条件扩散模型 (CONDITIONED DIFFUSION MODEL)

黄政宇
北京大学北京国际数学研究中心
北京大学国际机器学习研究中心

条件生成模型

➤ 应用：图像修复、去噪，蛋白质设计

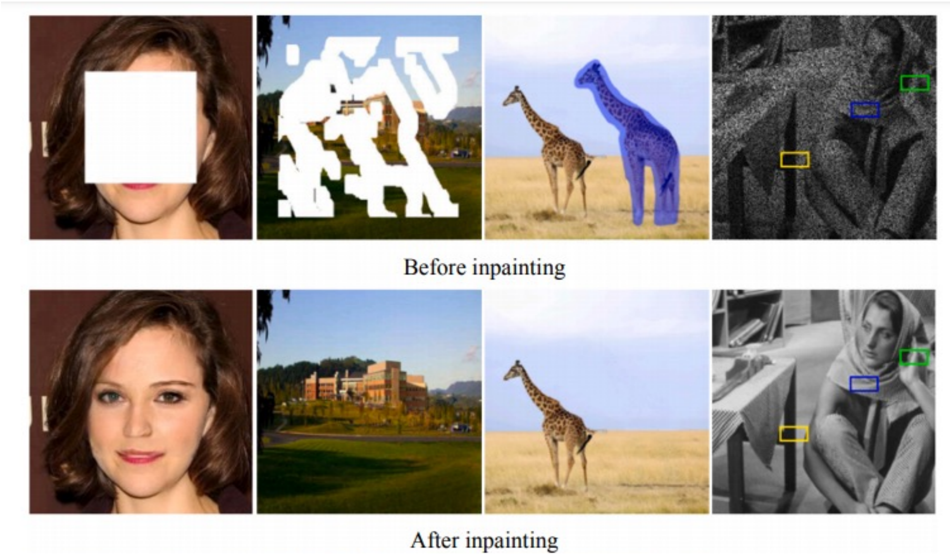

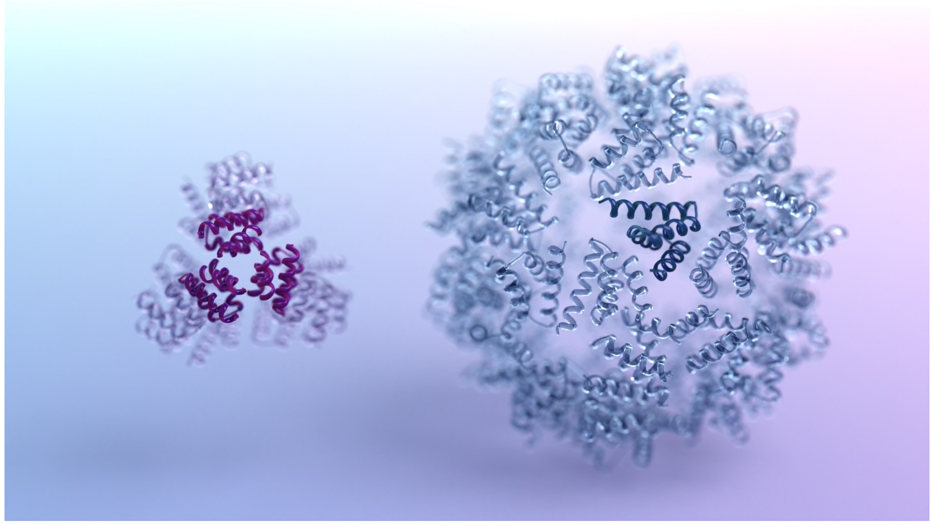

条件生成模型

- 扩散模型

$$q_0(x) \approx \frac{1}{N}\sum_i \delta(x - x^{*i})$$

生成$x \sim q_0(x)$

- 条件扩散模型

已知似然函数$q_0(y|x)$

目标：在扩散模型的基础上，$s_\theta(t,x) \approx \nabla \log q_{T-t}(x)$

生成 $x \sim q_0(x|y) \propto q_0(y|x)q_0(x)$

注：q_t在不同上下文中具有不同含义

条件生成模型

- 贝叶斯反问题

$$y = \mathcal{G}(x) + \eta$$

数据　模型　参数　数据噪音

后验分布：$q_0(x|y) \propto q_0(y|x) q_0(x)$

似然函数：$q_0(y|x) = \rho_\eta(y - \mathcal{G}(x))$

- 之前设置

$$q_0(x) = e^{-\Phi_0(x)}$$

新设置

$$q_0(x) \approx \frac{1}{N} \sum_i \delta(x - x^{*i})$$

条件生成模型

- 扩散后验采样 (diffusion posterior sampling)

