Robustness of Persistent Homology when Noising and Denoising 3D Images

Aakash Karlekar Ebru Dagdelen Manav Arora Matthew Illingworth Jonathan Jaquette Linda J. Cummings Lou Kondic

Department of Mathematical Sciences, New Jersey Institute of Technology

Introduction

Porous media are ubiquitous in nature and industry, arising in a multitude of processes including environmental cleanup, oil recovery, and CO_2 storage. Understanding and optimizing such processes, which typically involve flow through porous media, requires characterizing the complex internal pore structure of these materials. Techniques from Topological Data Analysis, especially persistent homology, are very helpful for this task. However, when working with real-world data — specifically 3D images of porous media — we face significant computational challenges because of the complexity of the datasets and the presence of experimental noise, which can obscure key topological features and increase computational costs.

We propose denoising such images using Gaussian convolution, with the goal of smoothing the data and reducing noise, while preserving true topological features. We demonstrate the method using simulated image datasets, to which we add noise to mimic real experimental data, then denoise. To assess the effectiveness of our method, we use several topological measures to compare the original and denoised datasets. Finally, we discuss the optimal denoising approach with Gaussian convolution that makes these measures closest to the original, noise-free data. We further look to extend our analyses to other approaches to denoising, such as those based on machine learning.

Topological Data Analysis

In the context of analyzing porous media, digital representations are typically either binary or grayscale. In the present work, we start with 3D grayscale images of porous media. Through a thresholding process based on the grayscale level of each voxel in the image, we construct a nested sequence of geometric layers called a topological filtration (see Fig. 1).

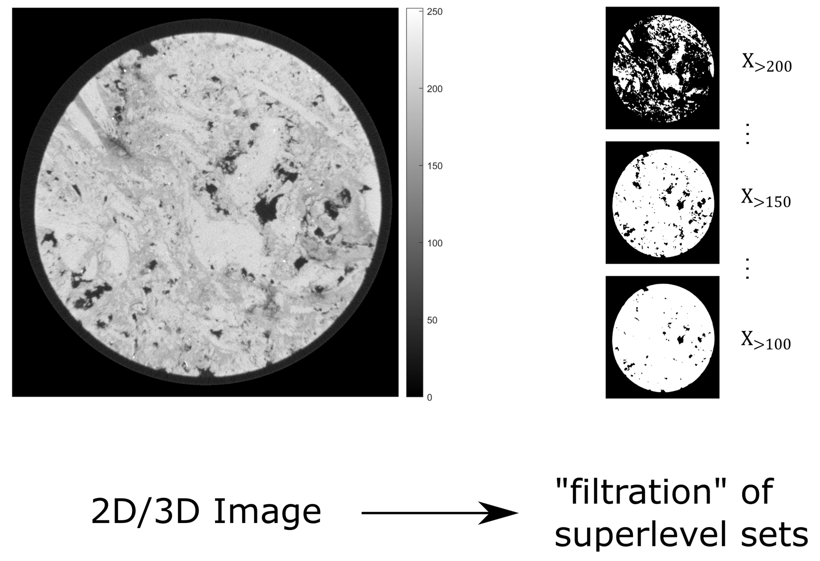

Figure 1. Superlevel thresholding visualization (2D Slice)

We can study the topological properties of each of these filtration layers by computing its relevant homology groups. These groups are most simply described in terms of Betti numbers, of which there are three for a 3D sample, $\beta_0, \beta_1, \beta_2$, corresponding to the numbers of connected components, handles, and cavities found within each layer, respectively [4, 1].

PH looks to synthesize this information by quantifying the topological features that persist across many layers of a filtration. As we progress through a filtration, the threshold value at which a feature appears is referred to as its birth number. Similarly, when we encounter a threshold value at which a feature disappears, we refer to it as the feature's death number.

