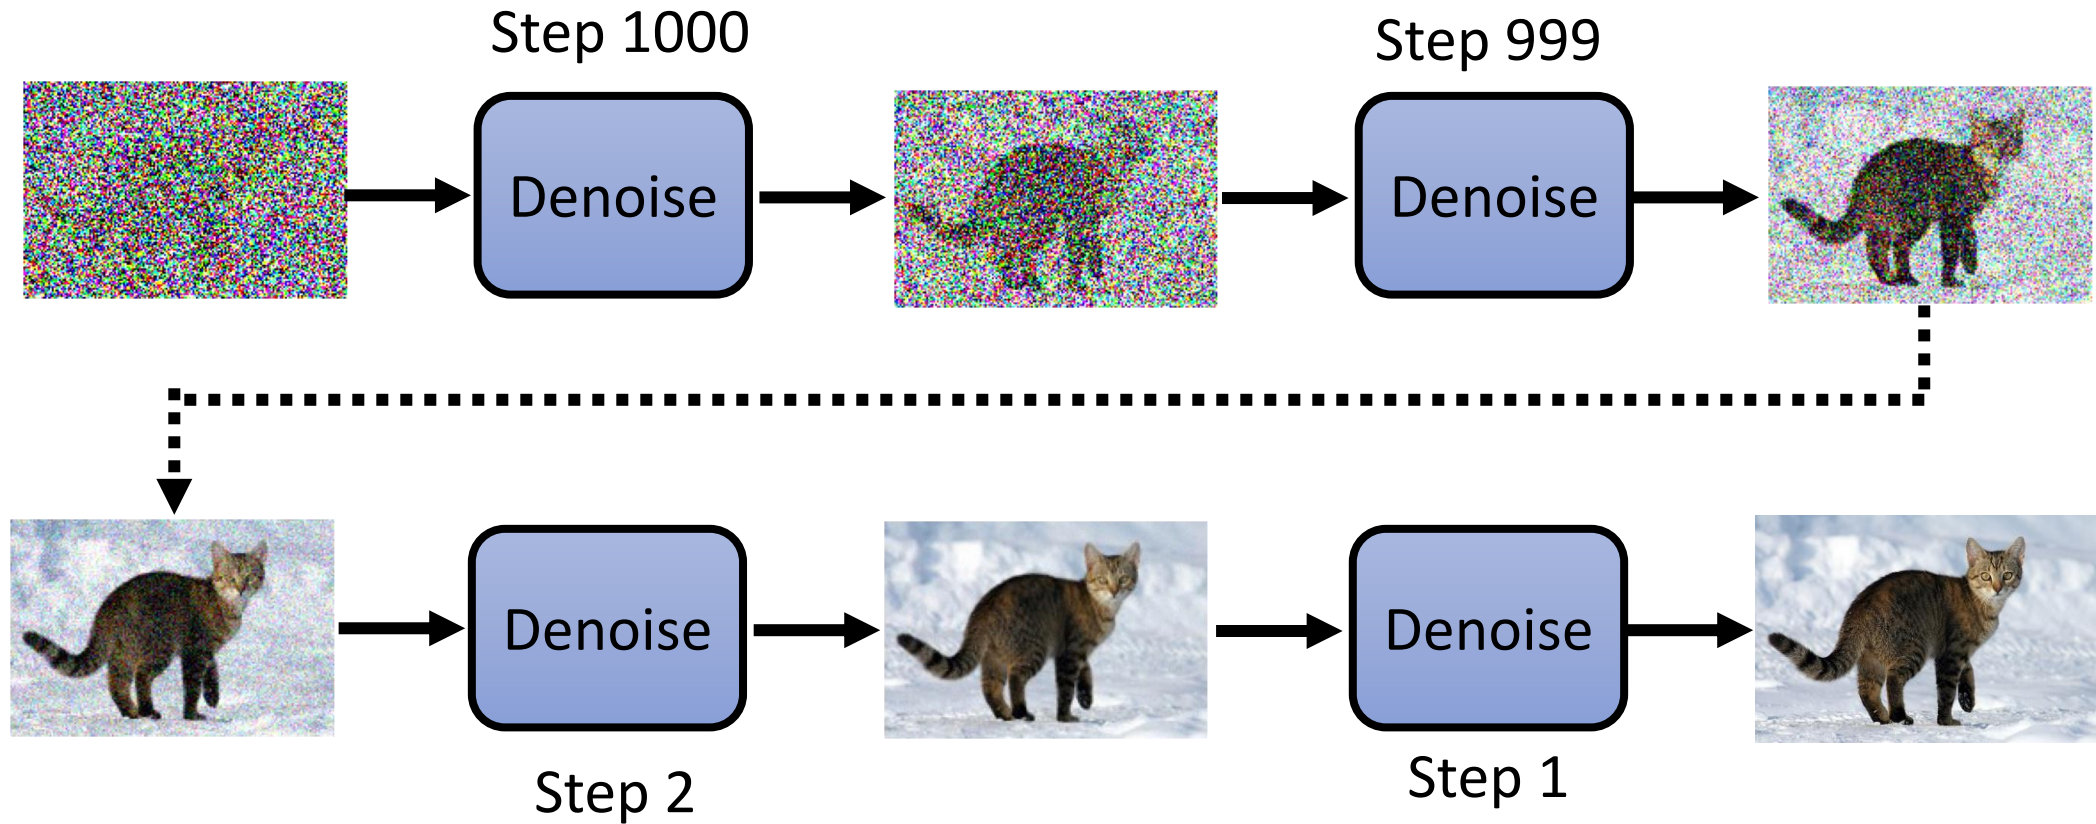


Diffusion Model

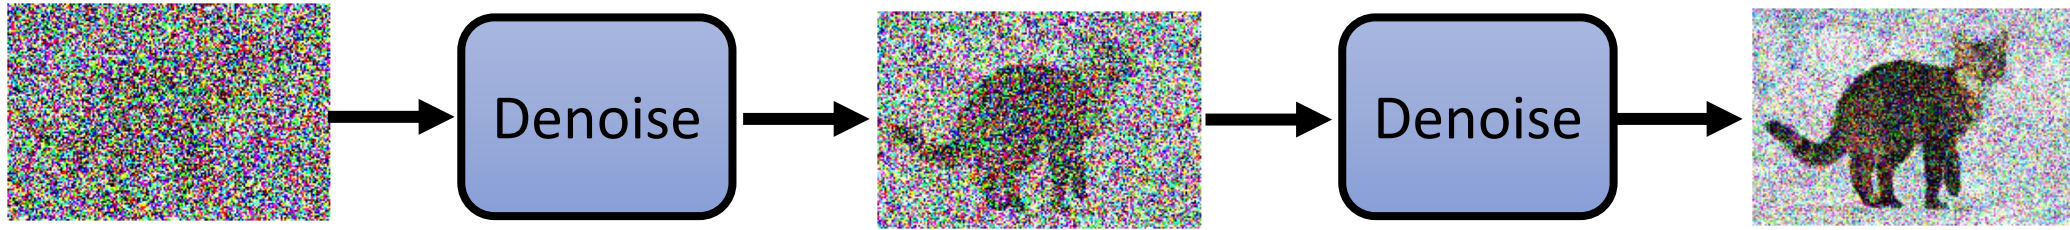
Denoising Diffusion Probabilistic Models (DDPM)

