Multiplicative Denoising with Uncertainty Quantification for Synthetic Aperture Radar

Jonathan Lindbloom* Anne Gelb Matthew Parno

Dartmouth College

April 14, 2022

SIAM UQ22

This work is supported by ONR MURI #N00014-20-1-2595.

Motivation

- SAR is operational irrespective of weather and daylight, making it an excellent source of information to complement optical imagery
- SAR scenes routinely captured by satellites such as Sentinel-1

Sentinel-1, courtesy of ESA

Speckle Noise

- Coherent imaging systems such as synthetic aperture radar (SAR) suffer from multiplicative *speckle noise*

$$\boldsymbol{f} = K\boldsymbol{u} \odot \boldsymbol{v}$$

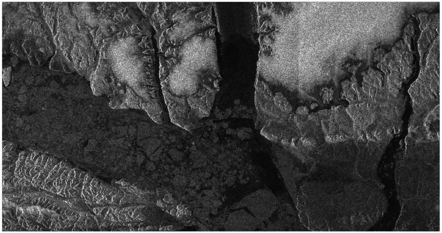

Copernicus Sentinel data 2016, processed by ESA.

- For single-look complex (SLC) SAR data, we can apply the *product model* with the simple forward operator $K = I$
- **Goal:** recover the latent radar cross-section (RCS) with uncertainties

Bayesian Formulation

- We adopt a Bayesian approach to recovering the RCS

$$\underbrace{\pi(\boldsymbol{u} \mid \boldsymbol{f})}_{\text{posterior}} \propto \overbrace{\pi(\boldsymbol{f} \mid \boldsymbol{u})}^{\text{likelihood}} \times \underbrace{\pi(\boldsymbol{u})}_{\text{prior}}$$

- Physical considerations motivate the assumption that the speckle noise is *i.i.d.* Gamma distributed, where the parameters depend on the number of looks L used to form the observed image $\boldsymbol{f}$

MAP Estimate

$$\boldsymbol{u}^\star = \arg\max_{\boldsymbol{u}} \; L\left(\sum_{i,j} \log u_{i,j} + \frac{f_{i,j}}{u_{i,j}}\right) - \log \pi(\boldsymbol{u})$$

Markov Random Field (MRF) Priors

- We may place a prior on the recovery directly, or perhaps some sparsifying transform A of the reconstruction with

$$\pi(\boldsymbol{u}) \propto e^{-\Phi(A\boldsymbol{u})}$$

- Our posterior inherits the conditional independence structure of our prior, since the noise model is applied pixel-wise. For example, the posterior under anisotropic TV satisfies

$$\pi(u_{i,j} \mid u_{\neg i,j}, \boldsymbol{f}) = \pi(u_{i,j} \mid u_{i-1,j}, u_{i+1,j}, u_{i,j-1}, u_{i,j+1}, f_{i,j})$$

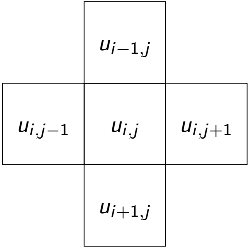

Gibbs Sampling

- To characterize the posterior $\pi(\boldsymbol{u} \mid \boldsymbol{f})$ we employ Markov chain Monte Carlo (MCMC), specifically *Metropolis-within-Gibbs (MwG) sampling*

Metropolis-within-Gibbs Sampling (Gaussian Proposal)

For each iteration $k = 1, \ldots, n_{\text{iterations}}$:

For each pixel $u = u_{p_1}, u_{p_2}, \ldots, u_{p_{mn}}$:

1. Propose a new pixel value $u' \sim \mathcal{N}(u, \sigma_u^2)$, new image $\boldsymbol{u}'$
2. Compute the acceptance ratio

$$\alpha = \min\left\{1, \frac{\pi(\boldsymbol{u}' \mid \boldsymbol{f})}{\pi(\boldsymbol{u} \mid \boldsymbol{f})}\right\} \tag{1}$$

3. Accept $\boldsymbol{u}'$ w.p. α, otherwise reject and the chain stays at $\boldsymbol{u}$

- The pixel-dependent σ_u^2 selected via the log-adaptation procedure of [Roberts & Rosenthal, 2009]

