EXERCISE · SLIDE 7 · QUESTION

Multimodal actions

Transfer: one route is more common

Find the squared-error minimizer u; compare it with the mode.

$$P(A = -2 \mid o) = \tfrac{3}{4}, \quad P(A = 2 \mid o) = \tfrac{1}{4}$$

- A action
- u fixed prediction to optimize

DL §3.8; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 7 · SOLUTION 1/1

Multimodal actions

Transfer: one route is more common

Solution 1 / 1

Choose the conditional mean

A is the random action; u is a deterministic prediction.

1. Squared error is minimized by $u^* = \mathbb{E}[A \mid o]$.
2. $u^* = \frac{3}{4}(-2) + \frac{1}{4}(2) = -\frac{3}{2} + \frac{1}{2} = -1$.
3. The mode is -2, while the minimum-risk mean is -1.

DL §3.8; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 7 · CHECK THE RESULT

Multimodal actions

Transfer: one route is more common

Find the squared-error minimizer u; compare it with the mode.

$$P(A = -2 \mid o) = \tfrac{3}{4}, \quad P(A = 2 \mid o) = \tfrac{1}{4}$$

A action

u fixed prediction to optimize

> **Check** $u^* = -1$; a mean need not be a mode.

DL §3.8; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 8 · QUESTION

Multimodal actions

A distribution and its mean answer different questions

Fixed observation; equally likely actions -2 and 2.
Compare: conditional mean vs a random draw.

Diffusion Policy §4.1; course exercise · Chi et al. (2023/2024)

Multimodal actions

A distribution and its mean answer different questions

Solution 1 / 1

Compare the operations

A is an action drawn from the two-point distribution.

1. $\mathbb{E}[A \mid o] = 0$: averaging returns one fixed value.
2. $A \sim p(\cdot \mid o)$: sampling returns -2 or 2, each with probability one half.
3. A correct distribution can still have an unusable mean.

Diffusion Policy §4.1; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 8 · CHECK THE RESULT

Multimodal actions

A distribution and its mean answer different questions

Fixed observation; equally likely actions -2 and 2.
Compare: conditional mean vs a random draw.

> **Check** Sampling preserves the two choices.

Diffusion Policy §4.1; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 12 · QUESTION

Forward diffusion

Check the conditional mean and covariance

Find the conditional mean and covariance of x_k.

$$x_k = \sqrt{1-\beta_k}\, x_{k-1} + \sqrt{\beta_k}\, z_k$$

x_{k-1} fixed when conditioned on

z_k independent N(0,I) vector

DDPM Eq.2; DL §3.8 · Ho, Jain & Abbeel (2020) · Goodfellow et al. (2016)

EXERCISE · SLIDE 12 · SOLUTION 1/1

Forward diffusion

Check the conditional mean and covariance

Solution 1 / 1

Conditioning fixes the input

z_k has mean zero and covariance I.

1. $\mathbb{E}[x_k \mid x_{k-1}] = \sqrt{1-\beta_k}\, x_{k-1}.$
2. $x_k - \mathbb{E}[x_k \mid x_{k-1}] = \sqrt{\beta_k} z_k.$
3. $\mathrm{Cov}(x_k \mid x_{k-1}) = \beta_k I.$

DDPM Eq.2; DL §3.8 · Ho, Jain & Abbeel (2020) · Goodfellow et al. (2016)

EXERCISE · SLIDE 12 · CHECK THE RESULT

Forward diffusion

Check the conditional mean and covariance

Find the conditional mean and covariance of x_k.

$$x_k = \sqrt{1-\beta_k}\, x_{k-1} + \sqrt{\beta_k}\, z_k$$

x_{k-1} fixed when conditioned on

z_k independent N(0,I) vector

> **Check** Mean $\sqrt{1-\beta_k}x_{k-1}$; covariance $\beta_k I$.

DDPM Eq.2; DL §3.8 · Ho, Jain & Abbeel (2020) · Goodfellow et al. (2016)

EXERCISE · SLIDE 13 · QUESTION

Forward diffusion

Transfer: variance is not the noise multiplier

Compute the realized x_k; show both coefficients.

$$\beta_k = 0.36, \quad x_{k-1} = 2, \quad z_k = -1$$

β_k variance

x_{k-1}, z_k given scalar values

DDPM Eq.2; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 13 · SOLUTION 1/1

Forward diffusion

Transfer: variance is not the noise multiplier

Solution 1 / 1

Compute coefficients before the state

The fresh noise z_k=-1 has been realized.