<https://arxiv.org/abs/2006.11239>

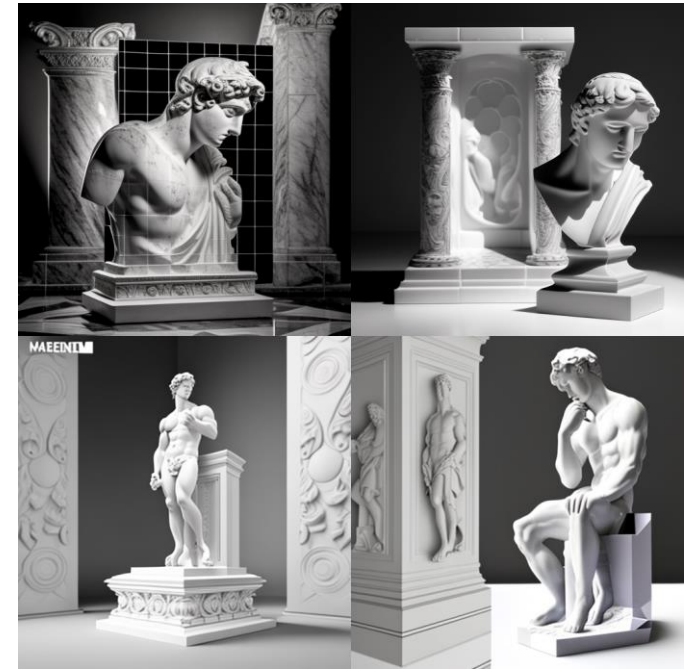
Diffusion Model 是如何運作的？

