Accelerated Blind Denoising of GPR Data via Deep Random Projections

Viven Sharma
Department of Earth and Environmental Sciences
University of Minnesota
Minneapolis, MN, USA
sharm760@umn.edu

Raunak S. Manekar
Department of Computer Science and Information Systems
BITS Pilani, K. K. Birla Goa Campus
Goa, India
raunakm@goa.bits-pilani.ac.in

Abstract

Ground Penetrating Radar (GPR) represents a critical non-invasive technology for subsurface imaging across diverse applications, including civil infrastructure assessment, archaeological surveys, and geological exploration. However, GPR data quality is inherently compromised by multiple noise sources, including electromagnetic interference, thermal noise, surface clutter, and system-related artifacts, which collectively obscure important subsurface features and complicate interpretation. This paper presents an application of Deep Random Projections (DRP), a computationally efficient variant of the Deep Image Prior framework, for blind denoising of real-world GPR data where neither noise characteristics nor ground truth clean signals are available. Our approach leverages the implicit regularization provided by convolutional neural network architectures while dramatically reducing computational requirements by freezing network weights and optimizing only a low dimensional input seed and batch normalization parameters. Extensive experiments on synthetic and field-collected GPR data demonstrate that DRP achieves denoising performance comparable or superior to state-of-the-art self-supervised methods, including S2S-WTV, while requiring two orders of magnitude fewer trainable parameters and achieving 5-10× speedup in processing time. We provide quantitative evaluations against high trace count reference images and discuss how the denoised radargrams should be interpreted in terms of subsurface structure. The method's ability to operate without training data or explicit assumptions about noise distributions makes it particularly suitable for practical GPR applications where obtaining clean reference data is infeasible.

1 Introduction

Ground Penetrating Radar (GPR) is a non-invasive geophysical method that uses radar pulses to image the subsurface [1, 19]. It is widely used in civil engineering for utility locating and structural assessment [1, 20, 27], in archaeology for mapping buried features [22, 23], and in environmental and hydrogeological studies for delineating contaminant plumes and hydrological pathways [20, 21]. More generally, GPR has become a standard tool for near-surface characterization and hazard assessment [16, 27], providing high-resolution radargrams (B-scans) by transmitting electromagnetic waves and recording reflections from interfaces with different dielectric properties [17, 19].

Published at the Machine Learning and the Physical Sciences Workshop (NeurIPS 2025).

In practice, GPR data are strongly contaminated by noise from instrument and environmental sources. Internal artifacts such as direct coupling between transmit and receive antennas and instrument ringing [2, 18] combine with external interference from nearby radio transmitters, power lines, and background noise [2, 17, 21]. The resulting noise is heterogeneous and poorly characterized, and clean reference data are rarely available in field settings, so GPR denoising is naturally posed as a blind inverse problem [9, 10, 14]. Classical approaches based on SVD-style filtering and multichannel spectral methods [9, 10, 16] or frequency–wavenumber (F-K) filtering [11] can be effective under specific assumptions on signal and noise, but they often struggle with the complex, non-stationary patterns present in real radargrams.

Recent advances in deep learning have opened new possibilities for image restoration and geophysical signal processing [5, 34]. Deep Image Prior (DIP) leverages the structure of an untrained convolutional network as an implicit regularizer and can achieve strong denoising performance without external training data [14, 28]. However, DIP requires optimizing all network weights over thousands of iterations, which is computationally expensive and prone to overfitting noise despite early stopping heuristics [14, 13, 28, 12]. Self-supervised approaches such as Self2Self [7] and S2S-WTV [4] further reduce reliance on labeled data, and classical methods like BM3D remain strong non-learning baselines [15]. In this work we adopt the Deep Random Projector (DRP) framework [12], which accelerates deep priors by freezing network weights, optimizing only a low-dimensional input seed and batch normalization parameters, and adding an explicit total variation prior. Our goal is to show that DRP provides a practical and scientifically credible solution for blind denoising of real GPR data, narrowing the gap between advanced deep models and field-ready workflows.

2 Methodology: DRP for GPR Denoising

2.1 Problem Formulation

We formulate GPR denoising as a classic inverse problem [32, 33, 34]. Given a noisy GPR B-scan (radargram) denoted as a 2D image $y \in \mathbb{R}^{H \times W}$, the objective is to recover the latent clean image x. We model the degradation process as

$$y = x + n, \tag{1}$$

where n represents the additive noise component whose characteristics are unknown [9, 14]. The task is to estimate x from y without explicit knowledge of n or access to clean reference data.

