Reflected Diffusion Models

Aaron Lou, Stefano Ermon

Diffusion Models Recap

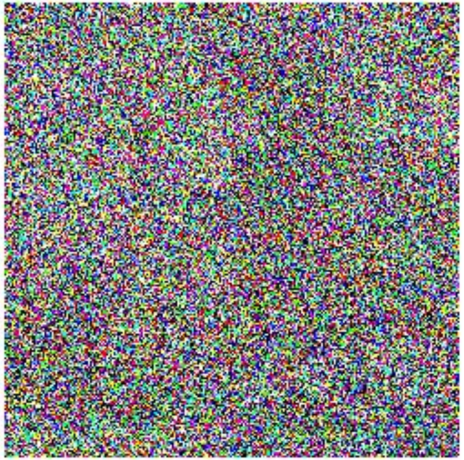

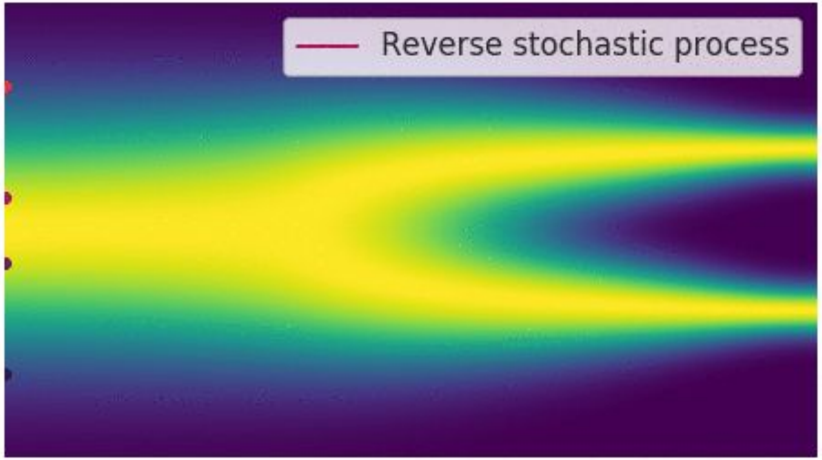

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t)\nabla_{\mathbf{x}} \log p_t(\mathbf{x})]dt + g(t)d\mathbf{B}_t$$

Minimize $\mathbb{E}_{\mathbf{x}_0 \sim p_0, \mathbf{x} \sim p_t(\cdot|\mathbf{x}_0)} ||\mathbf{s}_\theta(\mathbf{x}, t) - \nabla_x \log p_t(\mathbf{x}|\mathbf{x}_0)||^2$

Divergent Sampling

Not Learned Perfectly Discretized Random Variable

$$\mathbf{x}_{t-\Delta t} = \mathbf{x}_t - \left[\mathbf{f}(\mathbf{x}_t, t) - g(t)^2 \boxed{s_\theta(\mathbf{x}_t, t)}\right] \boxed{\Delta t} + g(t)\boxed{\mathbf{B}_{\Delta t}^{t}}$$

Guided Diffusion (2021)

Imagen (2022)

Thresholding

$$\mathbf{x}_{t-\Delta t} = \underbrace{\mathbf{x}_t - \left[\mathbf{f}(\mathbf{x}_t, t) - g(t)^2 s_\theta(\mathbf{x}_t, t)\right] \Delta t}_{\text{VAE Predicted Mean } \overline{\mathbf{x}}_{t-\Delta t}} + \underbrace{g(t)\mathbf{B}^t_{\Delta t}}_{\text{VAE Noise}}$$

```
x_recon = tf.clip_by_value(x_recon, -1., 1.)
```


Imagen (2022)

Reflected SDEs

$$\mathbf{x}_{t+\Delta t} = \text{proj}\left(\mathbf{x}_t + \mathbf{f}(\mathbf{x}_t, t)\Delta t\right) + g(t)\mathbf{B}^t_{\Delta t}$$

$$\Delta t \to 0$$

$$\mathrm{d}\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t)\mathrm{d}t + g(t)\mathrm{d}\mathbf{B}_t + \boxed{\mathrm{d}\mathbf{L}_t}$$

Reversing Reflected SDEs

$$\mathrm{d}\mathbf{x}_t = \left[\mathbf{f}(\mathbf{x}_t, t) - g(t)^2 \nabla_x \log p_t\right] \mathrm{d}t + g(t)\mathrm{d}\overline{\mathbf{B}}_t + \mathrm{d}\overline{\mathbf{L}}_t$$

$$\mathbb{E}_{\mathbf{x}_0 \sim p_0, \mathbf{x} \sim p_t(\cdot|\mathbf{x}_0)} \left\| \mathbf{s}_\theta(\mathbf{x}, t) - \nabla_x \log p_t(\mathbf{x}|\mathbf{x}_0) \right\|^2$$