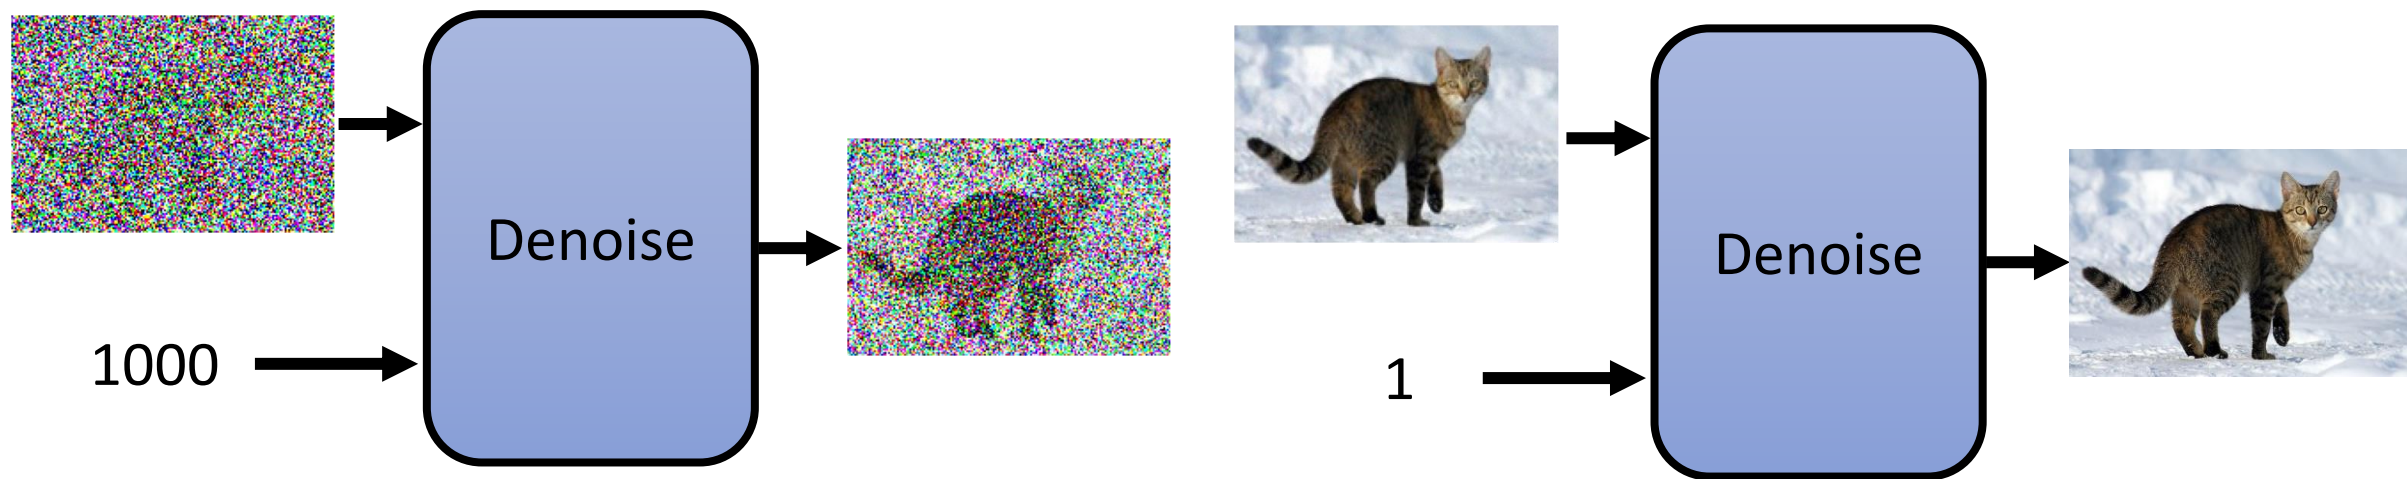
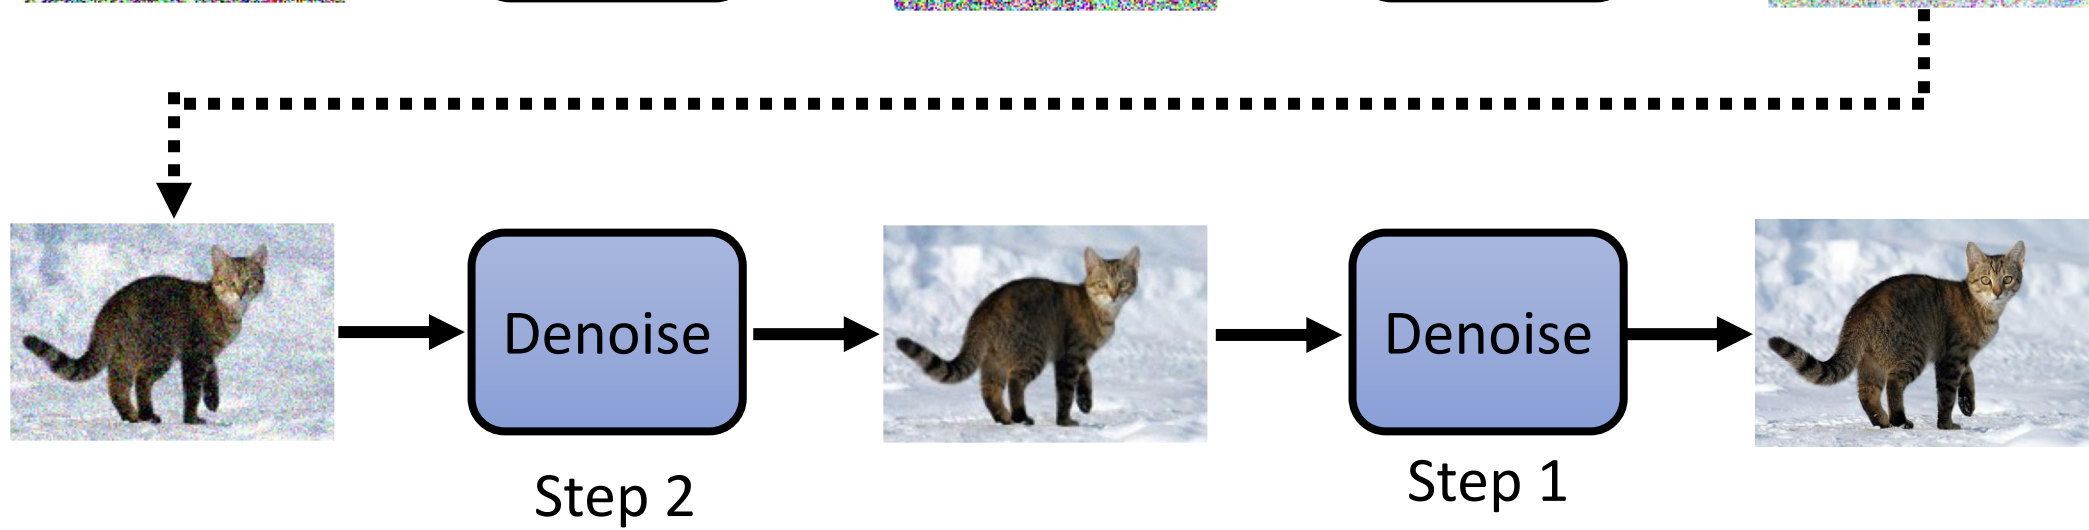
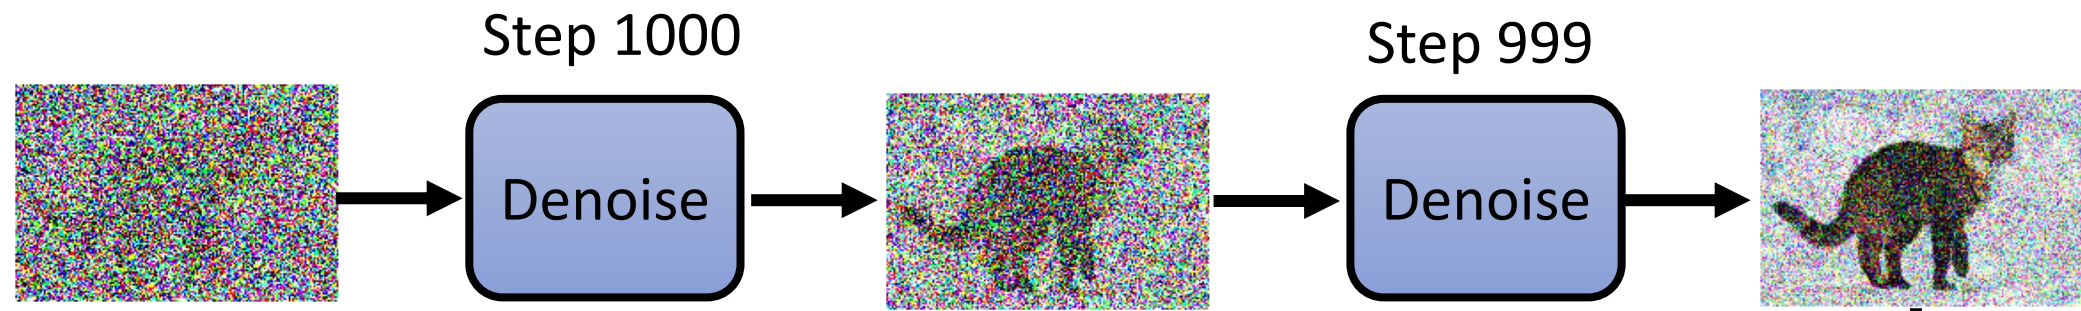


