Score-based Source Separation with Applications to Digital Communication Signals

Tejas Jayashankar[1], Gary C.F. Lee[1], Alejandro Lancho[1, 2], Amir Weiss[1], Yury Polyanskiy[1], Gregory Wornell[1]

[1]Massachusetts Institute of Technology (MIT) [2]Universidad Carlos III de Madrid

PROJECT WEBPAGE

Single-channel Source Separation

Goal: Separate a mixture of two signals into its constituent components.

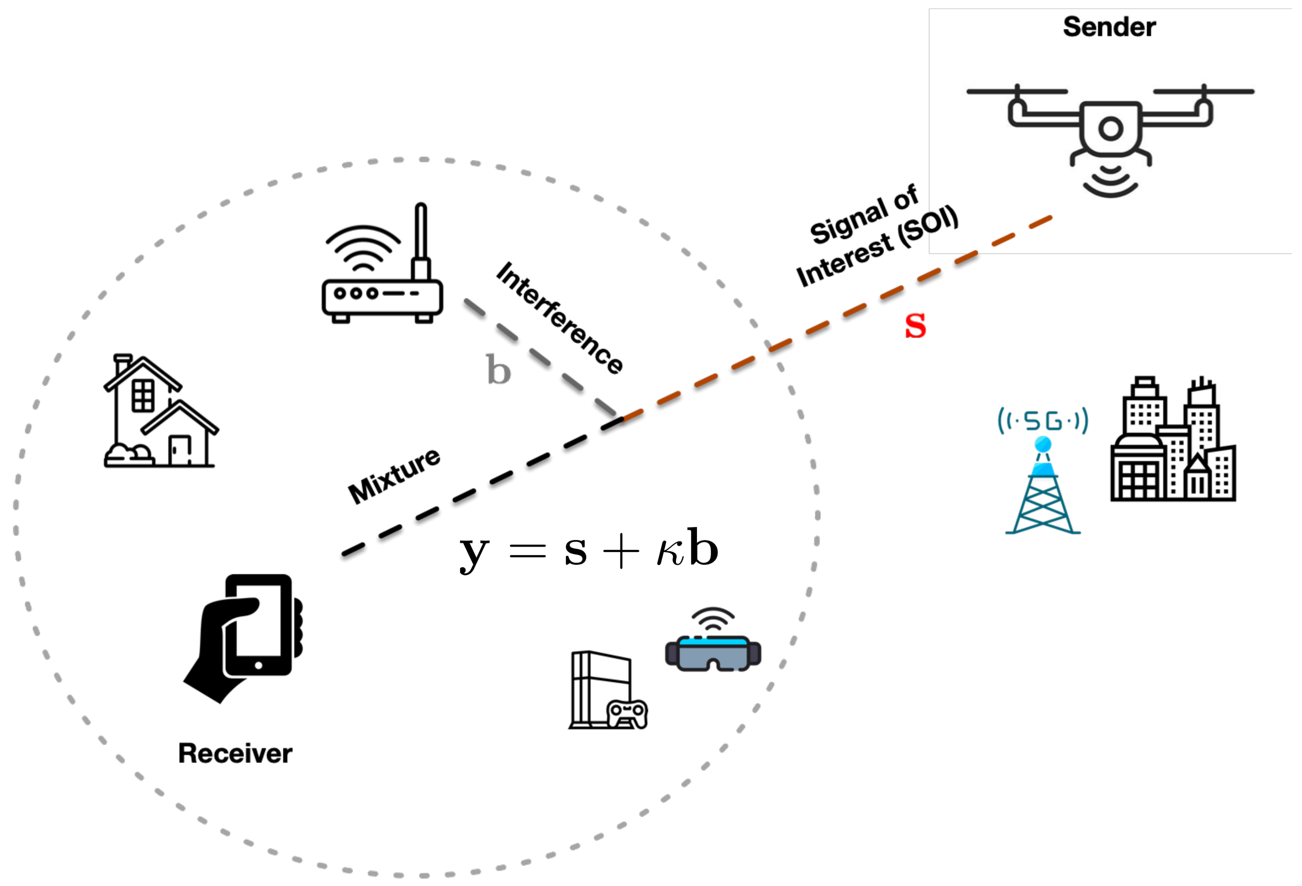

Motivation

- **Scalability:** Develop a library of pre-trained models for different wireless signals.
- **Automation:** Develop methods that can adapt to new wireless standards with unknown system parameters.