1. $\sqrt{1-\beta_k} = \sqrt{0.64} = 0.8$, and $\sqrt{\beta_k} = 0.6$.
2. $x_k = 0.8 \times 2 + 0.6 \times (-1) = 1.6 - 0.6 = 1$.
3. Using 0.36 as the noise multiplier would give the wrong variance.

DDPM Eq.2; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 13 · CHECK THE RESULT

Forward diffusion

Transfer: variance is not the noise multiplier

Compute the realized x_k; show both coefficients.

$$\beta_k = 0.36, \quad x_{k-1} = 2, \quad z_k = -1$$

β_k variance

x_{k-1}, z_k given scalar values

> **Check** $x_k = 0.8(2) + 0.6(-1) = 1.$

DDPM Eq.2; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 19 · QUESTION

Forward diffusion

Transfer: a two-component action

Compute x, then recover the sampled noise.

$$x_0 = \begin{bmatrix} 2 \\ -1 \end{bmatrix},\ \epsilon = \begin{bmatrix} -0.5 \\ 1 \end{bmatrix},\ \alpha = 0.8,\ \sigma = 0.6$$

$x_0, \epsilon \in \mathbb{R}^2$	given vectors
$x = \alpha x_0 + \sigma \epsilon$	noisy vector to find

DDPM Eq.4; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 19 · SOLUTION 1/1

Forward diffusion

Transfer: a two-component action

Solution 1 / 1

Apply the rule componentwise

x, x0 and epsilon are two-component column vectors.

---

1. $x_1 = 0.8(2) + 0.6(-0.5) = 1.3.$
2. $x_2 = 0.8(-1) + 0.6(1) = -0.2.$
3. $(x - 0.8x_0)/0.6 = (-0.5, 1)^T = \epsilon.$

DDPM Eq.4; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 19 · CHECK THE RESULT

Forward diffusion

Transfer: a two-component action

Compute x, then recover the sampled noise.

$$x_0 = \begin{bmatrix} 2 \\ -1 \end{bmatrix},\ \epsilon = \begin{bmatrix} -0.5 \\ 1 \end{bmatrix},\ \alpha = 0.8,\ \sigma = 0.6$$

$x_0, \epsilon \in \mathbb{R}^2$ given vectors

$x = \alpha x_0 + \sigma \epsilon$ noisy vector to find

Check $x = (1.3, -0.2)^T$.

DDPM Eq.4; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 20 · QUESTION

Forward diffusion

The marginal covariance also contains data spread

Find the mean and covariance of X_k.

$$X_k = \alpha_k X_0 + \sigma_k \epsilon, \quad X_0 \perp \epsilon$$

μ_0, Σ_0 clean-data mean and covariance

$\epsilon \sim \mathcal{N}(0, I)$ independent noise

DDPM Eq.4; DL §3.8 · Ho, Jain & Abbeel (2020) · Goodfellow et al. (2016)

EXERCISE · SLIDE 20 · SOLUTION 1/1

Forward diffusion

The marginal covariance also contains data spread

Solution 1 / 1

Average over examples and noise

X0 is now random; epsilon is independent and centered.

1. $\mathbb{E}[X_k] = \alpha_k \mathbb{E}[X_0] + \sigma_k \mathbb{E}[\epsilon] = \alpha_k \mu_0.$
2. $X_k - \mu_k = \alpha_k (X_0 - \mu_0) + \sigma_k \epsilon.$
3. $\mathrm{Cov}(X_k) = \alpha_k^2 \Sigma_0 + \sigma_k^2 I$; cross terms vanish by independence.

DDPM Eq.4; DL §3.8 · Ho, Jain & Abbeel (2020) · Goodfellow et al. (2016)

EXERCISE · SLIDE 20 · CHECK THE RESULT

Forward diffusion

The marginal covariance also contains data spread

Find the mean and covariance of X_k.

$$X_k = \alpha_k X_0 + \sigma_k \epsilon, \quad X_0 \perp \epsilon$$

μ_0, Σ_0 clean-data mean and covariance

$\epsilon \sim \mathcal{N}(0, I)$ independent noise

> **Check** $\mu_k = \alpha_k \mu_0$, $\Sigma_k = \alpha_k^2 \Sigma_0 + \sigma_k^2 I$.