Reverse Process

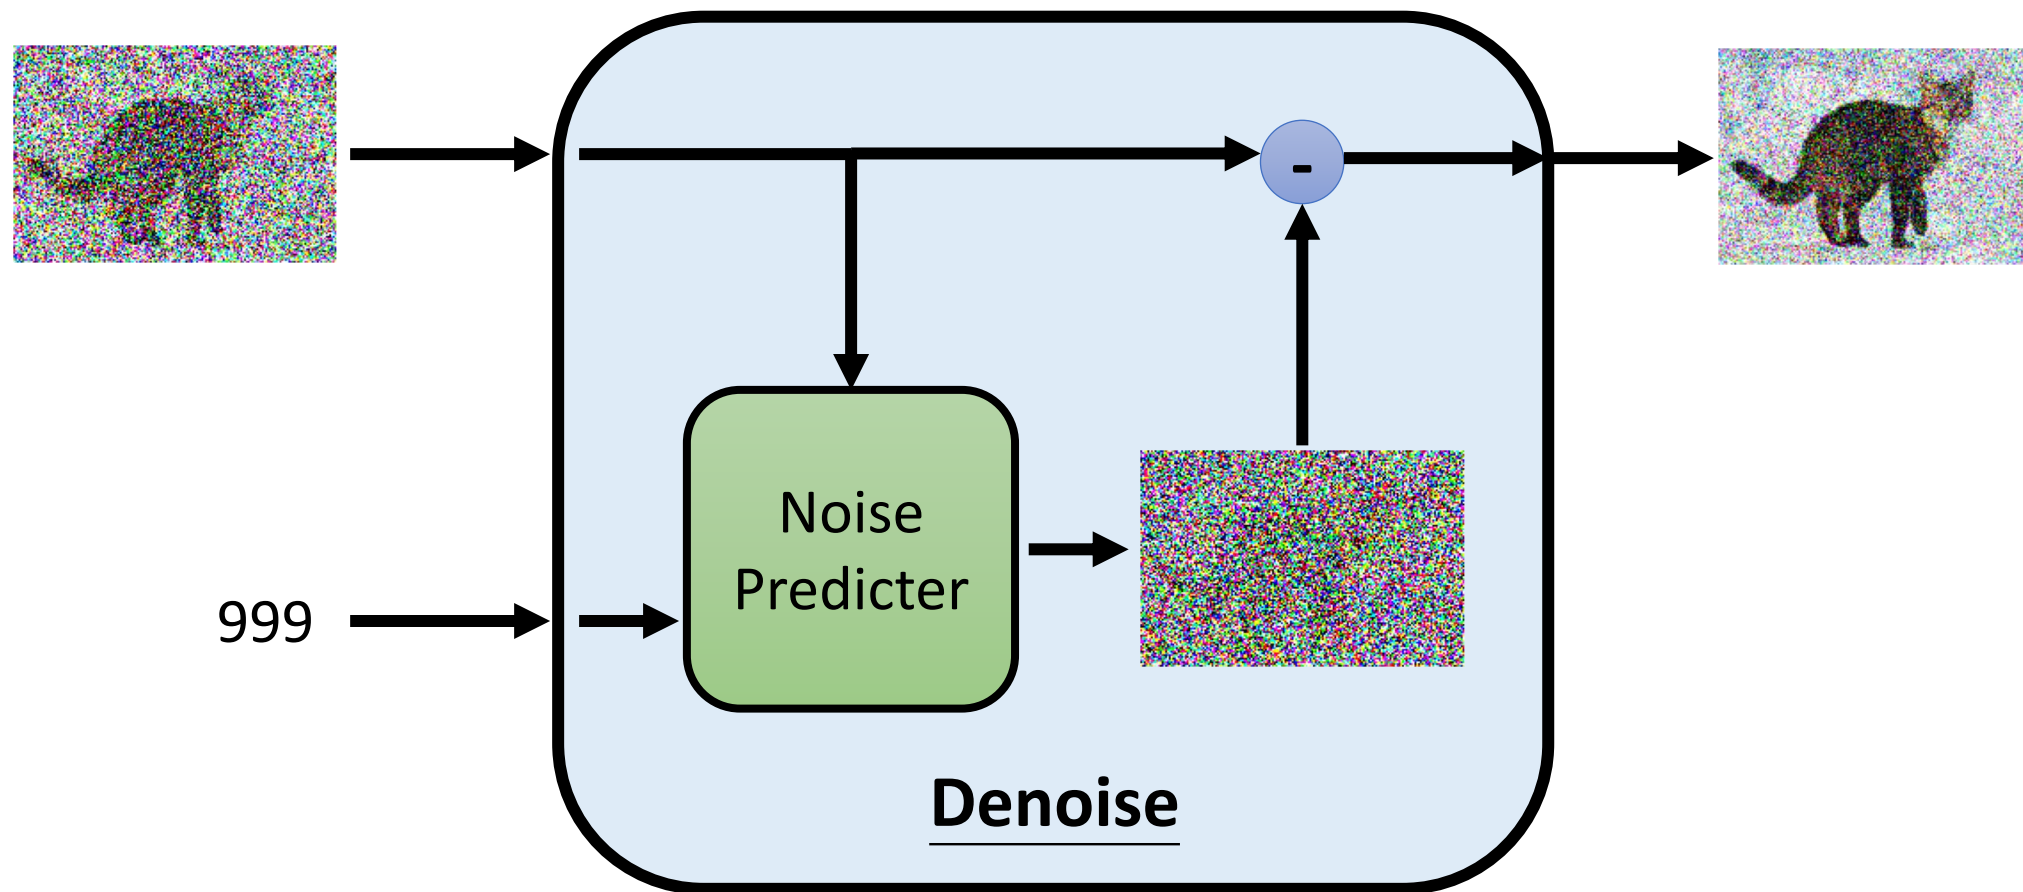


The sculpture is already complete within the marble block, before I start my work. It is already there, I just have to chisel away the superfluous material. - Michelangelo