假设 $X_0 \sim q_0(x)$，我们有前向和后向过程

$$dX_t = f(t)X_t dt + g(t)dW_t \ \ (0 \to T\text{，图片} \to \text{白噪音})$$

$$dY_t = \hat{f}(t, Y_t)dt + \hat{g}(t)dW_t \ \ (0 \to T\text{，白噪音} \to \text{图片})$$

我们有条件

$$y = \mathcal{G}(x) + \eta$$

且给定 X_0，X_t和 y 相互独立。按照贝叶斯观点，我们可以考虑X_t, y的条件概率$q_t(x|y)$，我们的目标是采样

$$q_0(x|y)$$

梯度诱导扩散模型

- 扩散后验采样(diffusion posterior sampling)

假设 $X_0 \sim q_0(x_0|y)$，直接使用生成模型

$$dX_t = f(t)X_t dt + g(t)dW_t$$

那么 $X_t \sim q_t(x_t|y) \propto q_t(y|x_t)q_t(x_t)$

OU后向过程条件生成

$$\partial_t \rho_t(x|y) = -\nabla \cdot [(x + \nabla \log q_{T-t}(x|y))\, \rho_t(x|y)]$$

贝叶斯法则

$$\nabla \log q_{T-t}(x|y) = \nabla \log q_{T-t}(y|x) + \nabla \log q_{T-t}(x)$$

梯度诱导项　　分数函数

梯度诱导扩散模型

- Tweedie 公式近似

似然函数

$$\nabla \log q_t(y|x) = \nabla \log \int q_{t0}(x_0|x) q_0(y|x_0) dx_0$$

$$\approx \nabla \log q_0(y|\hat{x}_0(t,x))$$

其中

$$q_{0t}(x|x_0) = \mathcal{N}(x; \lambda_t x_0, \sigma_t^2 I)$$

做如下近似

$$q_{t0}(x_0|x) \approx \delta(x_0 - \hat{x}_0(t,x))$$

$$\hat{x}_0(t,x) = \mathbb{E}_{q_{t0}(x_0|x)}[x_0] = \frac{1}{\lambda_t}(x + \sigma_t^2 \nabla \log q_t(x))$$

- 诱导导数计算

$$\nabla \log q_0(y|\hat{x}_0(t,x)) = \nabla_x \log \rho_\eta(y - \mathcal{G}(\hat{x}_0(T-t,x)))$$

梯度诱导扩散模型

- 扩散后验采样

后向过程条件生成 $(0 \to T)$

$$\partial_t \rho_t(x|y)$$
$$= -\nabla \cdot [(x + \nabla \log q_{T-t}(x) + \nabla \log q_0(y|\hat{x}_0(t,x)))\, \rho_t(x|y)]$$