DDPM Eq.4; DL §3.8 · Ho, Jain & Abbeel (2020) · Goodfellow et al. (2016)

EXERCISE · SLIDE 22 · QUESTION

Forward diffusion

Does this terminal schedule match the prior?

Can the terminal law equal N(0,1)? Check its moments.

$$\alpha_K = 0.8,\ \sigma_K = 0.6, \quad \mathbb{E}[X_0] = 5,\ \mathrm{Var}(X_0) = 1$$

X_0, X_K scalar clean and noisy data

$N(0,1)$ proposed generation prior

DDPM §4; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 22 · SOLUTION 1/1

Forward diffusion

Does this terminal schedule match the prior?

Solution 1 / 1

Use a necessary moment check

The added standard noise is independent of clean data.

---

1. $\mathbb{E}[X_K] = 0.8 \times 5 = 4.$
2. $\mathrm{Var}(X_K) = 0.8^2(1) + 0.6^2 = 1.$
3. A standard normal has mean zero. Matching its variance does not repair the mean mismatch.

DDPM §4; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 22 · CHECK THE RESULT

Forward diffusion

Does this terminal schedule match the prior?

Can the terminal law equal N(0,1)? Check its moments.

$$\alpha_K = 0.8,\ \sigma_K = 0.6, \quad \mathbb{E}[X_0] = 5,\ \mathrm{Var}(X_0) = 1$$

X_0, X_K scalar clean and noisy data

$N(0,1)$ proposed generation prior

> **Check** Variance 1 alone does not establish N(0,1).

DDPM §4; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 27 · QUESTION

The score

Compute an anisotropic score

Compute s(x). Is its direction parallel to −x?

$$\mu = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Sigma = \begin{bmatrix} 4 & 0 \\ 0 & 1 \end{bmatrix}, x = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

Σ covariance, not standard deviations

$s(x) = -\Sigma^{-1}(x - \mu)$ score to find

DL §3.9.3; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 27 · SOLUTION 1/1

The score

Compute an anisotropic score

Solution 1 / 1

Apply the precision matrix

The inverse covariance is the precision matrix.

1. $\Sigma^{-1} = \mathrm{diag}(1/4, 1)$.
2. $s(x) = -\mathrm{diag}(1/4, 1)(2, 1)^T = (-1/2, -1)^T$.
3. The narrower second direction receives stronger correction per unit displacement.

DL §3.9.3; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 27 · CHECK THE RESULT

The score

Compute an anisotropic score

Compute s(x). Is its direction parallel to −x?

$$\mu = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Sigma = \begin{bmatrix} 4 & 0 \\ 0 & 1 \end{bmatrix}, x = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

Σ covariance, not standard deviations

$s(x) = -\Sigma^{-1}(x - \mu)$ score to find

> **Check** $s(x) = (-0.5, -1)^T$.

DL §3.9.3; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 30 · QUESTION

The score

What does a stationary point tell us?

Find s(0). Is x=0 a mode? Explain.

$$p(x) = \tfrac{1}{2}\mathcal{N}(x; -2, 1) + \tfrac{1}{2}\mathcal{N}(x; 2, 1)$$

- x scalar state
- $s(x)$ gradient of log mixture density

DL §3.9.6; course exercise · Goodfellow et al. (2016)

The score

What does a stationary point tell us?

Solution 1 / 1

Opposing component directions cancel

At x=0, symmetry makes both responsibilities one half.

1. $s_1(0) = -(0+2) = -2$, $s_2(0) = -(0-2) = 2$.
2. $s(0) = \frac{1}{2}(-2) + \frac{1}{2}(2) = 0$.
3. Here $p''(0) = 3\mathcal{N}(0; 2, 1) > 0$: zero is a local density minimum.

DL §3.9.6; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 30 · CHECK THE RESULT

The score

What does a stationary point tell us?

Find s(0). Is x=0 a mode? Explain.

$$p(x) = \tfrac{1}{2}\mathcal{N}(x; -2, 1) + \tfrac{1}{2}\mathcal{N}(x; 2, 1)$$

x scalar state

$s(x)$ gradient of log mixture density

> **Check** At x=0, score 0 does not imply a mode.