2.2 The Deep Random Projector Framework

The Deep Random Projector (DRP) method adapts the Deep Image Prior concept to drastically improve computational efficiency [12]. The core idea is to solve an optimization problem of the form

$$\min_{z,\theta_{BN}} \mathcal{L}\big(y, G_\theta(z)\big) + \lambda \mathcal{R}\big(G_\theta(z)\big), \tag{2}$$

where $G_\theta(z)$ is the output of a CNN with weights θ and input seed z, $\mathcal{L}$ is a data fidelity term measuring the similarity between the network output and the noisy observation y, and $\mathcal{R}$ is a regularization term that imposes a prior on the solution.

Optimization target. Unlike DIP, which optimizes the full set of network weights θ, DRP freezes these weights after a random initialization. Instead, the optimization is performed over the much lower-dimensional input seed z and the affine parameters of the batch normalization layers, θ_{BN}. In practice, we initialize z as a random tensor with the same spatial dimensions as the desired output and treat its entries as learnable parameters updated by gradient descent. The batch normalization scale and shift parameters are also updated, while all convolutional kernels remain fixed. This approach circumvents the inhomogeneity in learning problem, the issue of exploding or vanishing gradients across different layers that significantly slows down the convergence of DIP [12]. By shifting the optimization target from tens of thousands of network weights to a compact input representation and a small set of normalization parameters, DRP dramatically reduces the number of trainable parameters (for our U-Net architecture, from roughly 8×10^4 in a typical DIP model to about 6.6×10^2 in DRP), leading to faster convergence and lower memory requirements.

Network architecture. To further reduce the computational cost of each iteration, which is dominated by the forward and backward passes through the network, DRP employs a shallow U-Net style architecture [25]. In our implementation we use two downsampling and upsampling stages, with a reduced number of feature channels compared to the deeper U-Nets commonly used for DIP. A deeper network with more channels generally increases representational capacity but also amplifies the risk of overfitting noise and increases runtime. Our experiments showed that this shallower architecture provides an effective trade-off: it is expressive enough to capture the main GPR reflectivity patterns while significantly reducing the per-iteration cost and stabilizing optimization.

Explicit regularization. The data-fitting term $\mathcal{L}$ is chosen as the mean squared error,

$$\mathcal{L}(y, G_\theta(z)) = \|y - G_\theta(z)\|_2^2, \tag{3}$$

which encourages the output to explain the observed data. Because DRP greatly reduces the implicit regularization normally provided by optimizing a large network, we incorporate an explicit prior based on Total Variation (TV),

$$\mathcal{R}\big(G_\theta(z)\big) = \mathrm{TV}\big(G_\theta(z)\big), \tag{4}$$

which penalizes the sum of the magnitudes of gradients in the image [8]. The TV prior is particularly well-suited for GPR data as it promotes solutions with piecewise-constant regions and preserves sharp edges, which correspond to coherent linear reflectors and hyperbolic diffraction patterns that are of primary interest in geophysical interpretation [9, 33]. The final optimization objective for DRP-based GPR denoising is

$$\min_{z, \theta_{BN}} \|y - G_\theta(z)\|_2^2 + \lambda \mathrm{TV}\big(G_\theta(z)\big). \tag{5}$$

The components of DRP work in a symbiotic relationship. The primary mechanisms for speedup are freezing the weights and using a shallow network, which fundamentally weakens the strong implicit prior that arises from optimizing a deep, over-parameterized network in the original DIP framework [14]. This reduction in implicit regularization could lead to a degradation in performance if used alone. The explicit TV regularizer is therefore not merely an optional enhancement but a crucial component that compensates for this loss. DRP effectively trades the computationally expensive implicit regularization of DIP for a combination of a lightweight generator and an efficient explicit prior, striking a new balance between performance and practicality.

3 Experimental Results

Implementation details. We implemented the DRP framework with a shallow U-Net architecture [25]. Unless otherwise noted, all experiments use the same architecture for DRP, DIP, and S2S-WTV so that differences can be attributed to the optimization scheme rather than network capacity. Optimization was performed using the Adam optimizer with a learning rate of 0.1 [6]. The weight for the TV regularization term, λ, was set to 0.45. These hyperparameters were chosen based on preliminary experiments and kept constant for all tests; small perturbations (±5%) did not noticeably affect performance, suggesting that DRP is not overly sensitive to fine tuning of these values. For DRP we optimized the input seed z and batch normalization parameters for 250–1000 iterations depending on the experiment, while for DIP and S2S-WTV we optimized the full network weights for up to 1000 iterations with early stopping when we observed overfitting.