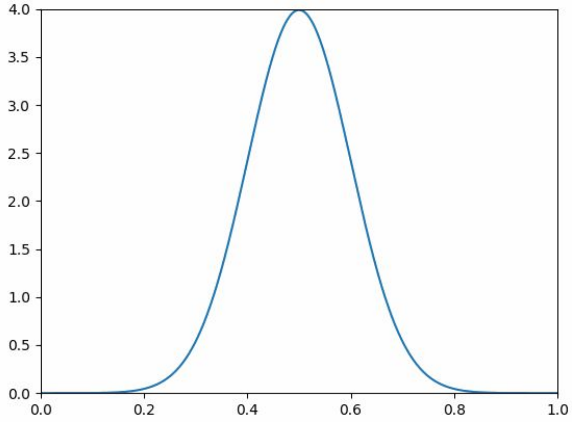

New Kernel

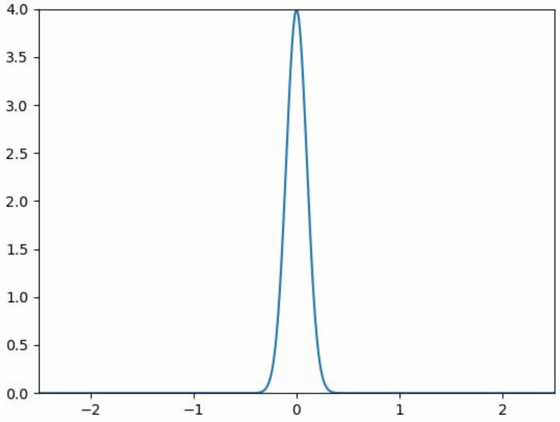

Gaussian Kernel

Results

Method	Inception score (↑)
NCSN++ [3]	9.89
Subspace Diffusion [13]	9.99
Ours	**10.42**

Method	CIFAR-10 BPD (↓)	ImageNet-32 BPD (↓)
ScoreFlow [10]	2.86	3.83
(with importance sampling)	*2.83*	*3.76*
VDM [15]	2.70	——
(with learned noise)	*2.65*	*3.72*
Ours	**2.68**	**3.74**

Results (cont.)

Thresholding with high weight

Ours with high weight

Results (cont.)

Sampled with far fewer steps. ~100 compared to ~1000

Generalizing Geometry

Our approach can generalize to other types of geometries, such as high dimensional simplices.

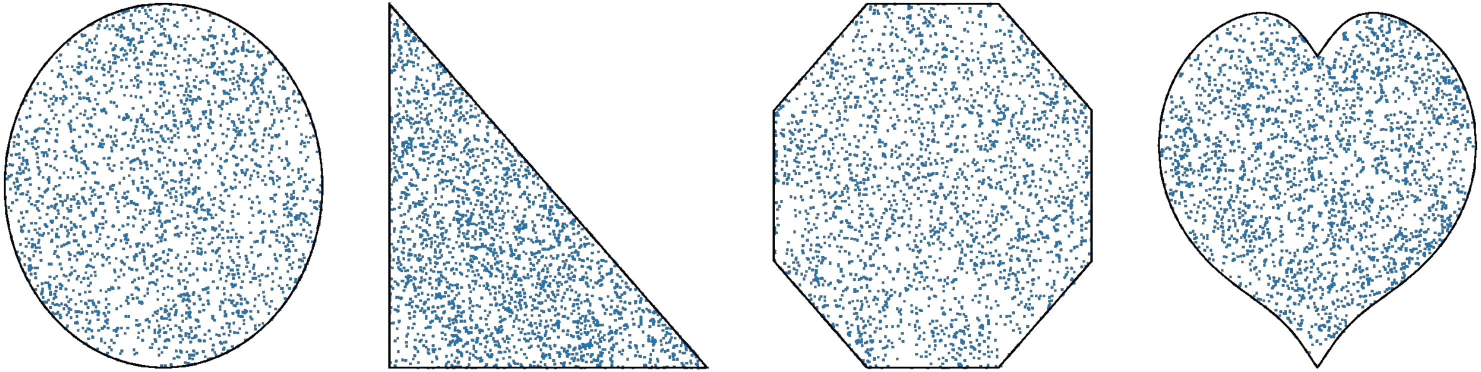

Takeaways and Conclusion

- Common sampling tricks have nice theoretical explanations.
- We can build a general framework that respects our data constraints.
- Method allows for further extensions (e.g. general geometries).

Check out our full paper for more information!

Page

Code

@aaron_lou
@StefanoErmon