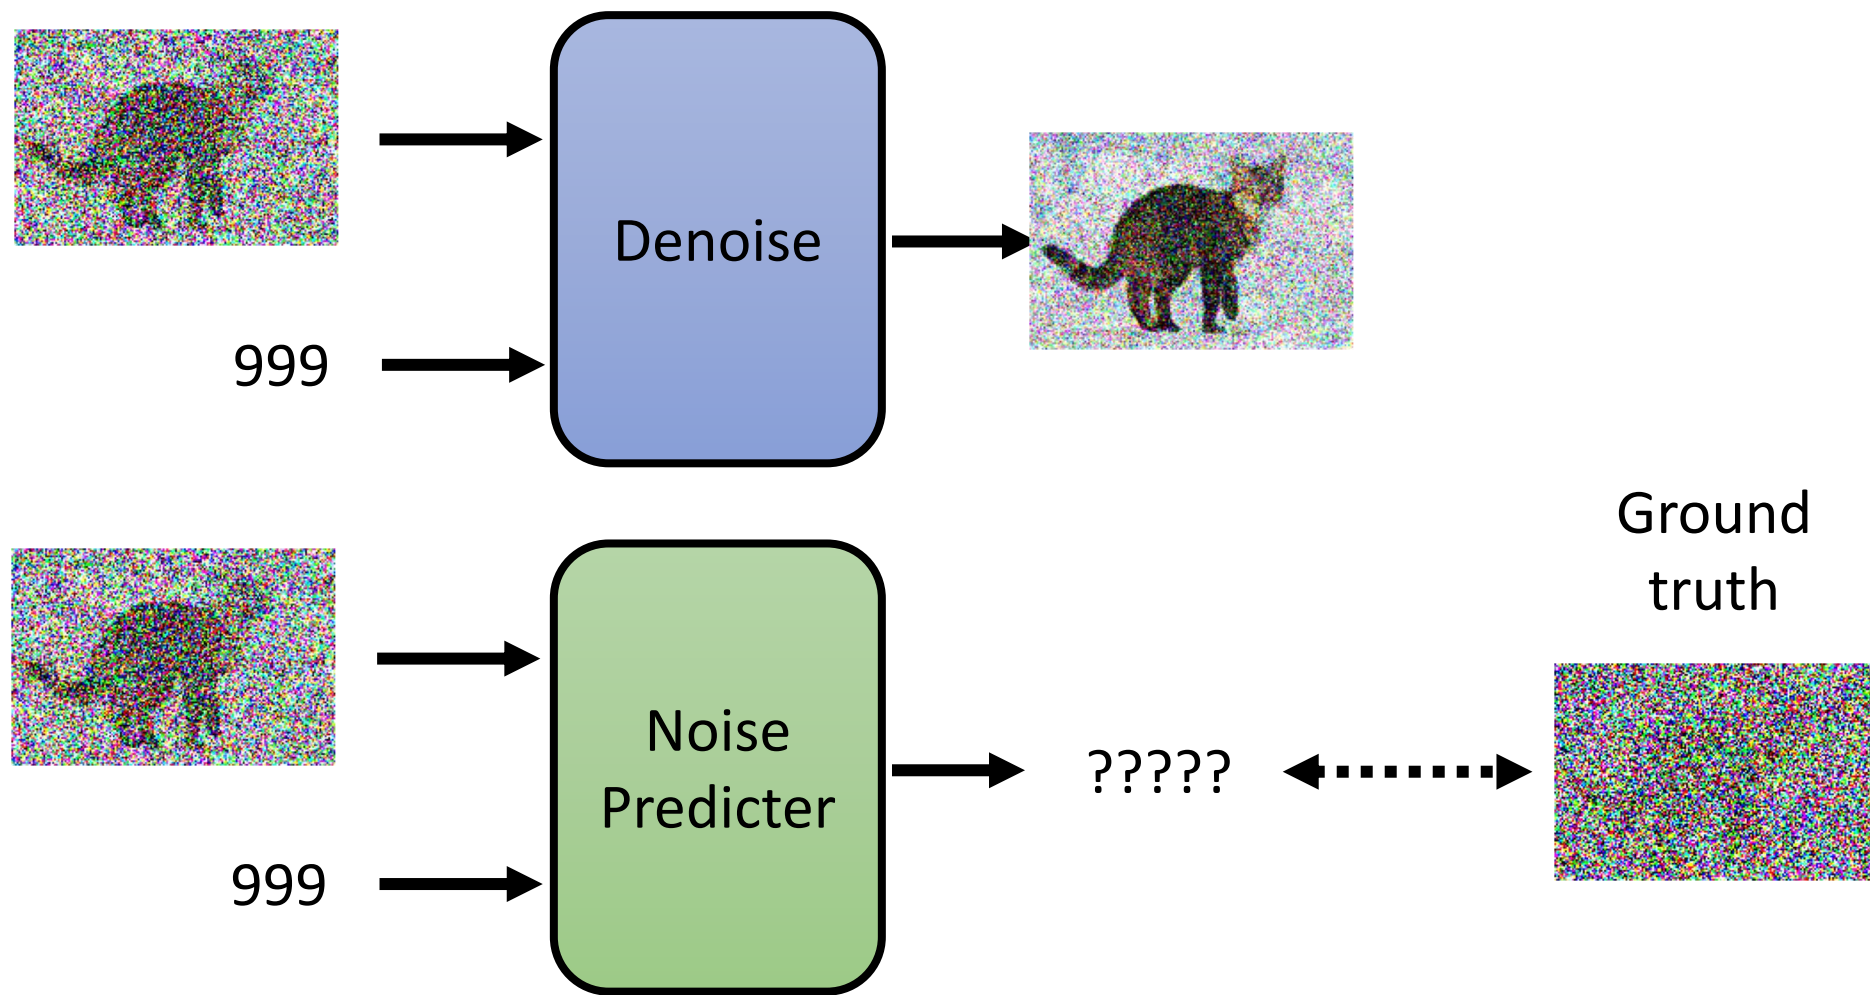


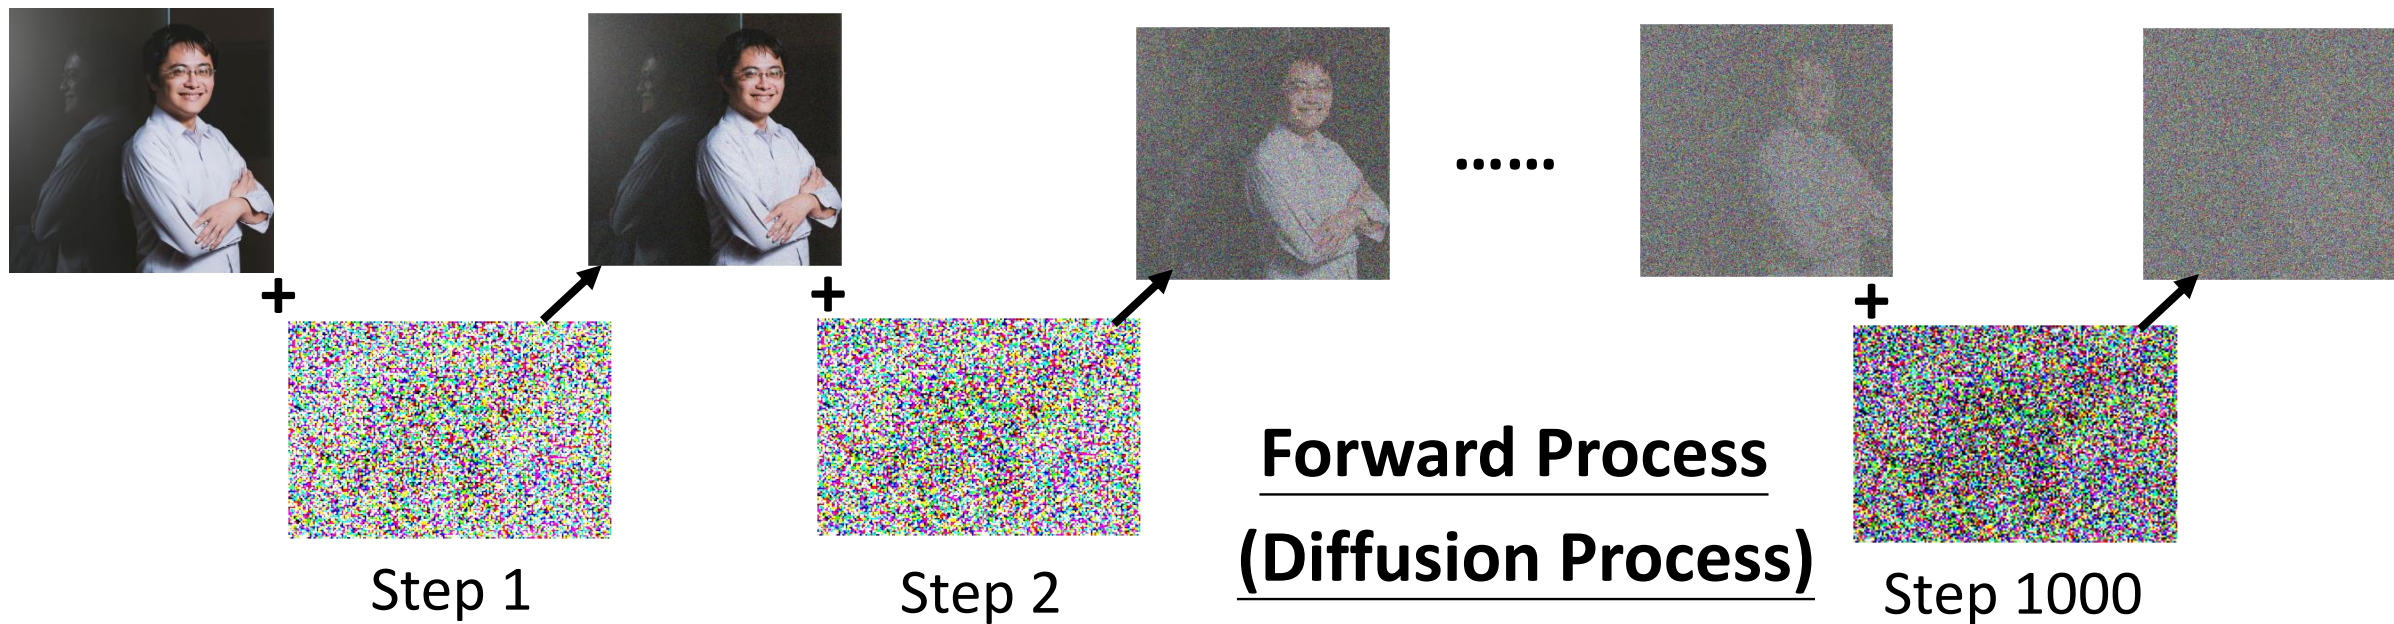
