Conditional diffusion models – The equation festival is over… *for the moment*

Eric Paquet

NRC

21 February 2025

The condition should be *hammered* at each time step – no cat was harmed with this algorithm

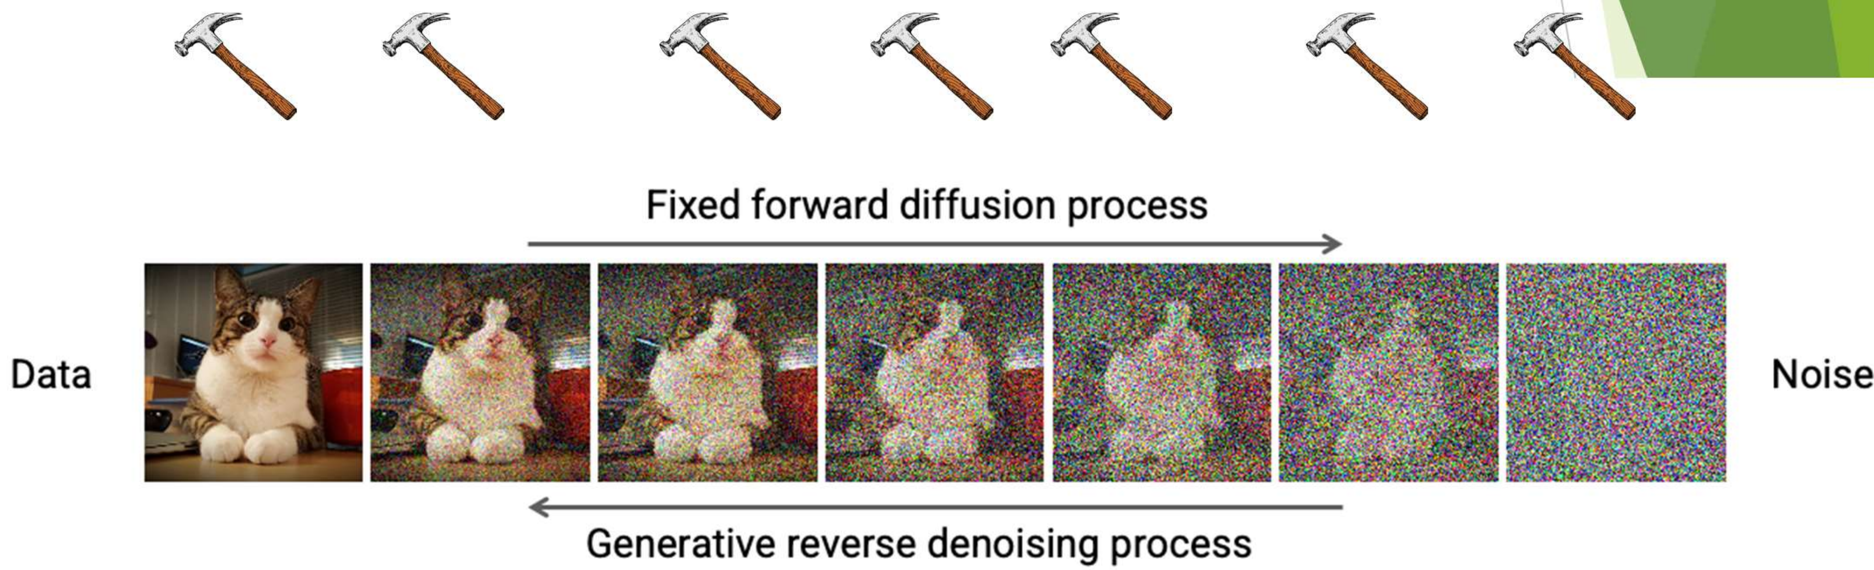

Replacement method

- Part of the input bandwidth is occupied by the encoded condition.
- The condition must be systematically hammered at each time step.
- When training the network, the noise of the condition should match the noise of the data for the corresponding time step.
- When generating data, the condition should be reintroduced at each time step with the correct level of noise.

Replacement method for approximate conditional sampling

Input: $\mathbf{x}_{\mathcal{M}}^{(0)}$ (motif)
// Forward diffuse motif
$\breve{\mathbf{x}}_{\mathcal{M}}^{(1:T)} \sim q(\mathbf{x}_{\mathcal{M}}^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)})$

// Reverse diffuse scaffold
$\mathbf{x}^{(T)} \sim p_\theta(\mathbf{x}^{(T)})$
for $t = T, \ldots, 1$ **do**
    // *Replace* with forward diffused motif
    $\mathbf{x}^{(t)} \leftarrow [\breve{\mathbf{x}}_{\mathcal{M}}^{(t)}, \mathbf{x}_{\mathcal{S}}^{(t)}]$

    // Propose next step
    $\mathbf{x}^{(t-1)} \sim p_\theta(\mathbf{x}^{(t-1)} \mid \mathbf{x}^{(t)})$
end for
Return $\mathbf{x}_{\mathcal{S}}^{(0)}$, $\mathbf{x}^{(1:T)}$

The replacement method is not the end of the story…

- It introduces irreducible errors that cannot be eliminated by making the denoising conditional density more expressive.
- https://arxiv.org/abs/2206.04119
- The approximation error is not tractable in general.
- But it is still the most common approach.

The solution: particle filtering

- In the spirit of a Monte Carlo method, generate K instances for each time step where the condition is enforced by the replacement method.
- Compute the unnormalised probabilities (weights) of the condition given the noisy data with replacement.
- Normalise the probabilities.
- Sample an outcome using the normalised probabilities (from the multinomial distribution).
- Propose the next step.

Particle filtering for conditionally sampling from unconditional diffusion models

Input: $\mathbf{x}_{\mathcal{M}}^{(0)}$ (motif), K (# particles)
// Forward diffuse motif
$\breve{\mathbf{x}}_{\mathcal{M}}^{(1:T)} \sim q(\mathbf{x}_{\mathcal{M}}^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)})$

// Reverse diffuse particles
$\forall k,\ \mathbf{x}_k^{(T)} \overset{i.i.d.}{\sim} p_\theta(\mathbf{x}^{(T)})$
for $t = T, \ldots, 1$ **do**
  // *Replace* motif
  $\forall k,\ \mathbf{x}_k^{(t)} \leftarrow [\breve{\mathbf{x}}_{\mathcal{M}}^{(t)}, \mathbf{x}_{\mathcal{S},k}^{(t)}]$

  // Re-weight based on $\breve{\mathbf{x}}_{\mathcal{M}}^{(t-1)}$
  $\forall k,\ w_k^{(t)} \leftarrow p_\theta(\breve{\mathbf{x}}_{\mathcal{M}}^{(t-1)} \mid \mathbf{x}_k^{(t)})$
  $\forall k,\ \tilde{w}_k^{(t)} \leftarrow w_k^{(t)} / \sum_{k'=1}^{K} w_{k'}^{(t)}$
  $\tilde{\mathbf{x}}_{1:K}^{(t)} \sim \texttt{Resample}(\tilde{w}_{1:K}^{(t)}, \mathbf{x}_{1:K}^{(t)})$

  // Propose next step
  $\forall k,\ \mathbf{x}_k^{(t-1)} \overset{indep.}{\sim} p_\theta(\mathbf{x}^{(t-1)} \mid \tilde{\mathbf{x}}_k^{(t)})$
end for
Return $\mathbf{x}_{\mathcal{S},1:K}^{(0)}$

But don’t worry…

- The equations festival will be back next week!