Naive MwG has Complexity $\mathcal{O}(mn)$

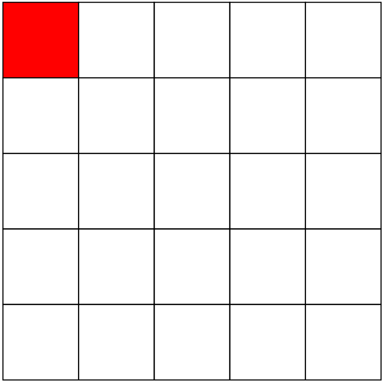

$$\pi(u_{i,j} \mid u_{\neg i,j}, \boldsymbol{f}) = \pi(u_{i,j} \mid u_{i-1,j}, u_{i+1,j}, u_{i,j-1}, u_{i,j+1}, f_{i,j})$$

Naive MwG has Complexity $\mathcal{O}(mn)$

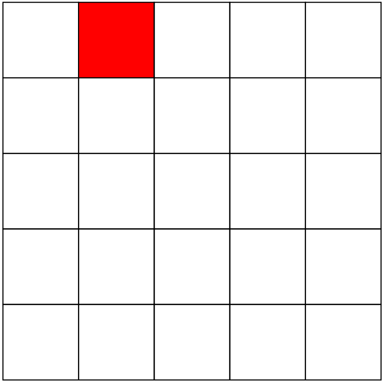

$$\pi(u_{i,j} \mid u_{\neg i,j}, \boldsymbol{f}) = \pi(u_{i,j} \mid u_{i-1,j}, u_{i+1,j}, u_{i,j-1}, u_{i,j+1}, f_{i,j})$$

Naive MwG has Complexity $\mathcal{O}(mn)$

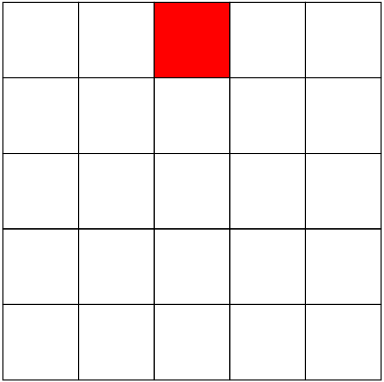

$$\pi(u_{i,j} \mid u_{\neg i,j}, \boldsymbol{f}) = \pi(u_{i,j} \mid u_{i-1,j}, u_{i+1,j}, u_{i,j-1}, u_{i,j+1}, f_{i,j})$$

Chromatic Gibbs Sampling

- Considering the conditional independence graph of the posterior, we can define a coloring and apply the MwG updates within each color simultaneously [Geman & Geman, 1984], [Brown et al., 2021]
- True complexity is $\mathcal{O}(\chi)$, independent of dimension
- We can leverage GPUs to carry out the updates using stencils implemented as CUDA kernels

Anisotropic TV MwG Updates, $\chi = 2$

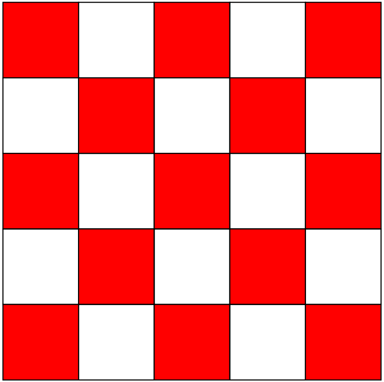

$$\pi(u_{i,j} \mid u_{\neg i,j}, \boldsymbol{f}) = \pi(u_{i,j} \mid u_{i-1,j}, u_{i+1,j}, u_{i,j-1}, u_{i,j+1}, f_{i,j})$$

Anisotropic TV MwG Updates, $\chi = 2$

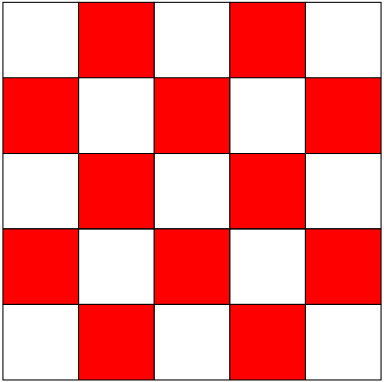

$$\pi(u_{i,j} \mid u_{\neg i,j}, \boldsymbol{f}) = \pi(u_{i,j} \mid u_{i-1,j}, u_{i+1,j}, u_{i,j-1}, u_{i,j+1}, f_{i,j})$$

