Guidance

Eric Paquet

11 March 2024

Score-based diffusion model

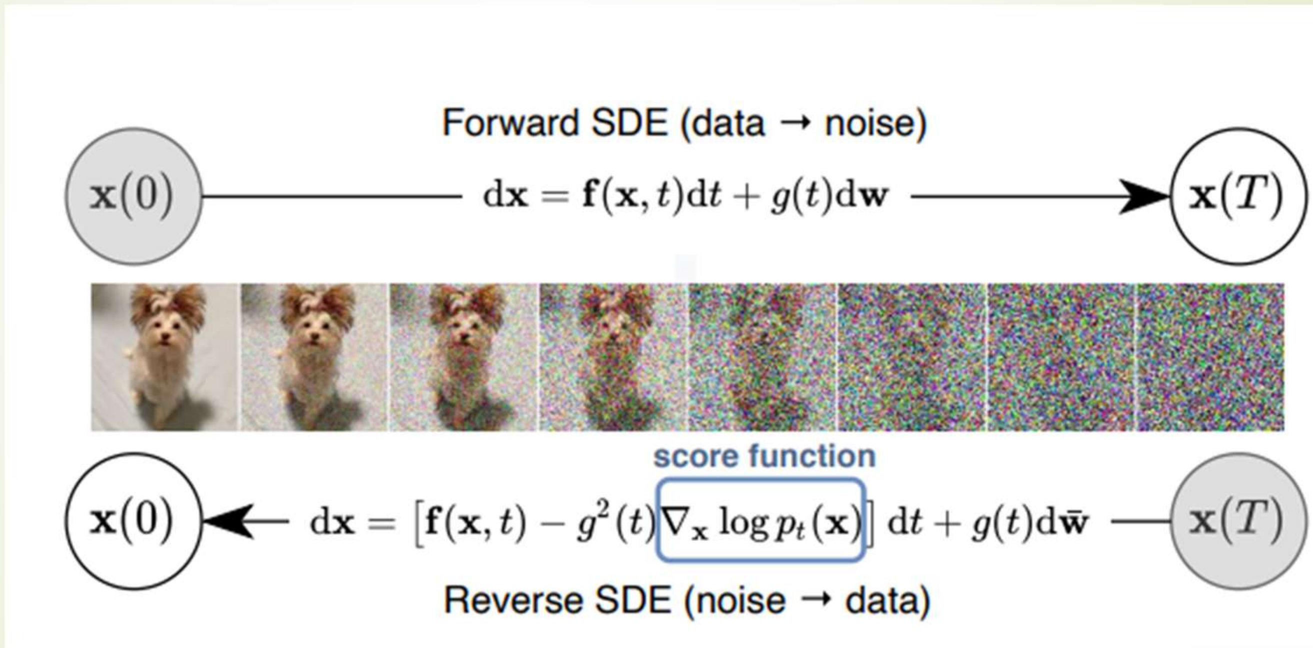

Vanilla conditioning

- Conditioning the data distribution alongside the timestep information, at each iteration
- A conditional model trained in this way may potentially learn to ignore or downplay any given conditioning information, as pointed out by the literature
- Training

$$p\left(\mathbf{x}_{0:T}\right)=p\left(\mathbf{x}_T\right)\prod_{t=1}^{T}p_\theta\left(\mathbf{x}_{t-1}\middle|\mathbf{x}_t\right)\Rightarrow$$

$$p\left(\mathbf{x}_{0:T}\middle|E\right)=p\left(\mathbf{x}_T\right)\prod_{t=1}^{T}p_\theta\left(\mathbf{x}_{t-1}\middle|\mathbf{x}_t,E\right)$$

$$\hat{\mathbf{x}}_\theta\left(\mathbf{x}_t\middle|E\right)\approx\mathbf{x}_0$$

$$\hat{\varepsilon}_\theta\left(\mathbf{x}_t\middle|E\right)\approx\varepsilon_0$$

$$\hat{\mathbf{s}}_\theta\left(\mathbf{x}_t\middle|E\right)\approx\nabla_{\mathbf{x}_t}\log p\left(\mathbf{x}_t\middle|E\right)$$