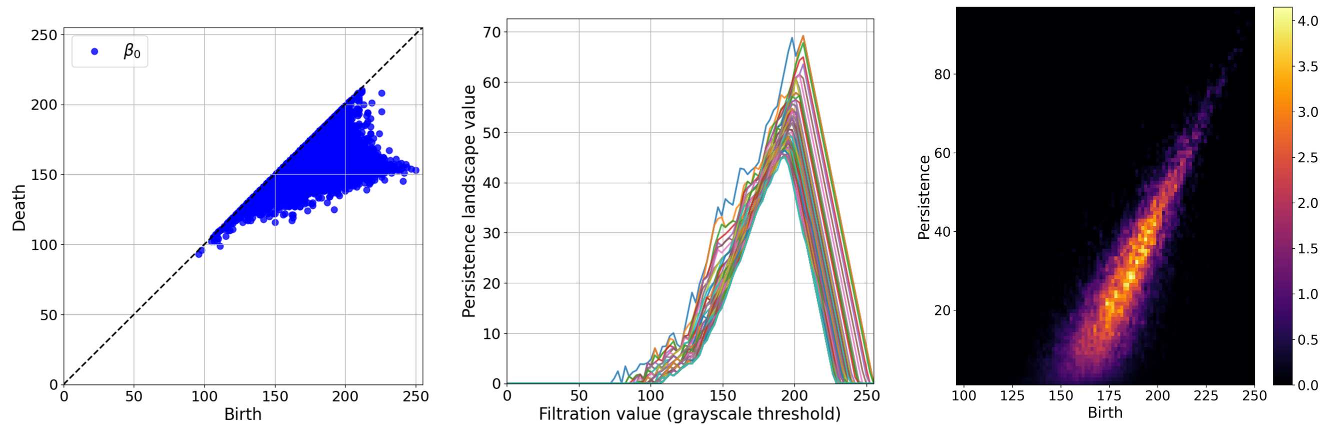

(a) PD (b) PL (c) PI

Figure 2. Persistence homology outputs

Topological Data Analysis (cont.)

We can then visualize the persistence of features as points in a 2D graph, where one axis represents each feature's birth number and the other its death number. By plotting the corresponding birth-death pairs for the features we encounter in our filtration, we obtain a **Persistence Diagram (PD)** (Fig. 2a) for a given homology group. We can use a PD to generate **Persistence Landscapes** (Fig. 2b) and **Persistence Images (PIs)** (Fig. 2c) which are vectorized representations of the topological composition of our porous material.

Methodology

Our methodology, outlined in Fig. 3, starts with a simulated dataset intended to represent a 3D image of a porous medium, which we refer to as the **original dataset**. We then add controlled static noise from a known distribution to this dataset to produce a **noisy dataset**. Next, we apply our smoothing method on the noisy dataset, referring to the resulting dataset after smoothing as the **denoised dataset**.

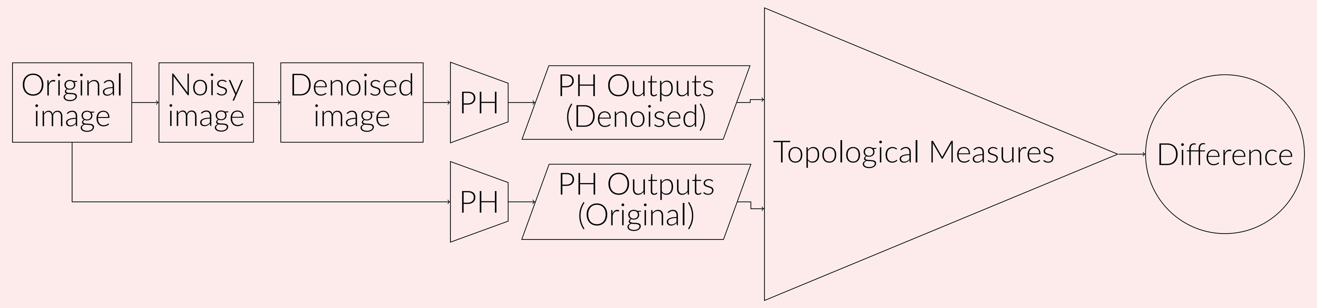

Figure 3. Flowchart of methodology.

We identify several topological measures that we use to quantify the effectiveness of our denoising algorithm, using the fact that we know the original state of our data before noise was introduced.

Datasets

To analyze our denoising methods, we apply our testing process to several datasets shown in Figure 4: A synthetic 3D image generated using Fourier series, which we refer to as the Fourier dataset; a dataset consisting of random, overlapping spheres generated using the PuMA library which we refer to as the PuMA dataset [2]; and a Worly noise based dataset which we refer to as the cellular dataset [3].