DL §3.9.6; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 31 · QUESTION

The score

Transfer: doubling standard deviation

Compare standard deviations 1 and 2. Compute both scores.

$$\mu = 1, \quad x = 3, \quad s(x) = -(x - \mu)/v$$

- v variance
- $\sqrt{v}$ standard deviation

DL §3.9.3; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 31 · SOLUTION 1/1

The score

Transfer: doubling standard deviation

Solution 1 / 1

Convert spread into variance

The displacement x-mu equals 2 in both models.

1. First: $v = 1^2 = 1$, so $s(3) = -2$.
2. Second: $v = 2^2 = 4$, so $s(3) = -2/4 = -0.5$.
3. The wider log density changes less sharply at the same displacement.

DL §3.9.3; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 31 · CHECK THE RESULT

The score

Transfer: doubling standard deviation

Compare standard deviations 1 and 2. Compute both scores.

$$\mu = 1, \quad x = 3, \quad s(x) = -(x - \mu)/v$$

v variance

$\sqrt{v}$ standard deviation

> **Check** Doubling standard deviation divides the score by four.

DL §3.9.3; course exercise · Goodfellow et al. (2016)

EXERCISE · SLIDE 37 · QUESTION

Learning to denoise

Noise prediction gives a score estimate

Explain the minus sign and the factor 1/σ_k.

$$s_k(x) = -\frac{\mathbb{E}[\epsilon \mid x_k = x, k]}{\sigma_k} \approx -\frac{\epsilon_\theta(x, k)}{\sigma_k}$$

- s_k marginal noisy-data score
- $\sigma_k > 0$ noise standard deviation

DDPM §3.2; Vincent (2011) §4 · Ho, Jain & Abbeel (2020) · Vincent (2011)

Learning to denoise

Noise prediction gives a score estimate

Solution 1 / 2

Join regression to the score identity

At a fixed level, sigma_k is a positive constant.

1. $\nabla \log q_k = -\epsilon / \sigma_k$.
2. $s_k = \mathbb{E}[\nabla \log q_k \mid x_k] = -\mathbb{E}[\epsilon \mid x_k] / \sigma_k$.
3. The MSE optimum is $\epsilon_\theta^* = \mathbb{E}[\epsilon \mid x_k]$.

DDPM §3.2; Vincent (2011) §4 · Ho, Jain & Abbeel (2020) · Vincent (2011)

Learning to denoise

Noise prediction gives a score estimate

Solution 2 / 2

Noise regression is a form of denoising score matching

At fixed k, set s_theta=−epsilon_theta/sigma and c=∇log q_k=−epsilon/sigma.

1. $\mathbb{E}\|\epsilon_\theta - \epsilon\|^2 = \sigma_k^2 \mathbb{E}\|s_\theta - c\|^2$.
2. $\mathbb{E}\|s_\theta - c\|^2 = \mathbb{E}\|s_\theta - s_k\|^2 + C_k$.
3. Here C_k is the conditional target variance, independent of theta. Positive weights across k preserve the population optimum at each level.

DDPM §3.2; Vincent (2011) §4 · Ho, Jain & Abbeel (2020) · Vincent (2011)

EXERCISE · SLIDE 37 · CHECK THE RESULT

Learning to denoise

Noise prediction gives a score estimate

Explain the minus sign and the factor 1/σ_k.

$$s_k(x) = -\frac{\mathbb{E}[\epsilon \mid x_k = x, k]}{\sigma_k} \approx -\frac{\epsilon_\theta(x, k)}{\sigma_k}$$

s_k marginal noisy-data score

$\sigma_k > 0$ noise standard deviation

> **Check** $s_\theta = -\epsilon_\theta / \sigma_k$.

DDPM §3.2; Vincent (2011) §4 · Ho, Jain & Abbeel (2020) · Vincent (2011)

EXERCISE · SLIDE 39 · QUESTION

Learning to denoise

What can a zero noise prediction mean?

Find the noise for each clean hypothesis, then its mean.

$$X_0 \in \{-1, 1\}, \quad \alpha = \sigma = 1/\sqrt{2}, \quad X_k = 0$$

X_0 two equally likely clean values

ϵ noise conditioned on X_k=0

DDPM Eq.14; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 39 · SOLUTION 1/1

Learning to denoise

What can a zero noise prediction mean?