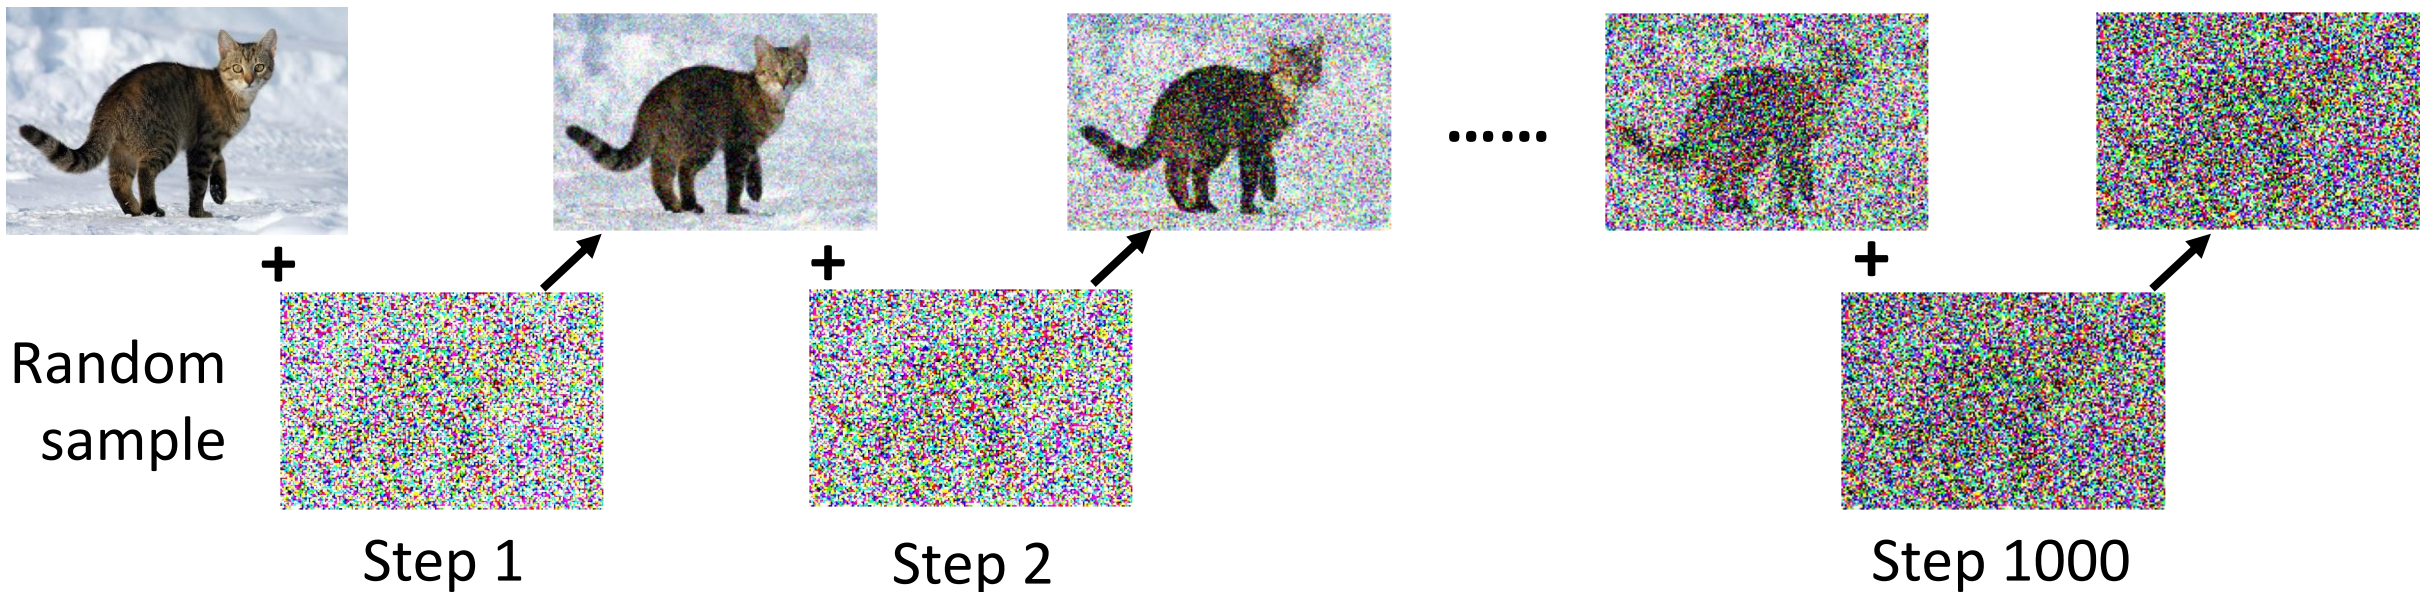