Digital Communication Signals

Signal Generation Model for digital radio-frequency (RF) signals:

$$s(t) = \sum_{p=-\infty}^{\infty} \sum_{\ell=0}^{L-1} c_{p,\ell}\, g(t - pT_s, \ell) \exp\{j2\pi \ell t / L\}$$

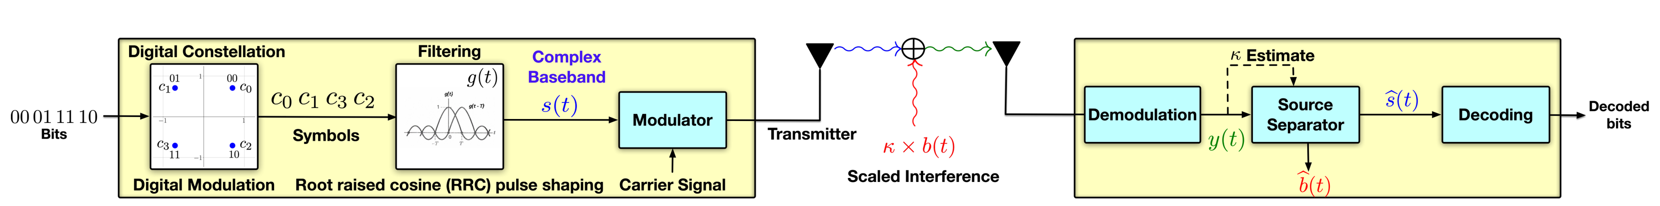

Challenges with RF Source Separation

- **Underlying discreteness** arising from discrete symbols modulated onto a continuous waveform.
- **Overlapping temporal structures in time and frequency.**
- SOI signal model is possibly unknown and **non-Gaussian interference.**

α-RGS: A new Bayesian-inspired Method

- Uses pre-trained diffusion models as plug and play priors inspired by α-posterior Generalized Bayes' theorem.
- Re-uses the training noise schedule without extensive hyperparameter fine-tuning. Only need to tune the learning rate.
- Uses randomized levels of Gaussian smoothing.

Backbone: Maximum a Posteriori (MAP) Estimation

$$\widehat{\mathbf{s}} = \underset{\mathbf{s}\in\mathcal{S} \text{ s.t } \mathbf{y}=\mathbf{s}+\kappa\mathbf{b}}{\arg\max}\; p_{\mathsf{s}|\mathsf{y}}(\mathbf{s}|\mathbf{y}),$$
$$= \underset{\mathbf{s}\in\mathcal{S}}{\arg\min} -\log P_{\mathsf{s}}(\mathbf{s}) - \log p_{\mathsf{b}}\left((\mathbf{y}-\mathbf{s})/\kappa\right)$$

Challenges:
- Combinatorial Problem
- Non-differentiable

α-posterior with Randomized Gaussian Smoothing

Our solution: Gradient-based algorithm that asymptotically approaches the modes of $P_{\mathsf{s}}(\mathbf{s})$ by estimating $\bar{\mathbf{s}} = \mathbf{s} + \epsilon, \epsilon \to 0$.

Gaussian Smoothing

Let $\{1-\alpha_t\}_{t=0}^{T}$ be increasing noise levels in (0, 1) and let $\mathbf{z}_s, \mathbf{z}_b \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

$$\tilde{\mathbf{s}}_t(\bar{\mathbf{s}}) \triangleq \sqrt{\alpha_t}\bar{\mathbf{s}} + \sqrt{1-\alpha_t}\mathbf{z}_s,$$
$$\tilde{\mathbf{b}}_u(\bar{\mathbf{s}}, \mathbf{y}) \triangleq \sqrt{\alpha_u}\left(\mathbf{y} - \bar{\mathbf{s}}\right)/\kappa + \sqrt{1-\alpha_u}\mathbf{z}_b$$

Generalized Bayes' with an α-posterior

Idea: Use a tempered likelihood to reweigh the contribution of the more complicated distribution.

$$\underbrace{p_{\mathsf{s}|\mathsf{y}}(\mathbf{s}|\mathbf{y};\omega)}_{\alpha\text{-posterior } (\alpha=\omega)} \propto \underbrace{p_{\mathsf{y}|\mathsf{s}}(\mathbf{y}|\mathbf{s})^{\omega}}_{\text{tempered likelihood}} \underbrace{P_{\mathsf{s}}(\mathbf{s})}_{\text{prior}}$$