Solution 1 / 1

Two explanations of one input

The symmetric likelihoods give equal posterior weights.

1. If $X_0 = 1$, then $0 = (1 + \epsilon)/\sqrt{2}$, so $\epsilon = -1$.
2. If $X_0 = -1$, the required noise is $\epsilon = 1$.
3. $\mathbb{E}[\epsilon \mid X_k = 0] = (-1 + 1)/2 = 0$.

DDPM Eq.14; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 39 · CHECK THE RESULT

Learning to denoise

What can a zero noise prediction mean?

Find the noise for each clean hypothesis, then its mean.

$$X_0 \in \{-1, 1\}, \quad \alpha = \sigma = 1/\sqrt{2}, \quad X_k = 0$$

X_0 two equally likely clean values

ϵ noise conditioned on X_k=0

> **Check** Zero conditional mean can average nonzero noises.

DDPM Eq.14; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 42 · QUESTION

Learning to denoise

How does predictor error reach the clean estimate?

Find $\Delta\hat{x}_0$. Compare $\alpha=0.8$ and $\alpha=0.1$, with $\sigma^2=1-\alpha^2$.

$$\hat{x}_0 = \frac{x - \sigma\hat{\epsilon}}{\alpha}, \quad \hat{\epsilon} = \epsilon^* + \delta$$

- ϵ^* optimal mean-noise predictor
- δ prediction error; alpha > 0

DDPM Eq.11; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 42 · SOLUTION 1/1

Learning to denoise

How does predictor error reach the clean estimate?

Solution 1 / 1

Subtract the two estimates

Both estimates use the same x, alpha and sigma.

1. $\Delta\hat{x}_0 = [x - \sigma(\epsilon^* + \delta)]/\alpha - [x - \sigma\epsilon^*]/\alpha.$
2. $\Delta\hat{x}_0 = -(\sigma/\alpha)\delta.$
3. At $\alpha = 0.8$, the factor is 0.75; at $\alpha = 0.1$, it is $\sqrt{0.99}/0.1 \approx 9.95$.

DDPM Eq.11; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 42 · CHECK THE RESULT

Learning to denoise

How does predictor error reach the clean estimate?

Find $\Delta\hat{x}_0$. Compare $\alpha=0.8$ and $\alpha=0.1$, with $\sigma^2=1-\alpha^2$.

$$\hat{x}_0 = \frac{x - \sigma\hat{\epsilon}}{\alpha}, \quad \hat{\epsilon} = \epsilon^* + \delta$$

- ϵ^* optimal mean-noise predictor
- δ prediction error; alpha > 0

> **Check** $\Delta\hat{x}_0 = -(\sigma/\alpha)\delta$.

DDPM Eq.11; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 44 · QUESTION

Learning to denoise

Checkpoint: trace an input to a score

Data → known corruption → noise regression → score
Which quantity is sampled, observed, and learned?

DDPM Algorithms 1–2; Vincent (2011) §4 · Ho, Jain & Abbeel (2020) · Vincent (2011)

EXERCISE · SLIDE 44 · SOLUTION 1/1

Learning to denoise

Checkpoint: trace an input to a score

Solution 1 / 1

Connect the four objects

At one training level, alpha and sigma are fixed.

---

1. $x_0 \sim p_0$, $\epsilon \sim \mathcal{N}(0, I)$, then $x = \alpha x_0 + \sigma\epsilon$.
2. Input (x, k); target ϵ; population optimum $\mathbb{E}[\epsilon \mid x, k]$.
3. $s_k(x) = -\mathbb{E}[\epsilon \mid x, k]/\sigma$.

DDPM Algorithms 1–2; Vincent (2011) §4 · Ho, Jain & Abbeel (2020) · Vincent (2011)

EXERCISE · SLIDE 44 · CHECK THE RESULT

Learning to denoise

Checkpoint: trace an input to a score

Data → known corruption → noise regression → score
Which quantity is sampled, observed, and learned?

> **Check** Training learns a conditional mean noise field.