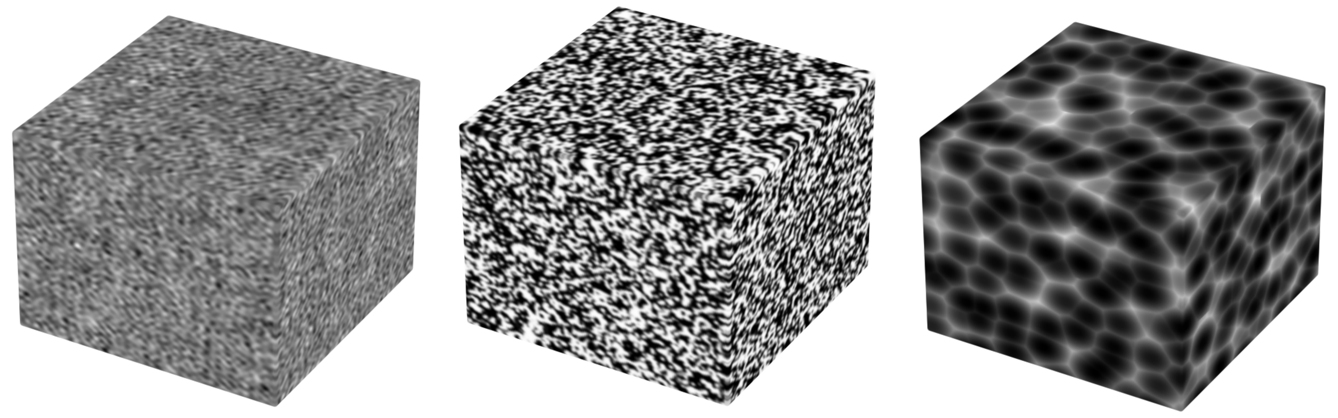

(a) Fourier (b) PuMA (c) Cellular

Figure 4. 3D views of the three datasets.

Noising

For many experimental samples, it is found that the overall noise approximately follows a normal distribution. We model the noise intensity value, g, using the univariate probability density function

$$\mathcal{N}(g; \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(g-\mu)^2}{2\sigma^2}\right) \tag{1}$$

where μ is the mean and σ^2 is the variance. In our analysis, we take $\mu = 0$ and choose a random standard deviation $\sigma \in [0, 10]$ for each generated dataset. We sample from this distribution 254^3 times independently—once for each voxel center (x, y, z)—to create a noise array of shape $254 \times 254 \times 254$. This noise array is then added to the original dataset voxel-wise to produce our noisy dataset.

Denoising

After producing our $254 \times 254 \times 254$ noisy image, we denoise it with our smoothing method, based on the discrete Gaussian convolution using Eq. (1). Here we once again take $\mu = 0$ and vary our denoising level σ on the interval $[0, 10]$.

While naive Gaussian smoothing provides a computationally inexpensive baseline for noise reduction, its symmetric nature inherently blurs high-frequency structural boundaries, leading to poor edge preservation. Furthermore, it requires empirical tuning of the parameter σ to adapt to different noise profiles. To mitigate these limitations, we developed a **convolutional U-Net** architecture designed to learn a robust, non-linear mapping from the noisy inputs to reconstruct the clean, original images.

The U-Net architecture leverages an encoder-decoder structure with skip connections, allowing the model to effectively capture both broad spatial contexts and fine-grained local geometries. This architecture enables the model to preserve edges and critical topological features more effectively than linear filters, while also eliminating the need for manual parameter tuning during inference.

The network was optimized using a Huber loss function, which dynamically balances the benefits of Mean Squared Error (MSE) and Mean Absolute Error (MAE). For small errors, the loss behaves as MSE—the Maximum Likelihood Estimator in the presence of Gaussian noise—which ensures stable convergence. For large errors, it transitions to an MAE penalty, which limits the influence of outliers and fundamentally aids in preserving sharp structural edges.

For each of the three morphological datasets, the model was trained and evaluated on a set of 450 image volumes. Each dataset was partitioned using a standard 64%–16%–20% split, yielding 288 volumes for training, 72 for validation to monitor overfitting, and 90 for the final hold-out testing phase.

