MAD: Manifold Attracted Diffusion

Dennis Elbrächter
Faculty of Mathematics
University of Vienna
Vienna, Austria

Giovanni S. Alberti
MaLGa Center
Department of Mathematics
University of Genoa
Genoa, Italy

Matteo Santacesaria
MaLGa Center
Department of Mathematics
University of Genoa
Genoa, Italy

A Case Study: Cryo Electron Microscopy

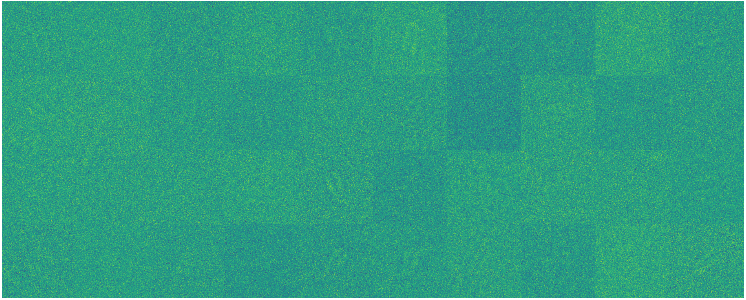

Real Data[1] (Extremely noisy)

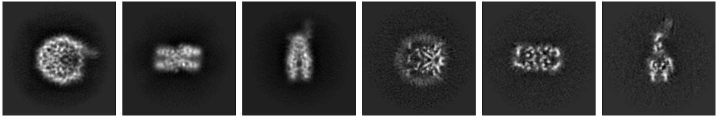

Example reconstructions from [Bacic et al., 2021]

[1][wwPDB Consortium, 2023]

The Manifold Hypothesis[2]

- Meaningful data distributions (images, physical fields) are concentrated near a low-dimensional manifold M embedded in a high-dimensional ambient space X.
- Noise typically manifests as off-manifold variations, while meaningful features correspond to on-manifold geodesics.

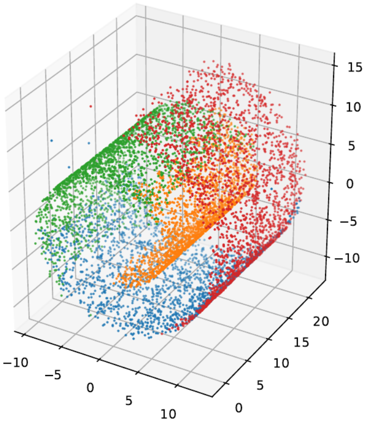

[Hertrich et al. 2024]

UniGe | MaLGa
Matteo Santacesaria

[2] [Bengio et al. 2013, Fefferman et al., 2016]

Diffusion Models[3]: Forward and Reverse Processes

Forward Process (Variance Exploding):

- Given a random variable $X_0 \in \mathbb{R}^d$ (the data distribution).
- Define a stochastic process $(X_t)_{t\in[0,T]}$ by progressively adding Gaussian noise:

$$X_t = X_0 + \mathcal{N}(0, \sigma(t)^2 I_d)$$

- This corresponds to the probability flow ODE:

$$dx = -\dot{\sigma}(t)\sigma(t)\nabla \log p_{X_t}(x)dt$$

  where p_{X_t} is the marginal probability density at time t.

Reverse Process (Inference):

- Start with pure noise $x_T \sim \mathcal{N}(0, \sigma(T)^2 I_d)$.
- Evolve the ODE backwards from T to 0 to generate a sample $x_0 \sim X_0$.
- The core challenge: We need to know the **score function** $\mathcal{S}p_{X_t}(x) = \nabla \log p_{X_t}(x)$.

[3][Song-Ermon 2019]

The Score Operator and The Singularity Problem

Let $M = \{p \in C^1(\mathbb{R}^d) : p(x) > 0 \ \forall x \in \mathbb{R}^d\}$. We define the score operator:

$$S : M \to C(\mathbb{R}^d)$$
$$p \mapsto \nabla \log p = \frac{\nabla p}{p}$$

The Singularity Problem

- Consider a simple Gaussian $g_{\sigma^2}(x) = (2\pi\sigma^2)^{-\frac{d}{2}} e^{-\frac{\|x\|^2}{2\sigma^2}}$.
- Its score is exactly computed as:

$$Sg_{\sigma^2}(x) = -\frac{x}{\sigma^2}$$

- As variance vanishes ($\sigma \to 0$), approaching a Dirac delta or a singular manifold:

$$\lim_{\sigma \to 0} Sg_{\sigma^2}(x) = \pm\infty$$

- Conclusion: If the true data distribution is supported on a low-dimensional manifold, the standard score S is ill-defined at the continuous limit.

Manifold Attracted Diffusion (MAD)

Our Method (MAD):

- A modified **inference procedure** (no change to training).
- We introduce an "Extended Score" H_0 that acts as a soft-threshold: it suppresses small variations (noise) while preserving large ones (data).

