Brain Hemorrhage Detection With DDPM Reconstruction

Dennis Loevlie, Xianmai Liang, Roger Brockett

Introduction

The work investigates diffusion models for anomaly detection and applies it to hemorrhage detection in slices of CT scans.

Hypotheses

The experiments are guided by the following hypotheses,

- A diffusion model trained on only "healthy" images will generate a "healthy" image even when provided noise created from an "unhealthy" image.
- The reconstruction error can be used to differentiate healthy brain scans from unhealthy.

Synthetic MNIST Dataset Generation

Given an image $\mathbf{X} \in [0,1]^{28\times 28}$, we generate a corrupted image $\tilde{\mathbf{X}}$ by overlaying a white splotch. First, we sample a splotch center $(c_x, c_y) \sim \mathcal{U}\{6, 22\}^2$ and base radius $r \sim \mathcal{U}\{2, 5\}$. We then generate $n \sim \mathcal{U}\{3, 7\}$ overlapping circular blobs, each with center offsets $\Delta x_i, \Delta y_i \sim \mathcal{U}\{-\lfloor r/2 \rfloor, \lfloor r/2 \rfloor\}$ and radius $r_i \sim \mathcal{U}\{\lfloor r/2 \rfloor, r\}$. The binary mask is:

$$\mathbf{M}(x,y) = \mathbb{1}\left[\bigcup_{i=1}^{n} \sqrt{(x - c_x - \Delta x_i)^2 + (y - c_y - \Delta y_i)^2} \leq r_i\right] \qquad (1)$$

The corrupted image is then $\tilde{\mathbf{X}} = (1 - \mathbf{M}) \odot \mathbf{X} + \mathbf{M}$.

The Model:

We trained a UNet with base channels 64 and channel multipliers (1, 2, 4) using AdamW (lr = 1e-4) for 20 epochs on MNIST. The diffusion process used 1000 timesteps with a linear β schedule from 0.0001 to 0.02.

Results:

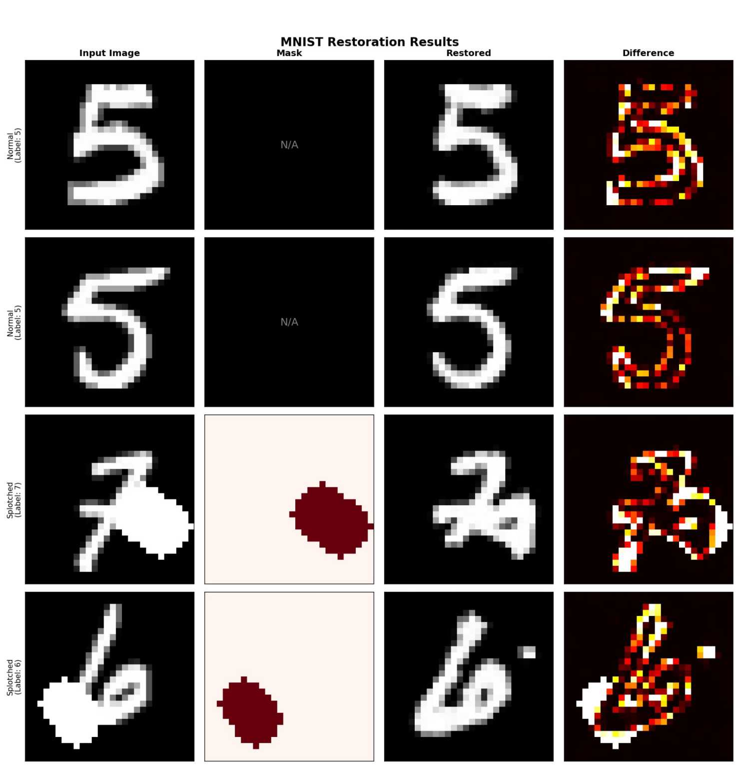

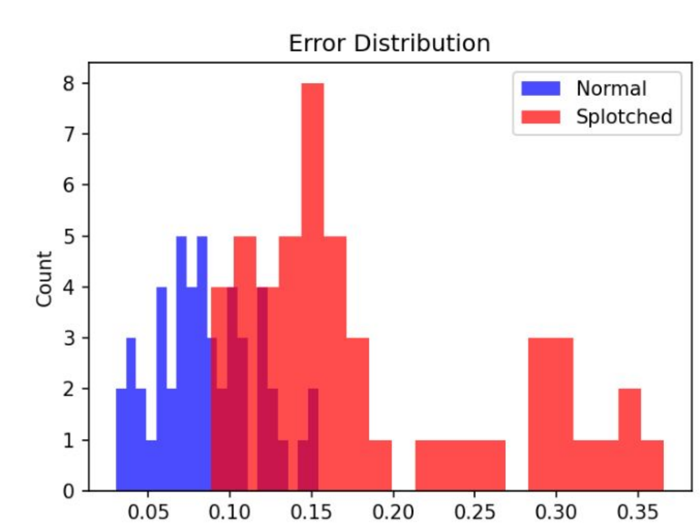

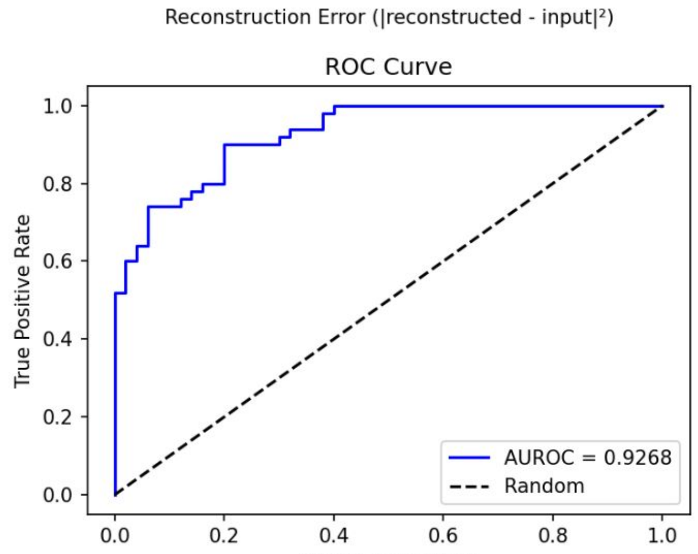

Hemorrhage Detection on CT Scans

The exploration aims to use the reconstruction error to identify hemorrhage in brain scans.

Setting	UNet (Scratch)	UNet (Pretrained)
Architecture	4-level encoder-decoder UNet	HuggingFace google/ddpm-ema-celebahq-256
# Levels	4 encoder + 4 decoder blocks	4 levels (~28 ResNet blocks total)
Channel dimensions	[64, 128, 256, 512]	[128, 128, 256, 256, 512, 512]
Input channels	1 (grayscale)	3 (RGB, duplicated from grayscale)
Total Parameters	~30M	~114M
Optimizer	AdamW	AdamW
Learning rate	1e-4	1e-4
# Training epochs	10	10

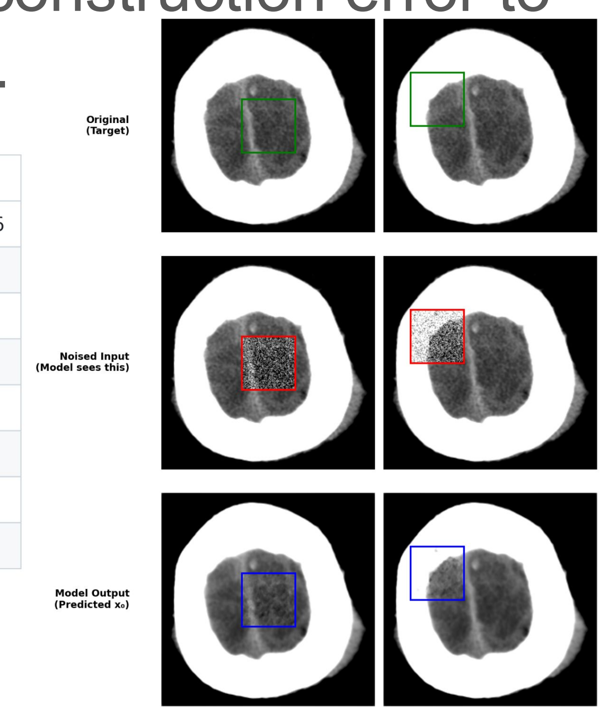

Examples:

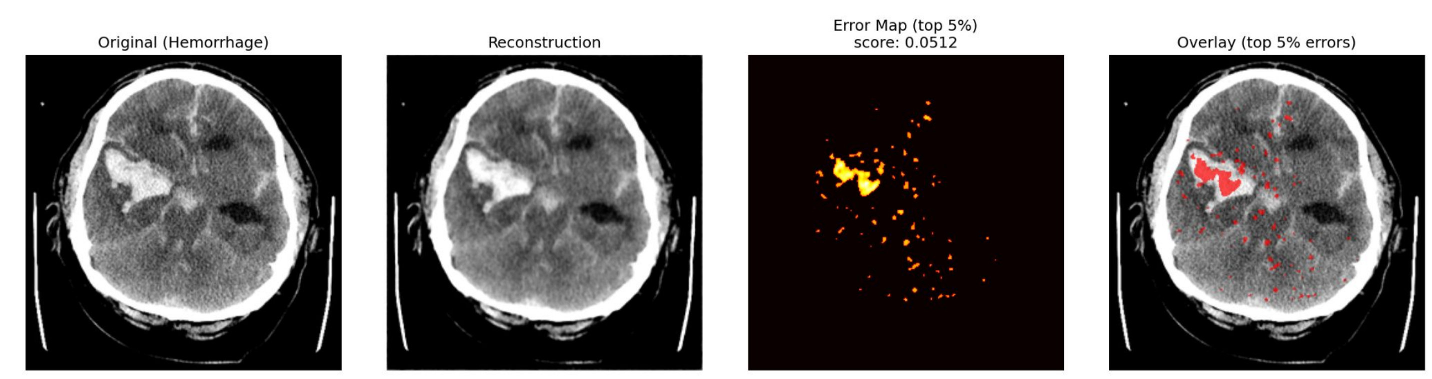

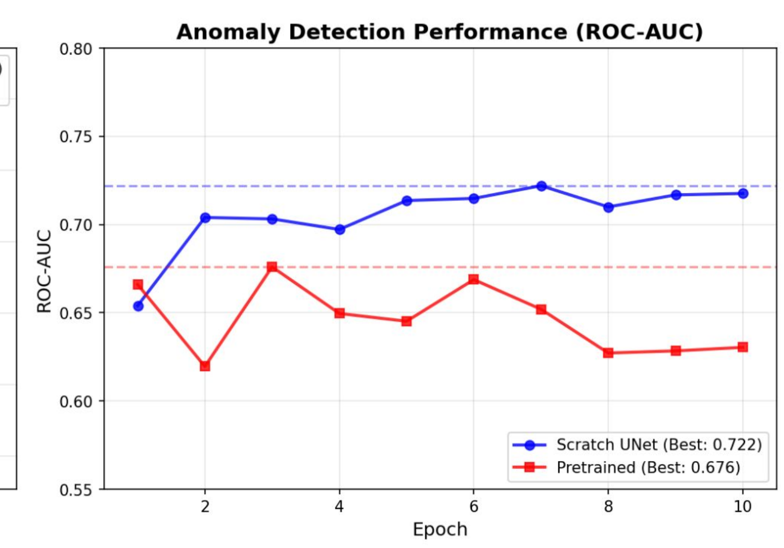

Conclusion

- Diffusion models are decent at image-level anomaly detection.
- They do not perform as well at pixel-level anomaly detection.
- Pretrained models, although they achieve lower validation reconstruction error, are worse at anomaly detection.

Future work

- Perform a more robust hyperparameter search
- Evaluate other generative models (VAE, GAN, etc.)
- Evaluate adding noise to the latent space

Citation

Behrendt, F., Bhattacharya, D., Krüger, J., Opfer, R., & Schlaefer, A. (2024). Patched Diffusion Models for Unsupervised Anomaly Detection in Brain MRI. *Medical Imaging with Deep Learning*, in *Proceedings of Machine Learning Research* 227:1019-1032.