For context, S2S-WTV [4] is a self-supervised denoising approach that trains a network directly on the noisy data by randomly masking pixels via dropout and reconstructing them from the remaining context. The approach adds a weighted TV term to the loss to preserve structural features in seismic style data. BM3D [15] is a classical denoising algorithm that groups similar patches into 3D stacks and performs collaborative filtering in a transform domain; we use it as a strong non-learning baseline. These baselines allow us to compare DRP not only against DIP-style deep priors but also against established self-supervised and classical methods.

Quantitative validation on synthetic GPR data. We first evaluated DRP on a synthetic GPR B-scan with a known ground truth, allowing for objective comparison using the peak signal-to-noise ratio (PSNR). The synthetic data simulate a subsurface with layered interfaces and diffractors, and we add realistic noise to obtain the noisy observation. We compare our DRP method against the original DIP [14], S2S-WTV [4], and BM3D [15]. All learning-based methods were run for 1000 optimization epochs on this dataset for a fair comparison of final performance.

The results, shown in Figure 1, highlight DRP's superior performance. DRP achieves the highest PSNR of 31.9 dB, representing a significant improvement over S2S-WTV (28.7 dB) and the original DIP (10.7 dB), and outperforming BM3D (17.7 dB). Notably, the standard DIP framework struggles with this task: its output has lower PSNR than the noisy input (approximately 16.5 dB), confirming that the original DIP optimization overfits noise even with early stopping. Visually, DRP's denoised image recovers subsurface reflectors with the least residual noise and most closely matches the ground truth, while DIP, S2S-WTV, and BM3D either leave significant noise or oversmooth important features.

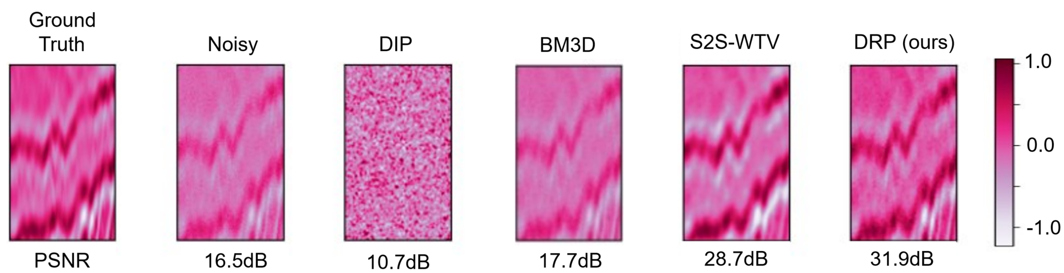

Figure 1: Comparative denoising of synthetic GPR data with 1000 epochs. From left to right: Ground Truth, Noisy Input, DIP, BM3D, S2S-WTV, and DRP (ours). DRP achieves the highest PSNR and visually recovers the signal with the least residual noise.

Blind denoising of field GPR data. To demonstrate DRP's practical utility, we applied it to a real-world field GPR dataset. The noisy input is a pre-stack radargram composed of only 256 stacked traces, which is insufficient to clearly resolve many subsurface features due to low signal-to-noise ratio. As a high quality reference, we use a post-stack image generated from over 65,000 traces, where stacking has significantly enhanced the signal-to-noise ratio. While this reference is not noise free, it serves as a reasonable proxy ground truth for quantitative evaluation.

Before discussing results, we briefly explain how to read the GPR images. In the radargrams (Figure 2), the horizontal axis corresponds to position along the survey line (or trace index), and the vertical axis corresponds to two-way travel time, which correlates with depth. Bright continuous bands or lines represent horizontal reflecting interfaces, such as layer boundaries, while characteristic hyperbolic patterns indicate reflections from point targets or discrete objects. A successful denoising method should preserve and sharpen these physically meaningful features while suppressing incoherent noise.

We ran DRP on the 256-trace GPR image in a blind fashion and compared its performance to S2S-WTV on the same data. Figure 2 shows that DRP rapidly converges to a clean result that closely resembles the high quality reference in just 250 iterations. The denoised image shows clear hyperbolic reflectors and layered structures, similar to those in the 65k-trace reference. In contrast, the S2S-WTV method exhibits slower convergence; even after 800 iterations, its output contains noticeable residual noise and some reflectors remain smeared.