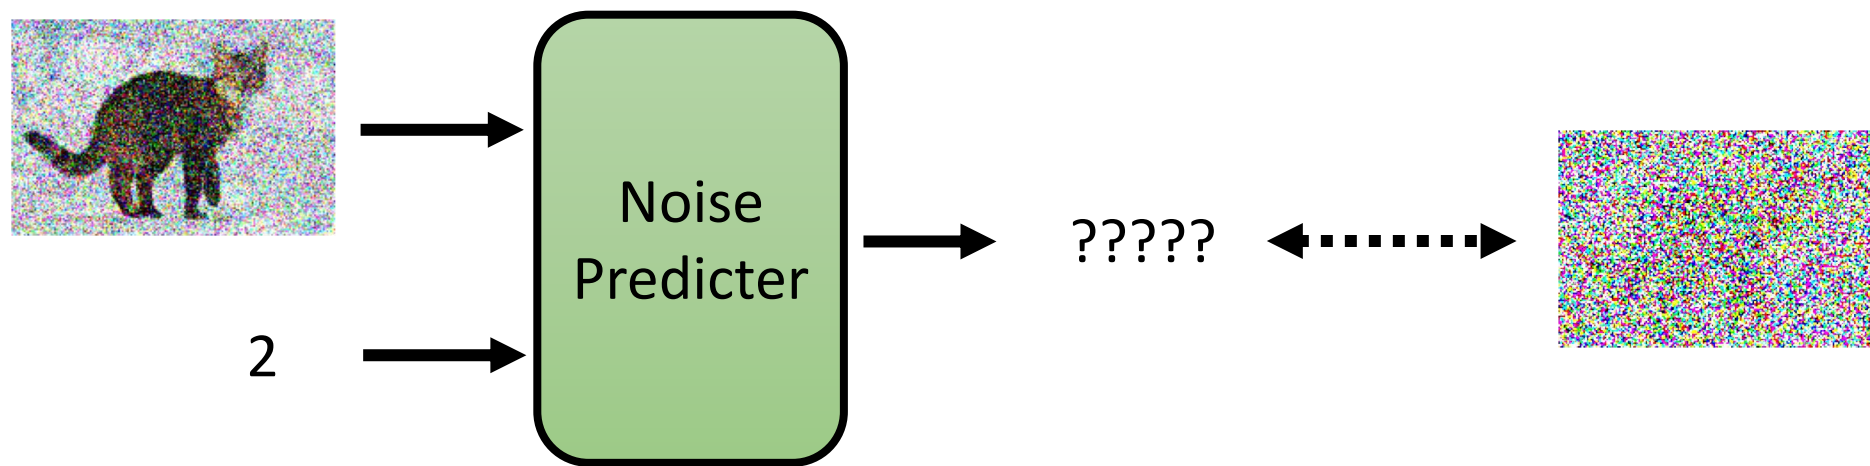
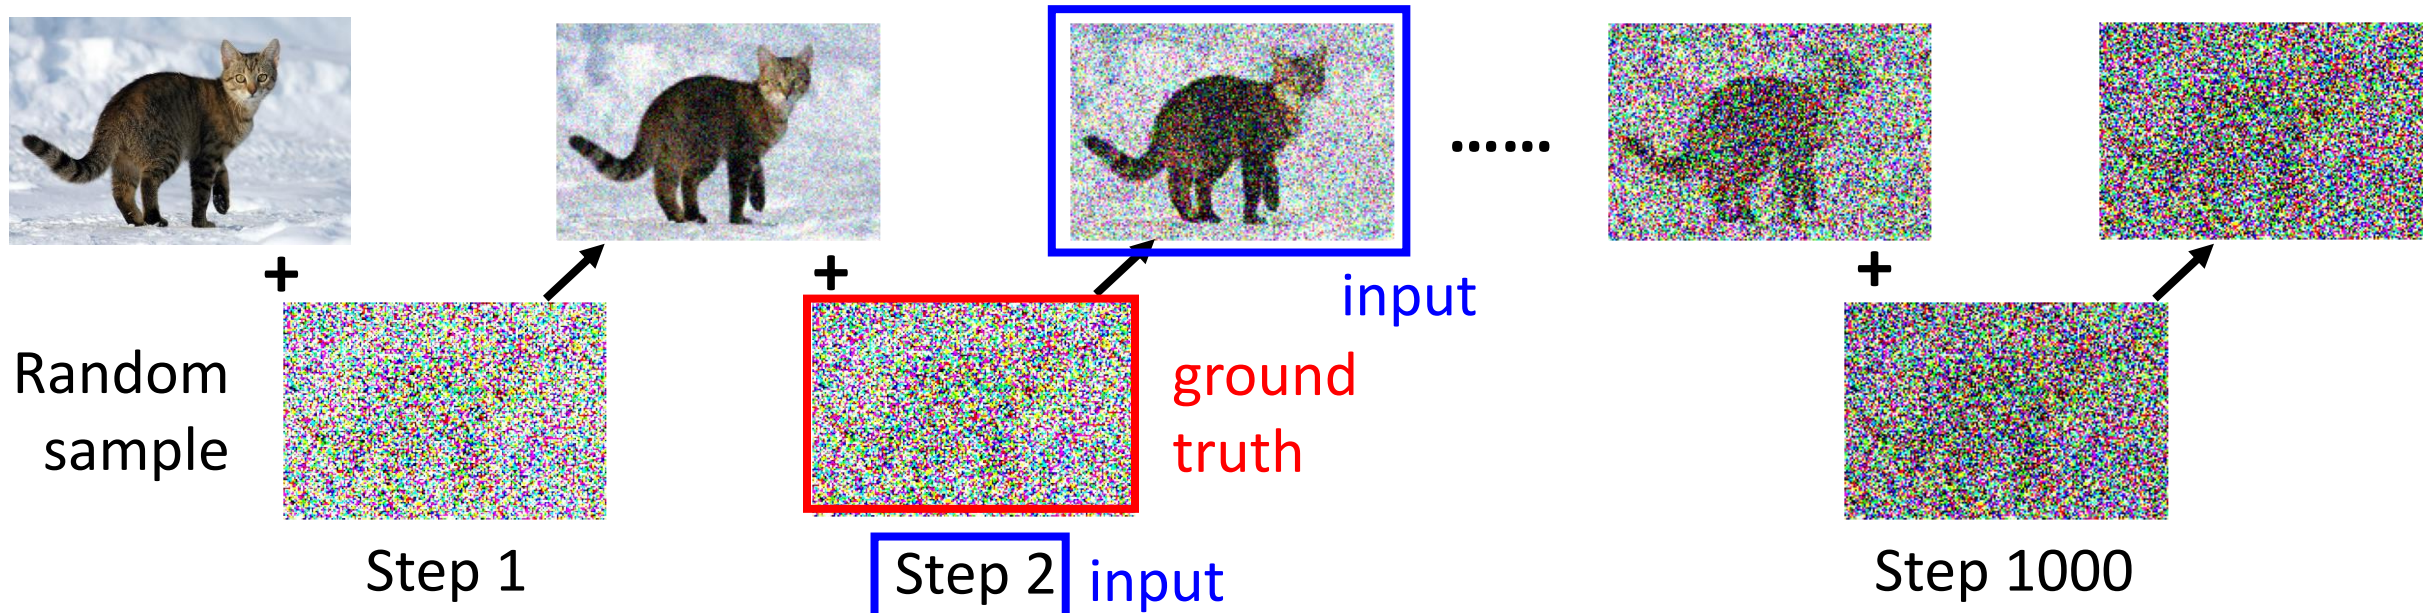
Denoise 模組內部實際做的事情



如何訓練 Noise Predictor







Text-to-Image

<https://laion.ai/blog/laion-5b/>

A cat in
the snow



Text-to-image
Generator



HW6

70k



ImageNet

1M



LAION

5.85B

Backend url:

<https://knn5.laion>

Index:

laion_5B

french cat

[Clip retrieval](#) works by converting the text query to a CLIP embedding, then using that embedding to query a knn index of clip image embeddings

Display captions

Display full captions

Display similarities

Safe mode

Hide duplicate urls

Hide (near) duplicate images

Search over

image

Search with multilingual clip



french cat



french cat



How to tell if your feline is french. He wears a b...