Topological Measures

We analyze the efficacy of our denoising methods through the lens of several normalized measures based on persistence statistics and various persistence representations. One topological measure we choose to illustrate here is the normalized L^2 distance between PIs defined as

$$PI_i := \frac{\|I_{\mathrm{PD}_i^o} - I_{\mathrm{PD}_i^d}\|_2}{\|I_{\mathrm{PD}_i^o}\|_2} \tag{2}$$

where $I_{\mathrm{PD}_i^o}$ and $I_{\mathrm{PD}_i^d}$ refer to the PIs of the i-th homology group corresponding to the original and denoised images respectively.

Results

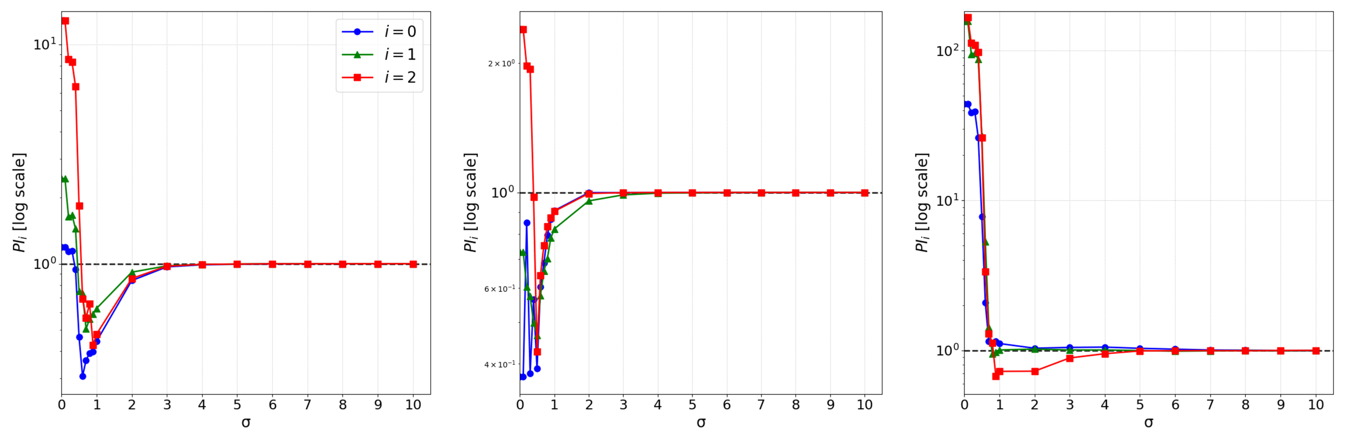

(a) Fourier dataset (b) PuMA dataset (c) Cellular dataset

Figure 5. Normalized L^2 distance PI_i between persistence images of the original and denoised datasets, for various denoising levels $\sigma \in [0, 5]$

- **Gaussian convolution:** Gaussian convolution successfully reduces noise and minimizes overall dataset differences at an optimal parameter of $\sigma \approx 0.5$ (see Fig. 5). However, it remains limited by the need for empirical tuning and its inherent trade-off between smoothing and structural preservation.
- **Robustness of Topological Measures:** Our analysis identifies several topological measures, such as PI_i (Defined in Eq. (2)), are highly robust in terms of sensitivity to the level of noising and denoising.

Results (cont.)