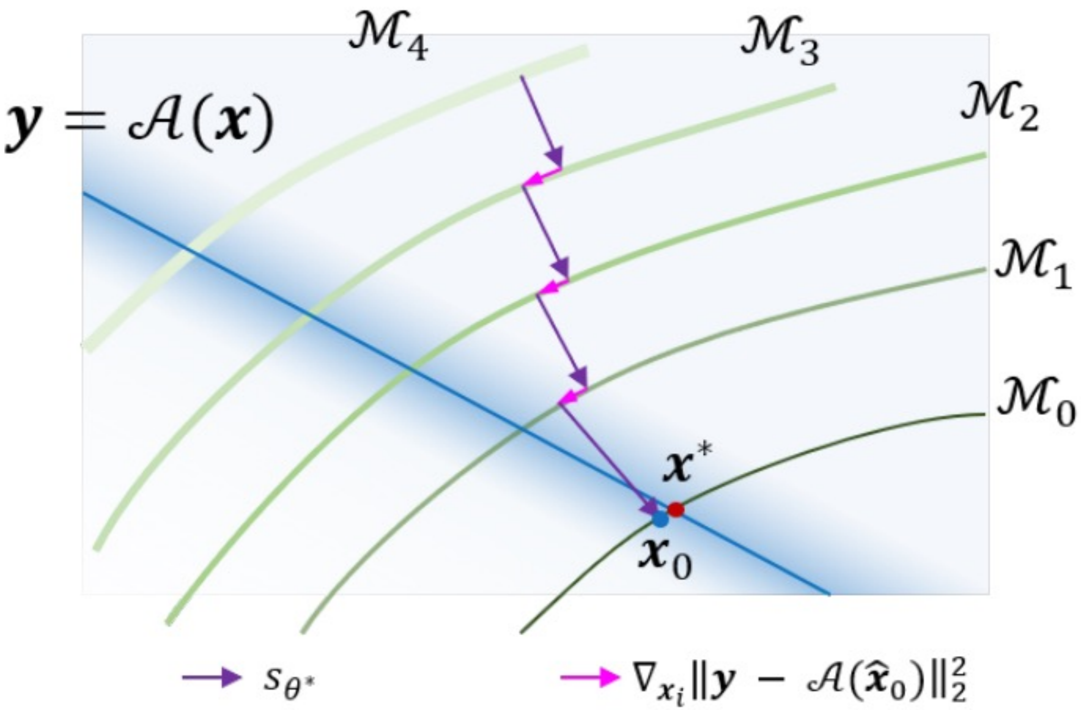

梯度诱导扩散模型

- 扩散后验采样

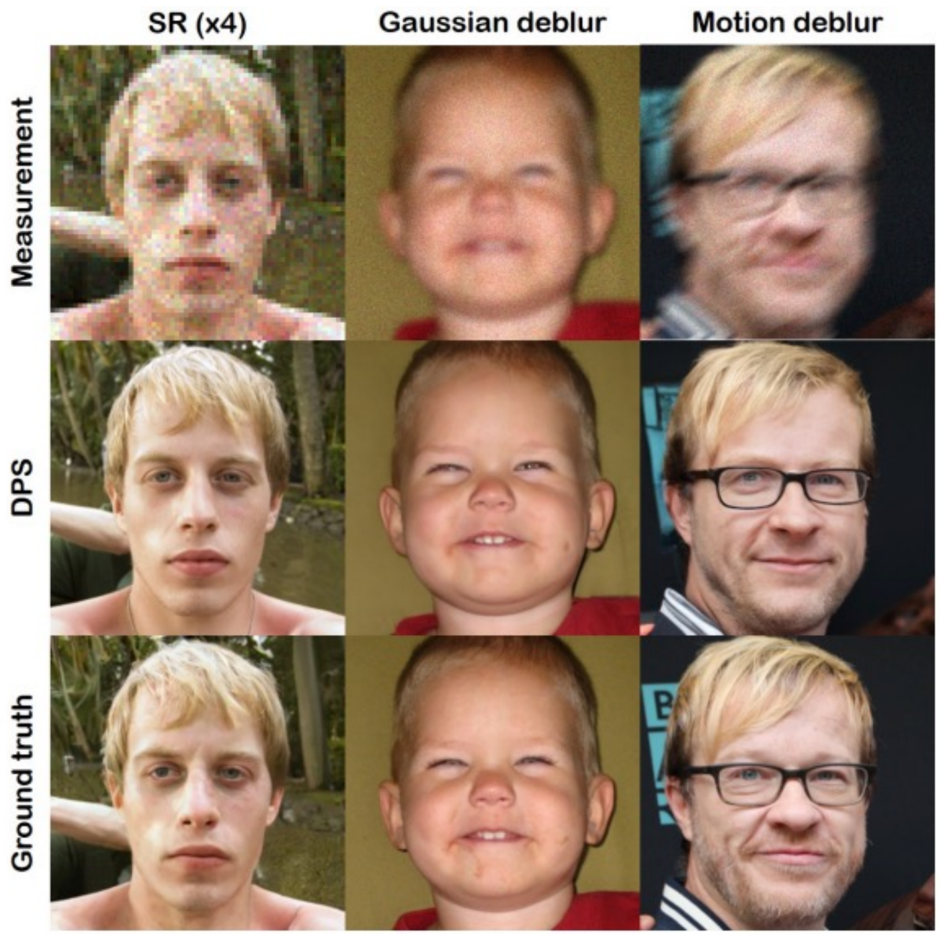

能不能修正
Tweedie 公式
近似的误差，
达到渐近精确
(asymptotically
exact)？

权重修正！

Chung, Hyungjin, et al. "Diffusion posterior sampling for general noisy inverse problems." *arXiv preprint arXiv:2209.14687* (2022)

序贯蒙特卡洛扩散模型

- 扩散后验采样(diffusion posterior sampling)