Winter cat



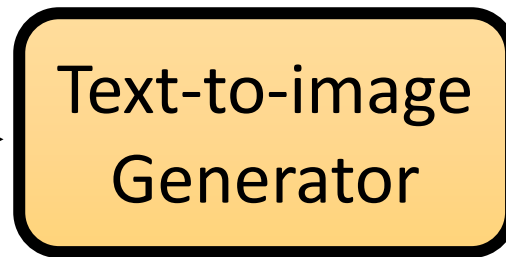
網友挑戰「加幾筆畫出最創意貓咪圖片」，第到份年之後我出



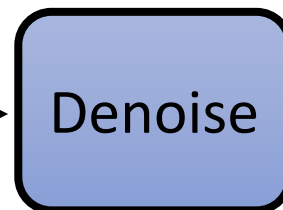
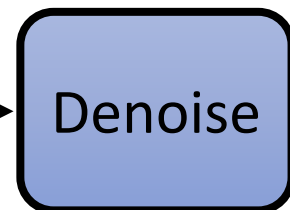
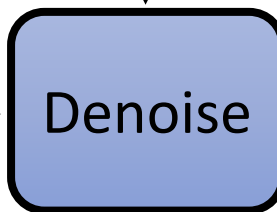
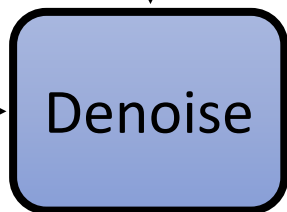
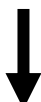
cat in a suit Georgian sells tomatoes

Text-to-Image

A cat in the snow



A cat in the snow



A cat in the snow

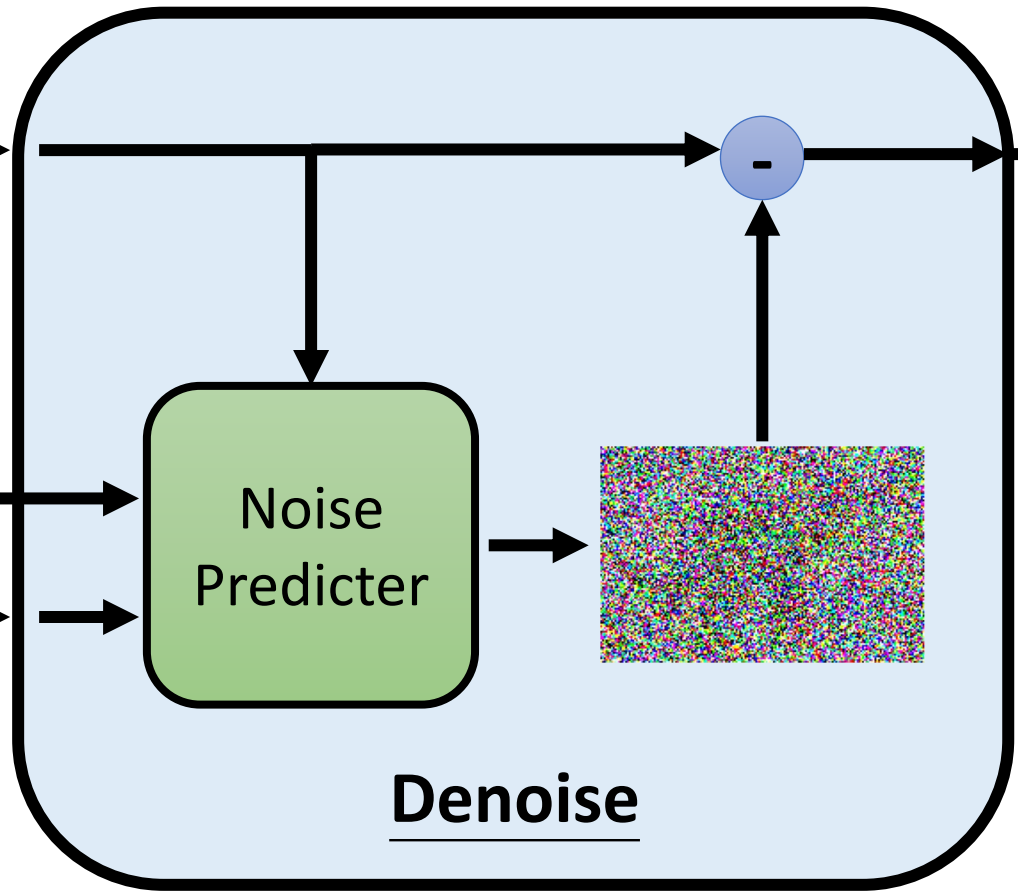
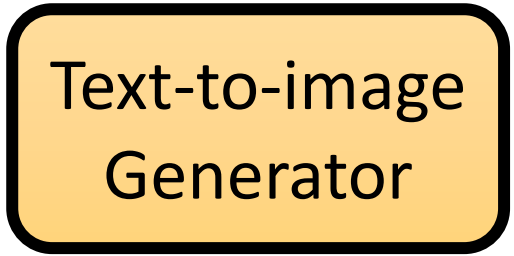


A cat in the snow



Text-to-Image

A cat in
the snow



A cat in the snow

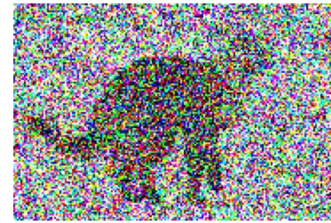
999

Noise
Predicter

Denoise

A cat in the snow

input

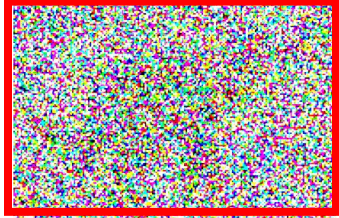
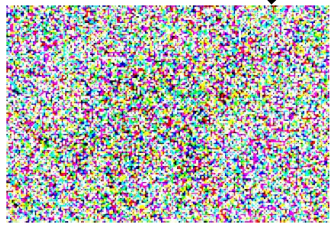


+

+

+

Random sample



Step 1

Step 2 input

Step 1000

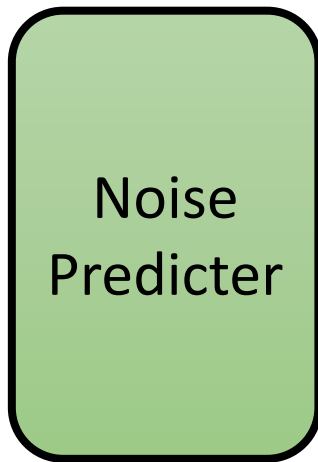
input

ground truth

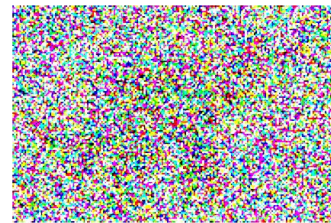


A cat in the snow

2



?????



Denoising Diffusion Probabilistic Models

Algorithm 1 Training

- 1: **repeat**
 - 2: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
 - 3: $t \sim \text{Uniform}(\{1, \dots, T\})$
 - 4: $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 5: Take gradient descent step on
$$\nabla_{\theta} \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\theta}(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon}, t)\|^2$$
 - 6: **until** converged
-

Algorithm 2 Sampling

- 1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 - 4: $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}$
 - 5: **end for**
 - 6: **return** \mathbf{x}_0
-

Diffusion Model

Denoising Diffusion Probabilistic Models (DDPM)

<https://arxiv.org/abs/2006.11239>