Definition of the Extended Score (H_0)

- We define a regularized operator H_γ for $\gamma > 0$ and measure p:

$$H_\gamma p = (1+\gamma)S(p * g_\gamma) + \gamma \frac{\partial}{\partial \gamma} S(p * g_\gamma)$$

- The extended score is the limit: $H_0 p = \lim_{\gamma \to 0} H_\gamma p$.

The Extended Score

Definition of the Extended Score (H_0)

$$H_0 p = \lim_{\gamma \to 0} \left[(1+\gamma) S(p * g_\gamma) + \gamma \frac{\partial}{\partial \gamma} S(p * g_\gamma) \right]$$

Key Properties

- **Consistency:** For distributions with a density, $H_0 p = Sp$.
- **Dirac Delta:** For $p = \delta_\mu$, it is well-defined: $H_0 \delta_\mu(x) = -(x - \mu)$. It points directly toward the mode.
- **Mixtures:** For $p = \sum c_i \delta_{\mu_i}$, $H_0 p(x)$ points toward the nearest μ_i.
- **No changes in training.**

MAD: Illustrative Examples in $\mathbb{R}^2$

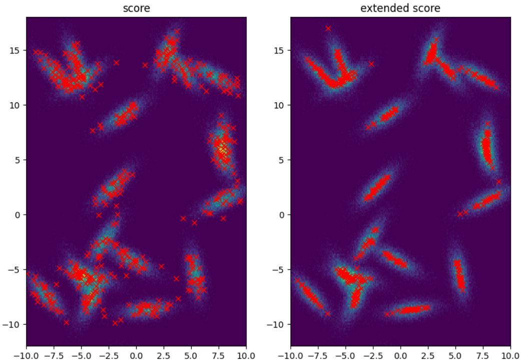

Tilted Gaussians

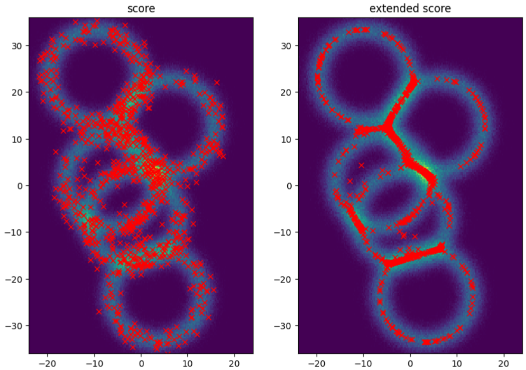

Radial Gaussians

UniGe | MaLGa

Matteo Santacesaria

MAD: Effects on FFHQ

Standard score (left) vs. Extended score with increasing b (right)

MAD: Denoising Results (Trained on Noisy Data)

Synthetic Shapes

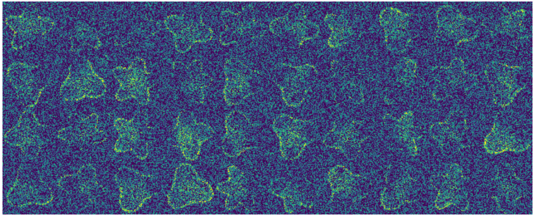

Standard Inference (reproduces noise)

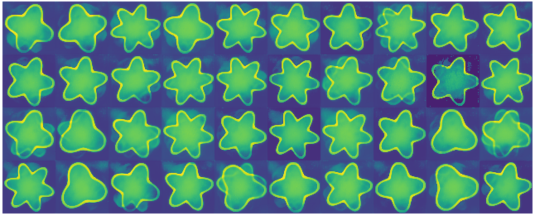

MAD Inference (recovers clean shapes)

Real Cryo-EM Data

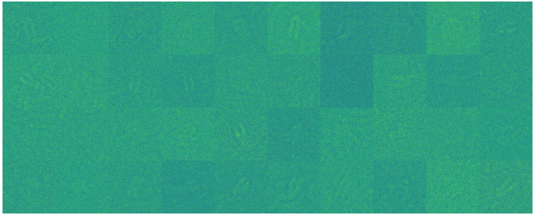

Training Data (Extremely noisy)

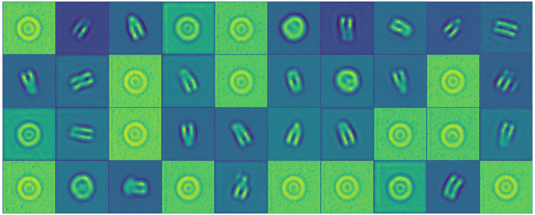

MAD Inference (recovers clean particles)

Conclusions

Additional results in the paper

- More theoretical results
- Pixel removal and shot noise on FFHQ
- Gaussian denoising on CIFAR-10

Take home message

- **MAD** allows us to implicitly learn and sample from the underlying clean manifold, even when trained on noisy data.

Open questions

- Conditional diffusion for solving inverse problems
- Extensions to other diffusion models (e.g. variance preserving)