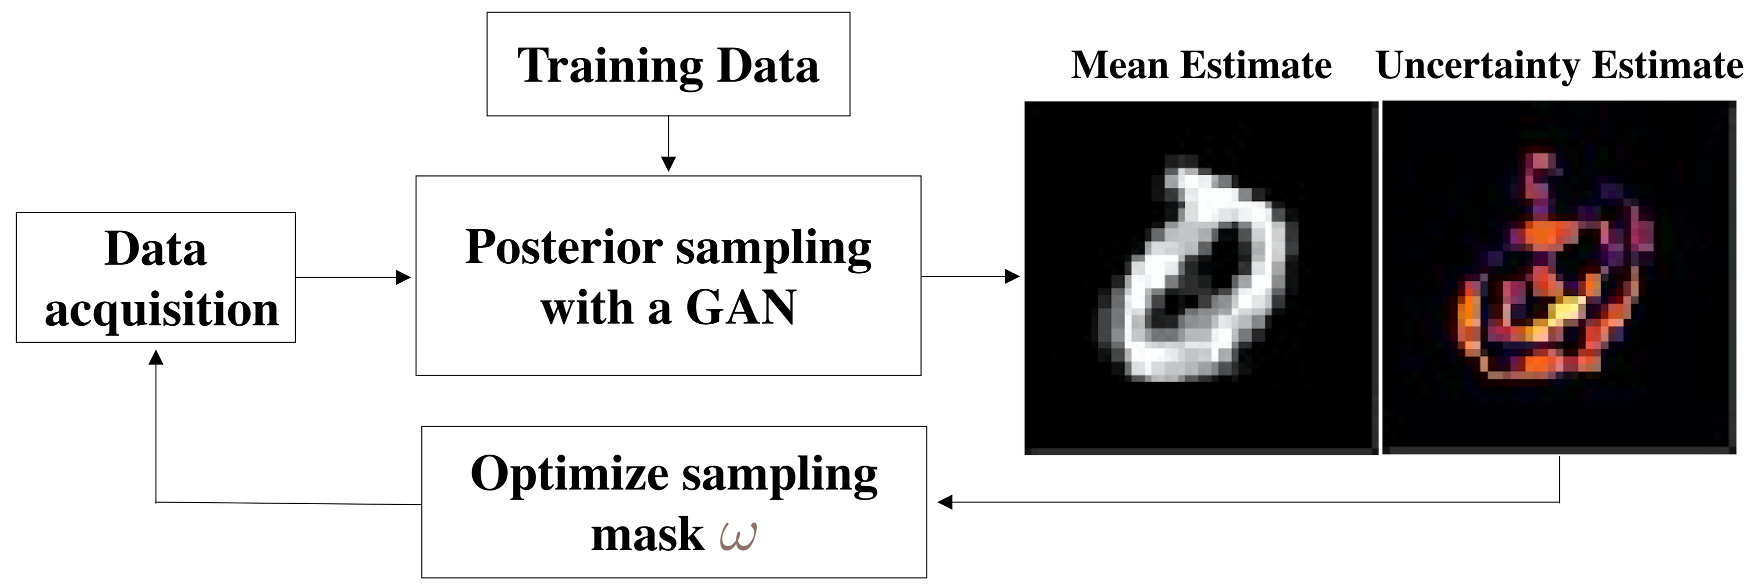

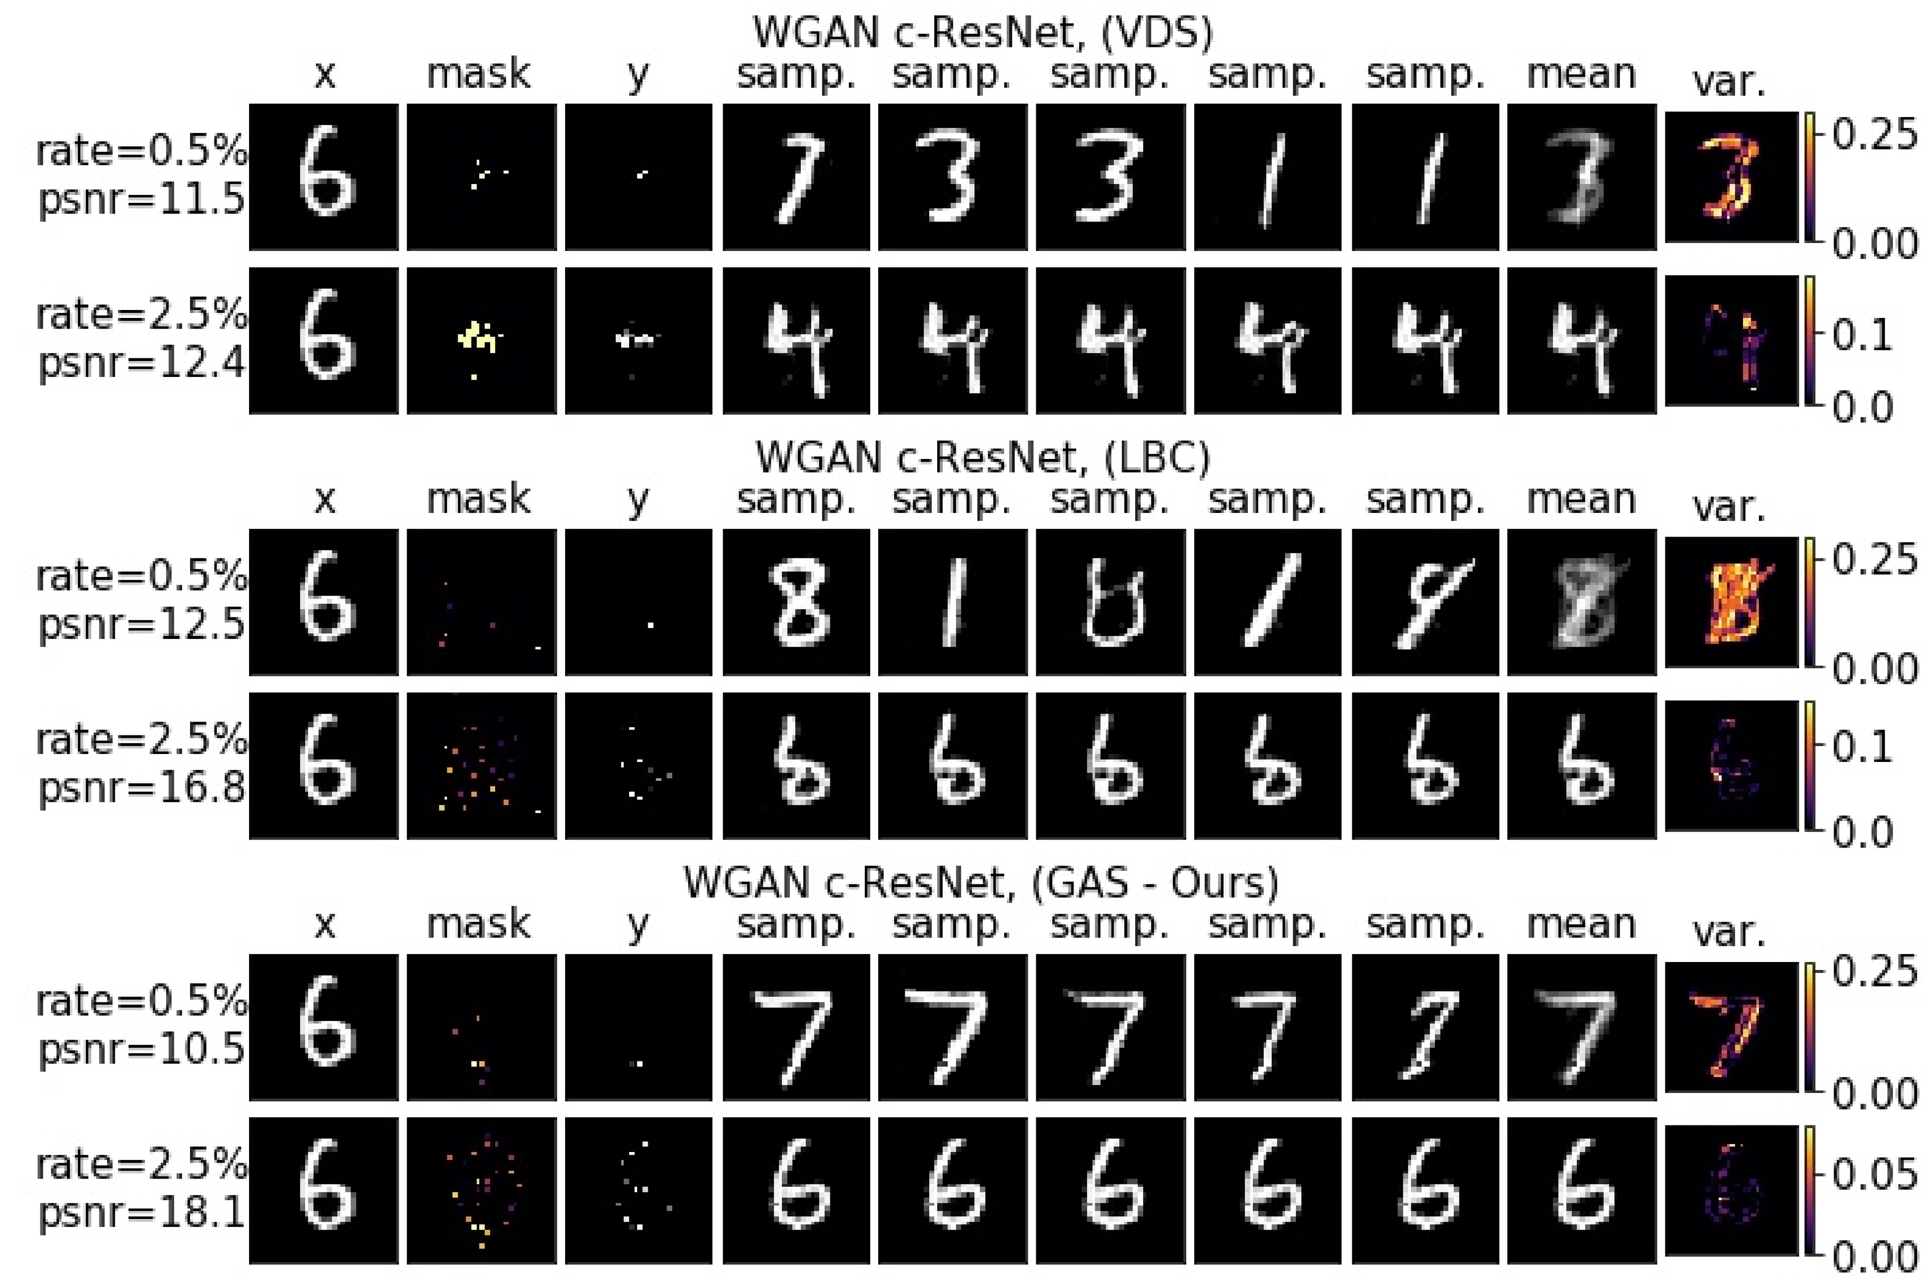

Fig. 2: Comparison of three sampling approaches on some test image from MNIST.

References

[1] J. Adler and O. Öktem. Deep bayesian inversion. *arXiv:1811.05910*, 2018.
[2] M. Belghazi, M. Oquab, and D. Lopez-Paz. Learning about an exponential amount of conditional distributions. In *NeurIPS*, 2019.
[3] O. Ivanov, M. Figurnov, and D. Vetrov. Variational autoencoder with arbitrary conditioning. In *ICLR*, 2019.
[4] K. H. Jin, M. Unser, and K. M. Yi. Self-supervised deep active accelerated MRI. *arXiv:1901.04547*, 2019.
[5] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.
[6] T. Sanchez, B. Gözcü, R. B. van Heeswijk, A. Eftekhari, E. Ilıcak, T. Çukur, and V. Cevher. Scalable learning-based sampling optimization for compressive dynamic mri. In *ICASSP*, 2020.
[7] Z. Zhang, A. Romero, M. J. Muckley, P. Vincent, L. Yang, and M. Drozdzal. Reducing uncertainty in undersampled MRI reconstruction with active acquisition. In *CVPR*, 2019.