To quantitatively validate the field results, we compute PSNR and mean squared error with respect to the 65k-trace reference. DRP attains a PSNR of 26.4 dB, higher than S2S-WTV (23.8 dB) and substantially above the raw 256-trace input (about 20 dB). DRP also reduces mean squared error by more than 50% compared to S2S-WTV. Importantly, key geological features such as a prominent hyperbola near 30 m offset and 50 ns two-way time are better resolved in the DRP output and closely match their appearance in the reference radargram. This combination of quantitative and qualitative evidence supports the scientific validity of DRP's blind denoising: the algorithm appears to reveal authentic subsurface signals rather than hallucinating structure.

4 Conclusion

This work demonstrates that Deep Random Projections provide an effective and efficient solution for blind denoising of real-world GPR data. By leveraging the implicit regularization of convolutional

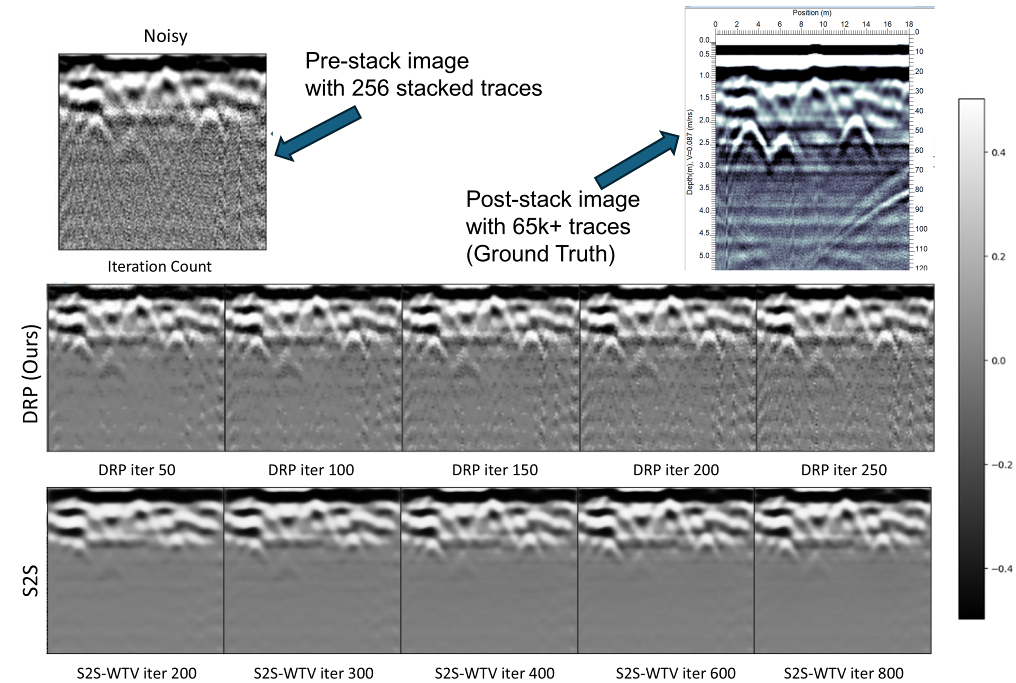

Figure 2: Comparative denoising of field GPR data. DRP (middle row) converges to a clean result closely matching the high trace reference (top right) in only 250 iterations. S2S-WTV (bottom row) remains relatively noisy even after 800 iterations. The top right panel shows the post-stack image formed by stacking over 65k traces.

architectures while dramatically reducing computational requirements through weight freezing and input optimization, DRP achieves denoising performance comparable to or better than state-of-the-art self-supervised methods with two orders of magnitude fewer trainable parameters and 5-10× speedup in processing time. The method's ability to operate without training data or explicit noise models addresses a critical gap in practical GPR processing workflows where obtaining clean reference data is typically infeasible.

Experiments on synthetic and field-collected GPR data confirm that DRP effectively removes diverse noise types while preserving important subsurface features, including layer interfaces, diffraction hyperbolae, and weak deep reflections. The computational efficiency enables near real-time processing suitable for field deployment and iterative acquisition strategies. Future work will explore extensions to 3D processing, incorporation of physics-based constraints such as wave equation priors, and adaptive regularization strategies based on local signal characteristics. More broadly, the results suggest that randomly projected deep priors combined with appropriate architectural biases may provide efficient alternatives to full network optimization in other geophysical imaging and signal enhancement applications.

References

[1] Yilmaz, Ö. (2001). *Seismic data analysis*. Society of Exploration Geophysicists.

[2] Kim, J.-H., Cho, S.-J., and Yi, M.-J. (2007). Removal of ringing noise in GPR data by signal processing. *Geosciences Journal*, 11(1), 75–81.

[3] Daniels, D. J. (2004). *Ground penetrating radar*. IET.