Toy Problem – Anisotropic TV, log-scale

(a) Ground truth

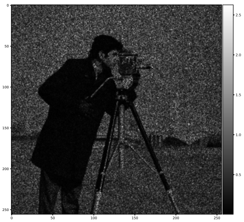

(b) Noisy image

(c) MAP point estimate

(d) Posterior mean

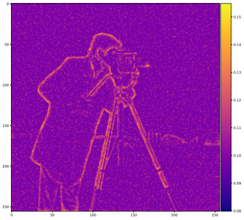

(e) Standard deviation

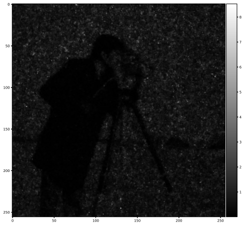

(f) Posterior sample

TV Prior, 1D Speckled Signal

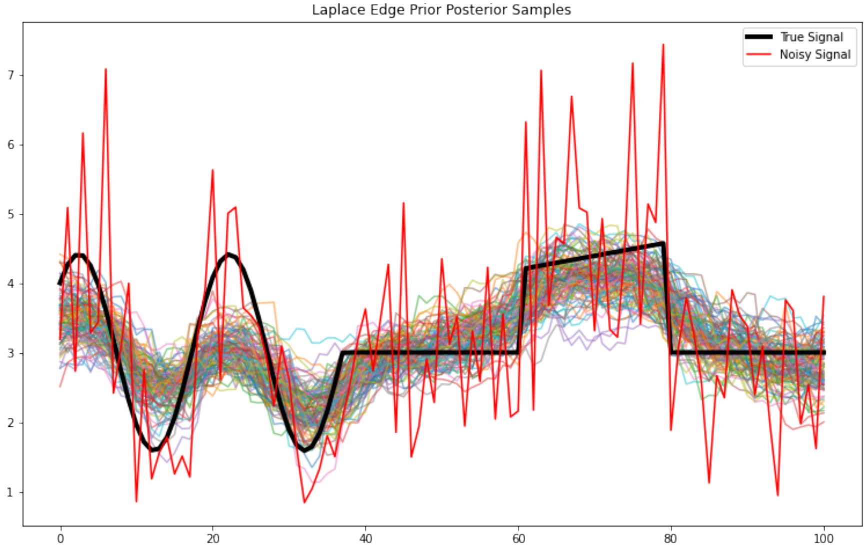

Cauchy Difference Prior, 1D Speckled Signal

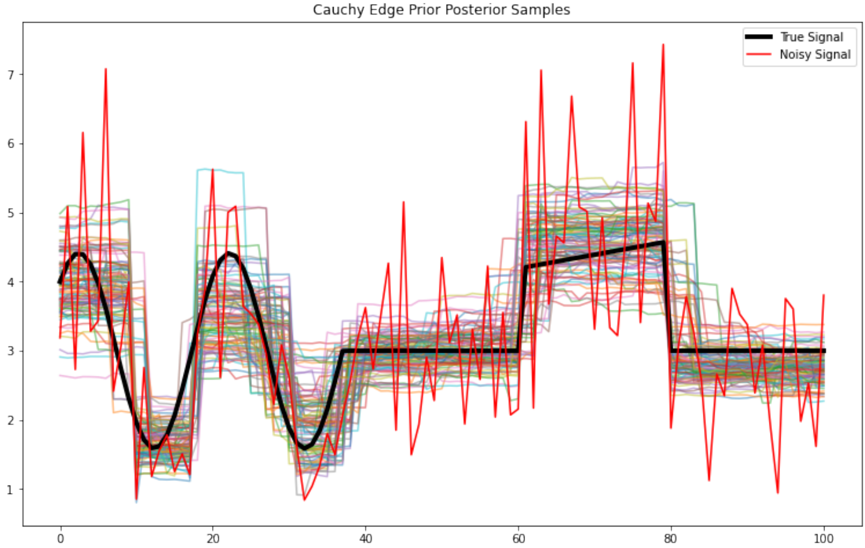

Performance

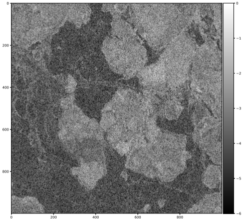

(a) SAR image (7.7)