Classifier guidance

- The conditional model trained in this way may potentially learn to ignore or downplay any given conditioning information
- One additional model for the classifier/predictor that must be trained
- Bayes' rule:

$$\nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E\right) = \nabla_{\mathbf{x}_t} \log\left(\frac{p\left(\mathbf{x}_t\right) p\left(E \middle| \mathbf{x}_t\right)}{p\left(E\right)}\right) =$$

$$\underbrace{\nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t\right)}_{\text{unconditional score}} + \underbrace{\nabla_{\mathbf{x}_t} \log p\left(E \middle| \mathbf{x}_t\right)}_{\text{adversarial gradient}} \Rightarrow$$

$$\nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E\right) = \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t\right) + \gamma \nabla_{\mathbf{x}_t} \log p\left(E \middle| \mathbf{x}_t\right)$$

Approximation of the classifier

- Taylor expansion (linear):

$$\log p_\phi(E|\mathbf{x}_t) \approx \log p_\phi(E|\mathbf{x}_t)\Big|_{\mathbf{x}_t=\mu} + (\mathbf{x}_t-\mu)\underbrace{\nabla_{\mathbf{x}_t}\log p_\phi(E|\mathbf{x}_t)\Big|_{\mathbf{x}_t=\mu}}_{\mathbf{g}} \Rightarrow$$

$$\log\left[p_\theta(\mathbf{x}_t|\mathbf{x}_{t+1})\,p_\phi(E|\mathbf{x}_t)\right] = -\frac{1}{2}(\mathbf{x}_t-\mu-\Sigma\mathbf{g})^T\,\Sigma(\mathbf{x}_t-\mu-\Sigma\mathbf{g})$$

- Training:

$$\hat{\varepsilon}(\mathbf{x}_t|E) \triangleq \varepsilon_\theta(\mathbf{x}_t) - \sqrt{1-\bar{\alpha}_t}\,\nabla_{\mathbf{x}_t}\log p_\phi(E|\mathbf{x}_t)$$

Classifier-free guidance

- Learning two separate models is computationally expensive
- We can learn both the conditional and unconditional models together as a unique conditional model
- The unconditional model can be queried by replacing the conditioning information, the energy, with fixed constant values, such as zero
- Greater control over the conditional generation procedure while requiring nothing beyond the training of a unique model

$$\left.\begin{aligned}\nabla_{\mathbf{x}_t} \log p\left(E \middle| \mathbf{x}_t\right) &= \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E\right) - \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t\right) + \\ \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E\right) &= \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t\right) + \gamma \nabla_{\mathbf{x}_t} \log p\left(E \middle| \mathbf{x}_t\right)\end{aligned}\right\} \Rightarrow \qquad \gamma > 1$$

$$\nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E\right) = \underbrace{\gamma \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E\right)}_{\text{conditional score}} + \underbrace{\left(1-\gamma\right) \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t\right)}_{\text{unconditional score}}$$

One neural network

- One extra label for the unconditional score

$$\nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E\right) = \underbrace{\gamma \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E\right)}_{\text{conditional score}} + \underbrace{\left(1-\gamma\right) \nabla_{\mathbf{x}_t} \log p\left(\mathbf{x}_t \middle| E=0\right)}_{\text{unconditional score}}$$

$$1-\gamma < 0$$

$$\hat{\varepsilon}_\theta\left(\mathbf{x}_t \middle| E\right) \quad \Bigg| \quad E \in \underbrace{\varnothing}_{E=0} \cup \left\{E_i\right\}_{i=1}^{N}$$

Conclusions

- The diffusion model not only prioritises the conditional score function but also moves in the direction away from the unconditional score function
- It reduces the probability of generating samples that do not use conditioning information, in favour of the samples that explicitly do
- Effect of decreasing sample diversity at the cost of generating samples that accurately match the conditioning information
- This is essentially performing random dropout of the conditioning information
- In short, business (almost) as usual!