假设 $X_0 \sim q_0(x)$，我们有前向和后向过程

$$dX_t = f(t)X_t dt + g(t)dW_t \ (0 \to T\text{，图片} \to \text{白噪音})$$

$$dY_t = \hat{f}(t, Y_t)dt + \hat{g}(t)dW_t \ (0 \to T\text{，白噪音} \to \text{图片})$$

我们有条件

$$y = \mathcal{G}(x) + \eta$$

考虑时间离散($t = 0,1,\cdots,T$)的条件分布

$$q(x_{0:T}|y) = \frac{q_0(y|x_0)q_T(x_T)\prod_{t=0}^{T-1} q_{t+1t}(x_t|x_{t+1})}{q(y)}$$

序贯蒙特卡洛扩散模型

- 重要性采样

考虑时间离散$(t=0,1,\cdots,T)$的条件分布

$$q(x_{0:T}|y)=\frac{q(y|x_0)q_T(x_T)\prod_{t=0}^{T-1}q_{t+1t}(x_t|x_{t+1})}{q(y)}$$

扩散模型

$$X_T\sim q_T(x_T)$$

$$X_t|X_{t+1}\sim q_{t+1t}(x_t|X_{t+1})\approx\mathcal{N}\left(x_t\middle|m_{t+1|t},\sigma_{t+1|t}^2\right)$$

$$t=T-1,\cdots,0$$

OU过程： $m_{t+1|t}=X_{t+1}+\left(X_{t+1}+2s_\theta(t+1,X_{t+1})\right),$

$\sigma_{t+1|t}^2=2$

计算权重

权重退化为0！ $w(X_{0:T})\propto q(y|X_0)$

序贯蒙特卡洛扩散模型

- 最优 (optimal)序贯蒙特卡洛

考虑时间离散($t = 0,1,\cdots,T$)的条件分布

$$q(x_{0:T}|y) = \frac{q(y|x_0)q_T(x_T)\prod_{t=0}^{T-1} q_{t+1t}(x_t|x_{t+1})}{q(y)}$$

$$= q_T(x_T|y)\prod_{t=0}^{T-1} q_{t+1t}(x_t|x_{t+1},y)$$

采样

$$X_T \sim q_T(x_T|y) \propto q_T(y|x_T)q_T(x_T)$$

$$X_t|X_{t+1} \sim q_{t+1t}(x_t|X_{t+1},y) \propto q_t(y|x_t)q_{t+1t}(x_t|X_{t+1})$$

$$t = T-1,\cdots,0$$

不需要权重！

序贯蒙特卡洛扩散模型

- Tweedie 公式近似

似然函数

$$q_t(y|x) \propto \int q_{t0}(x_0|x) q_0(y|x_0) dx_0$$

$$\approx q_0(y|\hat{x}_0(t,x))$$

其中

$$q_{0t}(x|x_0) = \mathcal{N}(x; \lambda_t x_0, \sigma_t^2 I)$$

做如下近似

$$q_{t0}(x_0|x) \approx \delta(x_0 - \hat{x}_0(t,x))$$

$$\hat{x}_0(t,x) = \mathbb{E}_{q_{t0}(x_0|x)}[x_0] = \frac{1}{\lambda_t}(x + \sigma_t^2 \nabla \log q_t(x))$$

序贯蒙特卡洛扩散模型

➤ 序贯蒙特卡洛

采样 $j = 1, \cdots, J$

$$X_T^j \sim X_T \sim q_T(x_T), \qquad w_T^j \propto q_0(y|\hat{x}_0(T, X_T))$$

对于 $t = T-1, \cdots, 0$ （隐去上标j）

$$X_t|X_{t+1} \sim \tilde{q}_{t+1,t}(x_t|X_{t+1}, y) \approx q_0(y|\hat{x}_0(t+1, x))q_{t+1t}(x|X_{t+1})$$

高斯提议核(OU过程)：

$$\tilde{q}_{t+1,t}(x|X_{t+1}, y) \approx \mathcal{N}\Big(x, X_{t+1} + 2\big(\nabla_x \log q_0(y|\hat{x}_0(t+1, X_{t+1})) + X_{t+1} + s_\theta(t+1, X_{t+1})\big), 2\Big)$$

梯度诱导项

权重：$w_t \propto \dfrac{q_0(y|\hat{x}_0(t, X_t))q_{t+1t}(X_t|X_{t+1})}{q_0(y|\hat{x}_0(t+1, X_{t+1}))\,\tilde{q}_{t+1,t}(X_t|X_{t+1}, y)}$