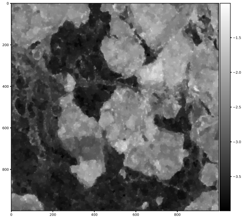

(b) TV MAP (43.2)

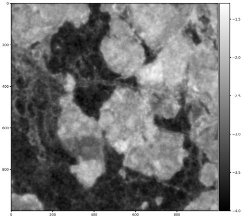

(c) Cauchy MAP (47.0)

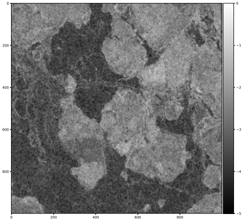

(d) Refined Lee (25.3)

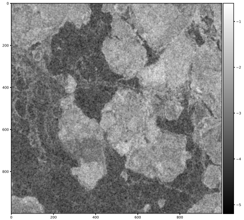

(e) TV mean (28.7)

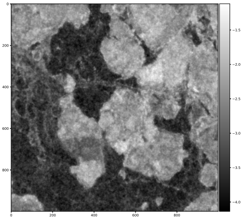

(f) Cauchy mean (42.3)

Performance - TV

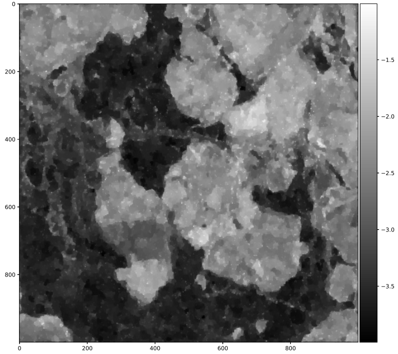

(a) TV MAP (43.2)

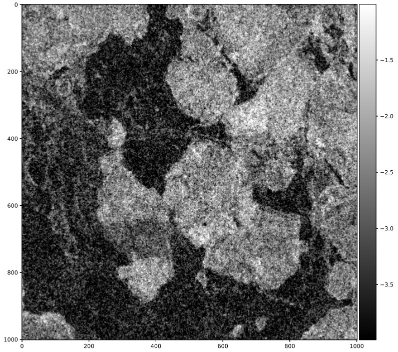

(b) TV sample (14.5)

Performance - Cauchy Difference

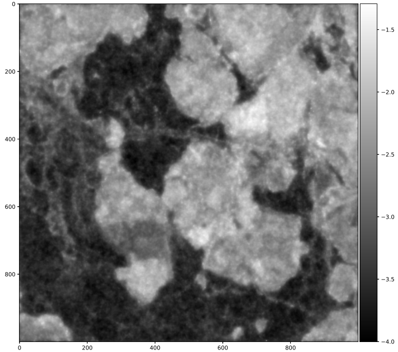

(a) Cauchy MAP (47.0)

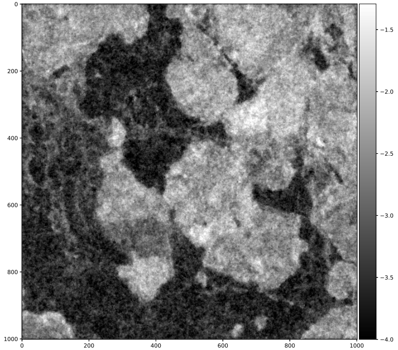

(b) Cauchy sample (34.6)