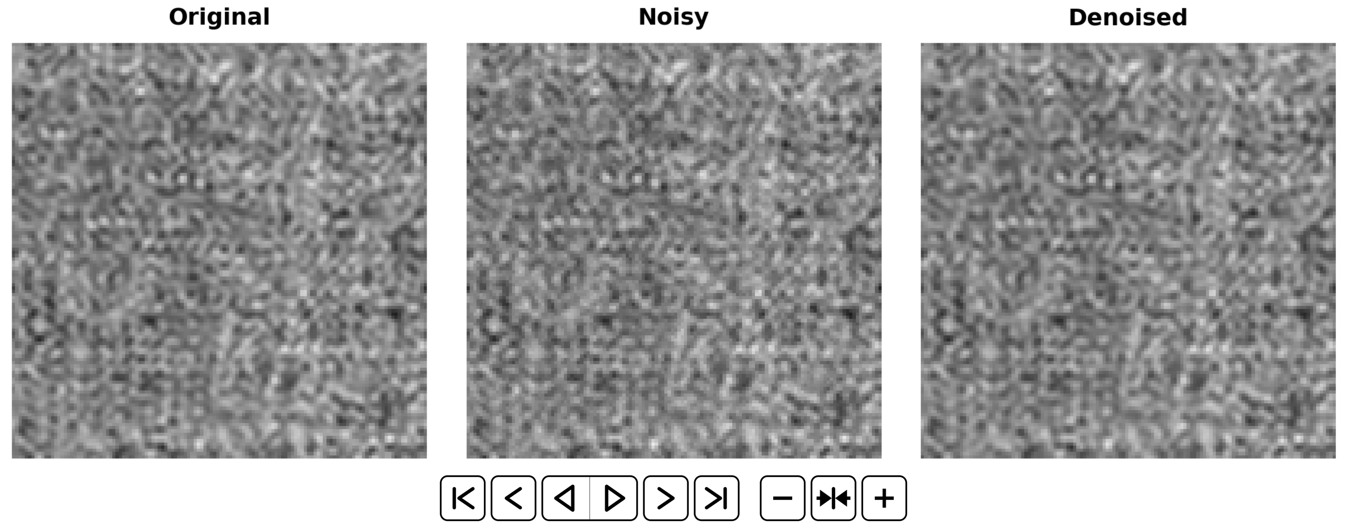

Figure 6. Animation of noising and denoising process (2D Slice).

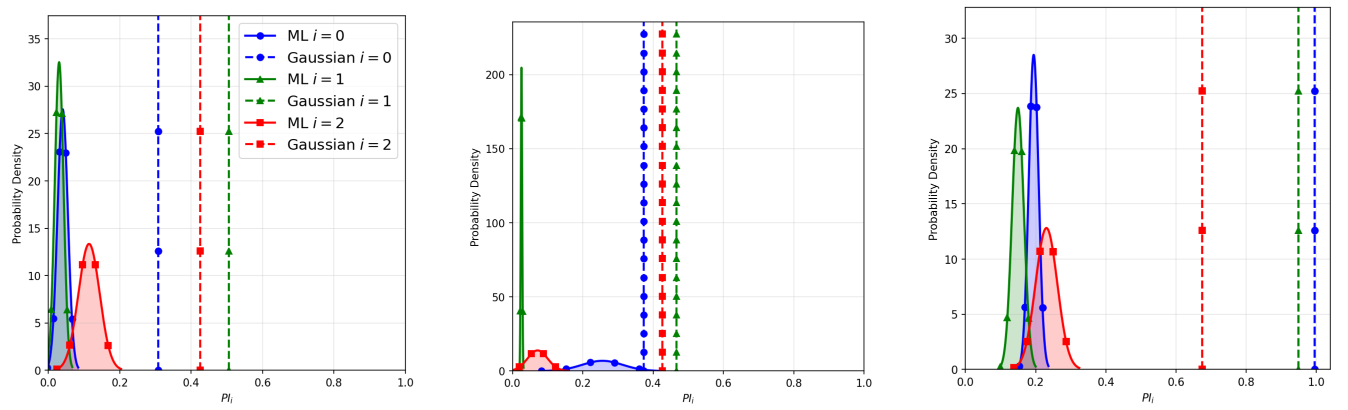

(a) Fourier (b) PuMA (c) Cellular

Figure 7. Denoising performance of our machine learning model across 90 test images for each of the datasets.

- **Machine Learning:** The proposed U-Net architecture outperforms the traditional Gaussian smoothing baseline (see Figs. 6 & 7) without requiring manual parameter tuning during inference.
- **Broader Implications:** These findings validate the use of deep learning as a promising way to preserve, critical fine-grained microstructures and topological features when denoising diverse morphological datasets.

References

[1] Herbert Edelsbrunner and John Harer. Persistent homology—a survey. In *Contemporary Mathematics*, volume 453, pages 257–282. American Mathematical Society, 2008.

[2] Joseph C Ferguson, Federico Semeraro, John M Thornton, Francesco Panerai, Arnaud Borner, and Nagi N Mansour. Update 3.0 to "puma: The porous microstructure analysis software",(pii: S2352711018300281). *SoftwareX*, 15:100775, 2021.

[3] Robert A. McLeod. pyfastnoisesimd. https://github.com/robbmcleod/pyfastnoisesimd. Accessed: 2025-08-04.

[4] James R. Munkres. *Elements of algebraic topology*. Perseus Publ, 1997.

Connect

LinkedIn (Aakash) LinkedIn (Ebru) CFSM Website