渐近精确：

$$\left\{X_{0:T}, \prod_{t=0}^{T-1} w_t\right\} \sim q_T(x_T)q_0(y|\hat{x}_0(0, x))\prod_{t=0}^{T-1} q_{t+1t}(x_t|X_{t+1})$$

序贯蒙特卡洛扩散模型

- 序贯蒙特卡洛

采样 $j = 1, \cdots, J$

$$X_T^j \sim X_T \sim q_T(x_T), \qquad w_T^j \propto q_0(y|\hat{x}_0(T, X_T))$$

对于 $t = T-1, \cdots, 0$ （隐去上标j）

$$X_t|X_{t+1} \sim \tilde{q}_{t+1,t}(x_t|X_{t+1}, y) \approx q_0(y|\hat{x}_0(t+1, x))q_{t+1t}(x|X_{t+1})$$

高斯提议核（OU过程）：

$$\tilde{q}_{t+1,t}(x|X_{t+1}, y)$$
$$\approx \mathcal{N}\left(x, X_{t+1} + 2\left(\nabla_x \log q_0(y|\hat{x}_0(t+1, X_{t+1})) + X_{t+1} + s_\theta(t+1, X_{t+1})\right), 2\right)$$

梯度诱导项

权重：$w_t \propto \dfrac{q_0(y|\hat{x}_0(t, X_t))q_{t+1t}(X_t|X_{t+1})}{q_0(y|\hat{x}_0(t+1, X_{t+1}))\,\tilde{q}_{t+1,t}(X_t|X_{t+1}, y)}$

渐近精确：

$$\left\{X_{0:T}, \prod_{t=0}^{T-1} w_t\right\} \sim q_T(x_T)q_0(y|\hat{x}_0(0, x))\prod_{t=0}^{T-1} q_{t+1t}(x_t|X_{t+1})$$

序贯蒙特卡洛扩散模型

➢ 序贯蒙特卡洛

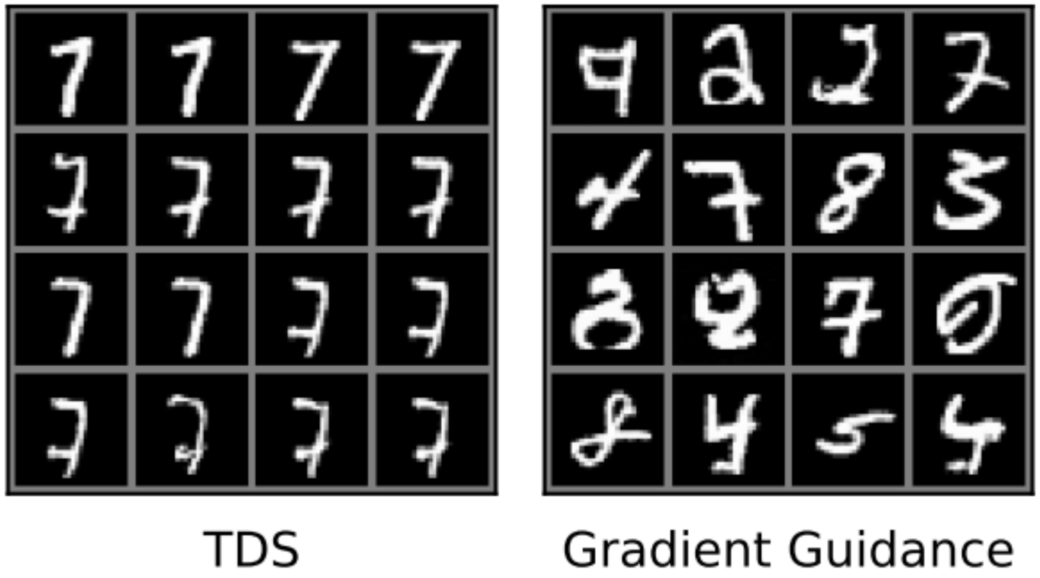

Wu, Luhuan, et al. "Practical and asymptotically exact conditional sampling in diffusion models." *Advances in Neural Information Processing Systems* 36 (2024).

扩展阅读

- 梯度诱导扩散模型

  Chung, Hyungjin, et al. "Diffusion posterior sampling for general noisy inverse problems." *arXiv preprint arXiv:2209.14687* (2022)

- 序贯蒙特卡洛扩散模型

  Wu, Luhuan, et al. "Practical and asymptotically exact conditional sampling in diffusion models." *Advances in Neural Information Processing Systems* 36 (2024).