Scales to Large Scenes

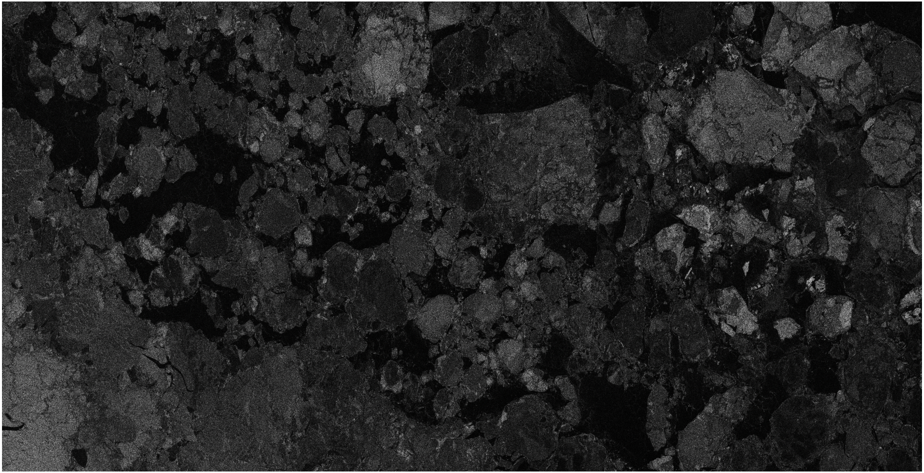

4058 x 7952 pixels ($\approx 3.2 \times 10^7$); 6.1×15.2 km

UQ – TV

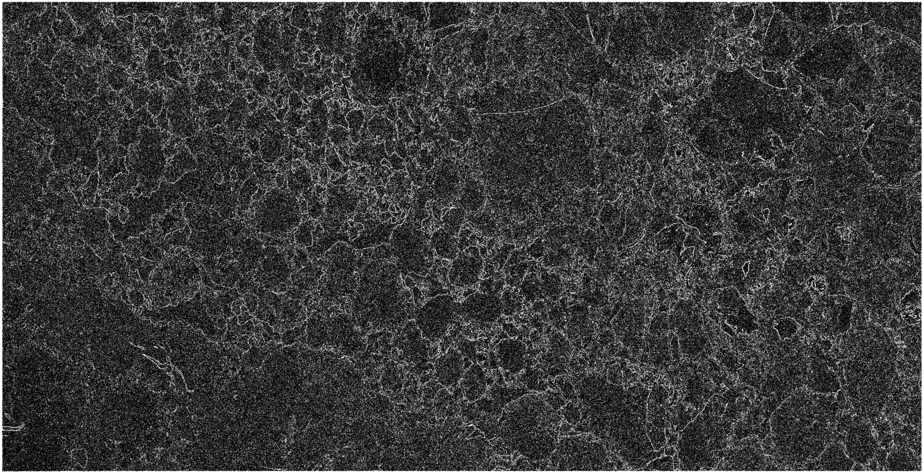

4058 x 7952 pixels ($\approx 3.2 \times 10^7$); 6.1×15.2 km

UQ – Cauchy

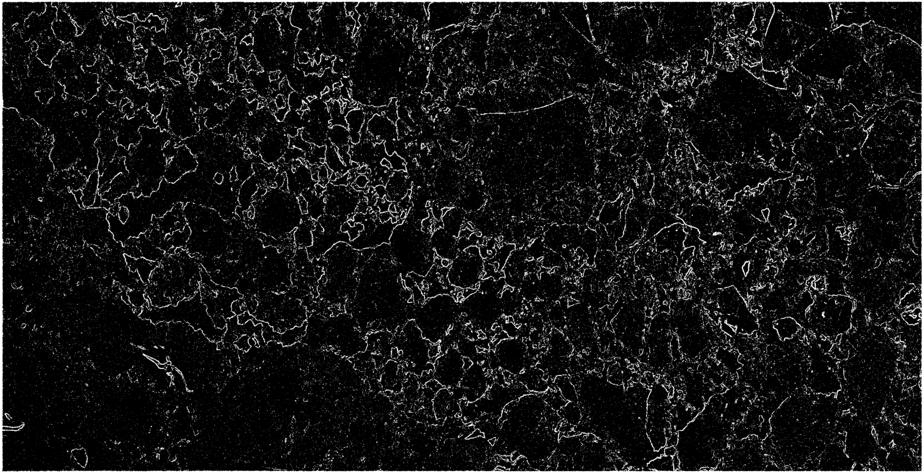

4058 x 7952 pixels ($\approx 3.2 \times 10^7$); 6.1×15.2 km

Current Work

- Auxiliary variable sampling methods for tackling multi-modality
- Making best use of available hardware
- Downstream applications for samples
- Quantitative evaluation of despeckling performance