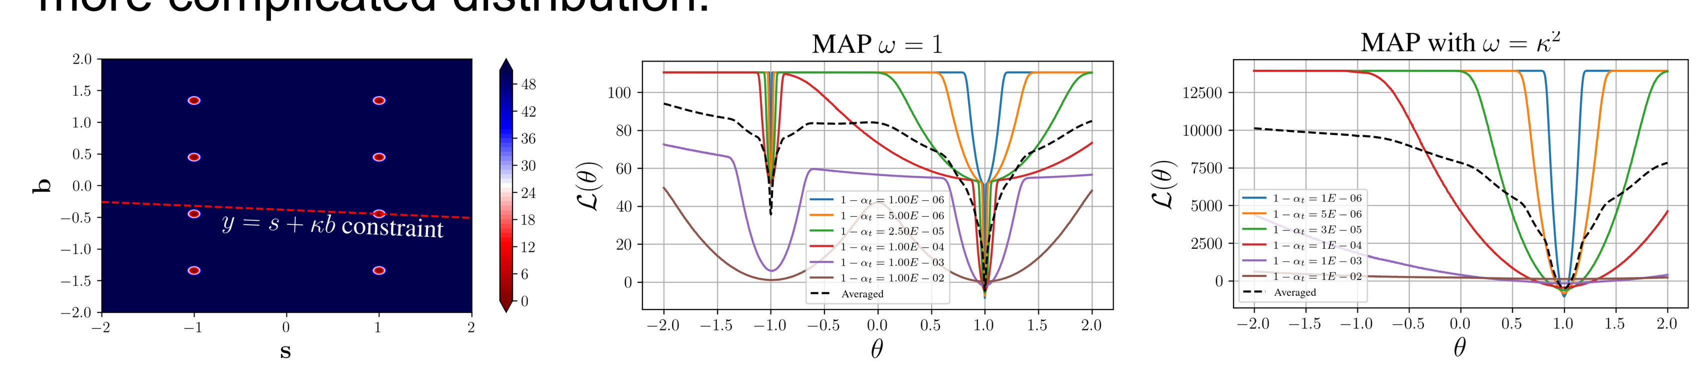

Visualization of problem as a 2D constraint

Asymptotic Loss Function

$$\mathcal{L}(\boldsymbol{\theta}) \triangleq -\mathbb{E}_{\mathsf{t},\mathsf{z}_s}\left[\log p_{\tilde{\mathsf{s}}_t}\left(\tilde{\mathbf{s}}_t(\boldsymbol{\theta})\right)\right] - \omega\mathbb{E}_{\mathsf{u},\mathsf{z}_u}\left[\log p_{\tilde{\mathsf{b}}_u}\left(\tilde{\mathbf{b}}_u(\boldsymbol{\theta},\mathbf{y})\right)\right]$$

Multiple Noise Level Update Rule

$$\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta}) \triangleq -\mathbb{E}_{\mathsf{t},\mathsf{z}_t}\left[\sqrt{\alpha_t}\, S_{\tilde{\mathsf{s}}_t}\left(\tilde{\mathbf{s}}_t(\boldsymbol{\theta})\right)\right] + \frac{\omega}{\kappa}\mathbb{E}_{\mathsf{u},\mathsf{z}_u}\left[\sqrt{\alpha_u} S_{\tilde{\mathsf{b}}_u}\left(\tilde{\mathbf{b}}_u(\boldsymbol{\theta},\mathbf{y})\right)\right]$$

- **Score from diffusion model:** $S_{\tilde{\mathsf{s}}_t}\left(\tilde{\mathbf{s}}_t(\boldsymbol{\theta})\right) \triangleq \nabla_{\mathbf{x}} \log p_{\tilde{\mathsf{s}}_t}(\mathbf{x})\Big|_{\mathbf{x}=\tilde{\mathbf{s}}_t(\boldsymbol{\theta})}$
- **Large noise levels:** Explore between modes.
- **Small noise levels:** Resolve the solution.

Experiments and Results

- **Pre-training:** Trained independent diffusion models on signals.
- **Experiment:** Separated mixtures using α-RGS with $\omega = \kappa^2$.
- **Primary Baselines:** 1) Matched Filtering (MF): optimal under Gaussian interference, 2) LMMSE: optimal when source and interference are Gaussian and 3) BASIS: Annealed Langevin dynamics with specially tuned noise schedule.
- **Metrics:** Bit error rate (BER) and mean squared error (MSE).