[4] Xu, Z., Luo, Y., Wu, B., and Meng, D. (2023). S2S-WTV: Seismic Data Noise Attenuation Using Weighted Total Variation Regularized Self-Supervised Learning. *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1–15.

[5] Saad, O. M., Oboué, Y. A. S. I., Bai, M., Samy, L., Yang, L., and Chen, Y. (2021). Self-Attention Deep Image Prior Network for Unsupervised 3-D Seismic Data Enhancement. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–14.

[6] Kingma, D. P., and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

[7] Quan, Y., Chen, Y., Pang, Y., and Ji, H. (2020). Self2Self with Dropout: Learning Self-Supervised Denoising from Single Image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

[8] Rudin, L. I., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1–4), 259–268.

[9] Oropeza, V., and Sacchi, M. (2011). Simultaneous seismic data denoising and reconstruction via multichannel singular spectrum analysis. *Geophysics*, 76(3), V25–V32.

[10] Liu, Y., Fomel, S., and Chen, Y. (2018). F-k-t domain seismic data denoising with a deep convolutional autoencoder. In *SEG Technical Program Expanded Abstracts 2018*.

[11] Song, P., Liu, Y., and Li, Z. (2019). F-k Filter Designs to Suppress Direct Waves for Bistatic Ground Penetrating Radar. *IEEE Transactions on Geoscience and Remote Sensing*.

[12] Li, T., Wang, H., Zhuang, Z., and Sun, J. (2023). Deep Random Projector: Accelerated Deep Image Prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

[13] Batson, J., and Royer, L. (2019). Noise2Self: Blind Denoising by Self-Supervision. In *International Conference on Machine Learning (ICML)*.

[14] Ulyanov, D., Vedaldi, A., and Lempitsky, V. (2018). Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[15] Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K. (2007). Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8), 2080–2095.

[16] Conyers, L. B. (2004). *Ground-penetrating radar for archaeology*. Rowman Altamira.

[17] Annan, A. P. (2009). Electromagnetic principles of ground penetrating radar. In *Ground penetrating radar: Theory and applications*. Elsevier.

[18] Jol, H. M. (Ed.). (2008). *Ground penetrating radar theory and applications*. Elsevier.

[19] Davis, J. L., and Annan, A. P. (1989). Ground-penetrating radar for high-resolution mapping of soil and rock stratigraphy. *Geophysical Prospecting*, 37(5), 531–551.

[20] Annan, A. P. (2005). GPR methods for hydrogeological studies. In *Hydrogeophysics*. Springer.

[21] Olhoeft, G. R. (2000). Maximizing the information return from ground penetrating radar. *Journal of Applied Geophysics*, 43(2–4), 175–187.

[22] Goodman, D., and Piro, S. (Eds.). (2013). *GPR remote sensing in archaeology*. Springer.

[23] Leckebusch, J. (2003). Ground-penetrating radar: a modern three-dimensional prospection method. *Archaeological Prospection*, 10(4), 213–240.

[24] Strubell, E., Ganesh, A., and McCallum, A. (2019). Energy and Policy Considerations for Deep Learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.

[25] Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015*.

[26] Duchi, J., Hazan, E., and Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7).

[27] Redman, J. D. (1994). Ground penetrating radar for geotechnical applications. In *Proceedings of the Symposium on the Application of Geophysics to Engineering and Environmental Problems*.

[28] Heckel, R., and Hand, P. (2019). Deep decoder: Concise image representations from untrained non-convolutional networks. In *International Conference on Learning Representations (ICLR)*.

[29] Sitzmann, V., Martel, J. N., Bergman, A. W., Lindell, D. B., and Wetzstein, G. (2020). Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*.

[30] Jagatap, G., and Hegde, C. (2019). Algorithmic guarantees for inverse imaging with untrained network priors. In *Advances in Neural Information Processing Systems*.

[31] Portniaguine, O., and Zhdanov, M. S. (1999). Focusing geophysical inversion data. *Geophysics*, 64(3), 874–887.

[32] Farquharson, C. G. (2008). Constructing piecewise-constant models in multidimensional minimum-structure inversions. *Geophysics*, 73(1), K1–K9.

[33] Li, Y., and Oldenburg, D. W. (1996). 3-D inversion of magnetic data. *Geophysics*, 61(2), 394–408.

[34] Kong, Q., Trugman, D. T., Ross, Z. E., Bianco, M. J., Meade, B. J., and Gerstoft, P. (2018). Machine Learning in Seismology: Turning Data into Insights. *Seismological Research Letters*, 90(1), 3–14.

[35] Dabov, Kostadin, et al. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8), 2080–2095, 2007.