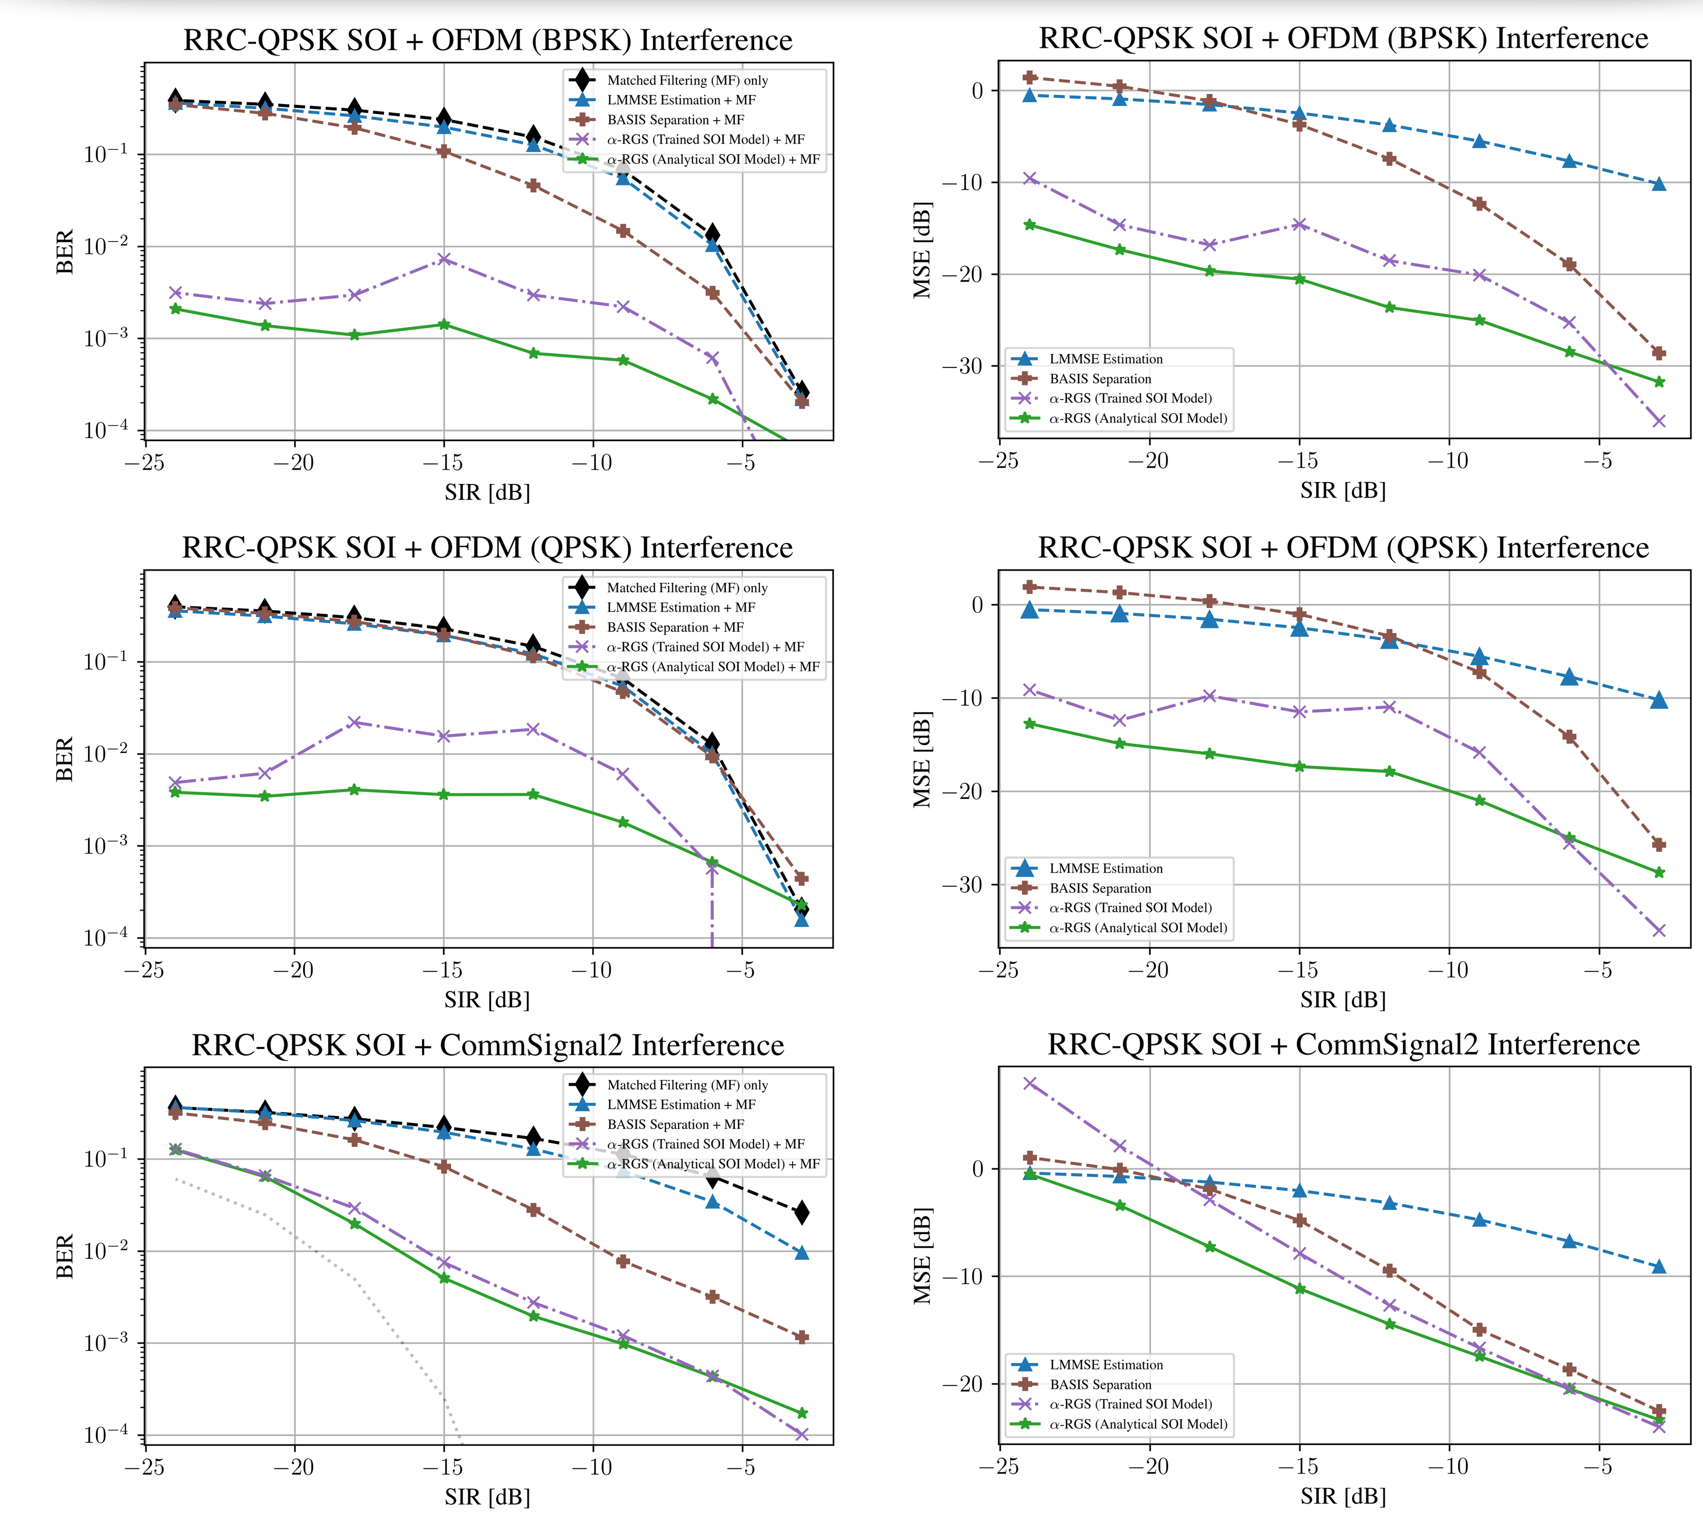

α-RGS achieves 95% reduction in BER over traditional and learning-based baselines.

References

[1] MIT RLE. Single-Channel RF Challenge. URL https://rfchallenge.mit.edu/challenge-1/. Accessed 2023-11-24.

[2] V. Jayaram and J. Thickstun. Source separation with deep generative priors. In International Conference on Machine Learning, pages 4724–4735. PMLR, 2020.

Research was sponsored by the United States Air Force Research Laboratory and the United States Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. Research was sponsored by the United States Air Force Research Laboratory and the United States Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.