DDPM Algorithms 1–2; Vincent (2011) §4 · Ho, Jain & Abbeel (2020) · Vincent (2011)

EXERCISE · SLIDE 51 · QUESTION

Reverse sampling

A reverse draw uses a standard deviation

Compute one draw and the variance across repeated draws.

$$x_{k-1} = \mu_\theta + \sqrt{\tilde{\beta}_k}\, z, \quad z \sim \mathcal{N}(0, I)$$

$\mu_\theta = 1.025$ reverse mean from the previous example

$\tilde{\beta}_k = 0.24609375,\ z = -1$ derived posterior variance; fresh draw

DDPM Algorithm 2; course exercise · Ho, Jain & Abbeel (2020)

Reverse sampling

A reverse draw uses a standard deviation

Solution 1 / 1

Continue the same forward schedule

beta=.36, a=.64, sigma_k=.8; therefore bar-a_k=.36.

1. $\bar{a}_{k-1} = .36/.64 = .5625.$
2. $\tilde{\beta}_k = .36(1 - .5625)/.64 = .24609375.$
3. $\sqrt{\tilde{\beta}_k} \approx .4960784;\ x_{k-1} = 1.025 - .4960784 = .5289216.$
4. Across fresh draws, $\mathrm{Var}(x_{k-1} \mid x_k) = .24609375$ in this Gaussian model.

DDPM Algorithm 2; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 51 · CHECK THE RESULT

Reverse sampling

A reverse draw uses a standard deviation

Compute one draw and the variance across repeated draws.

$$x_{k-1} = \mu_\theta + \sqrt{\tilde{\beta}_k}\, z, \qquad z \sim \mathcal{N}(0, I)$$

$\mu_\theta = 1.025$ reverse mean from the previous example

$\tilde{\beta}_k = 0.24609375,\ z = -1$ derived posterior variance; fresh draw

Check $x_{k-1} = 1.025 - \sqrt{0.24609375} \approx 0.5289.$

DDPM Algorithm 2; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 53 · QUESTION

Reverse sampling

Does a mean predictor force identical samples?

Same observation and trained weights.
Different initial Gaussian draws.
Must the final samples be identical?

DDPM Algorithm 2; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 53 · SOLUTION 1/1

Reverse sampling

Does a mean predictor force identical samples?

Solution 1 / 1

Follow the whole random process

The network is deterministic at a fixed input and level.

1. X_K is random, so $\mu_\theta(X_K, K)$ can vary.
2. For stochastic reverse steps, fresh z_k adds variation.
3. A conditional mean at each state does not equal the unconditional mean of every final sample.

DDPM Algorithm 2; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 53 · CHECK THE RESULT

Reverse sampling

Does a mean predictor force identical samples?

Same observation and trained weights.
Different initial Gaussian draws.
Must the final samples be identical?

> **Check** A conditional mean field is not one fixed output.

DDPM Algorithm 2; course exercise · Ho, Jain & Abbeel (2020)

EXERCISE · SLIDE 57 · QUESTION

Reverse sampling

Transfer: check deterministic sampling at the endpoint

Compute the deterministic clean-endpoint output. Is the whole generated distribution deterministic?

$$x_k = 0.9, \quad \alpha_k = 0.8, \quad \sigma_k = 0.6, \quad \hat{\epsilon} = 0.5$$

$\alpha_0 = 1, \sigma_0 = 0$ clean endpoint

x_0^{out} deterministic output to find

DDIM Eq.12; course exercise · Song, Meng & Ermon (2021)

EXERCISE · SLIDE 57 · SOLUTION 1/1

Reverse sampling

Transfer: check deterministic sampling at the endpoint

Solution 1 / 1

Set the destination to the clean endpoint

At level zero, signal coefficient is one and noise coefficient is zero.

---

1. $\hat{x}_0 = (0.9 - 0.3)/0.8 = 0.75.$
2. $x_0^{\text{out}} = 1\hat{x}_0 + 0\hat{\epsilon} = 0.75.$
3. The initial random state can still differ across generation runs; no fresh endpoint noise is added.

DDIM Eq.12; course exercise · Song, Meng & Ermon (2021)

EXERCISE · SLIDE 57 · CHECK THE RESULT

Reverse sampling

Transfer: check deterministic sampling at the endpoint

Compute the deterministic clean-endpoint output. Is the whole generated distribution deterministic?

$$x_k = 0.9, \quad \alpha_k = 0.8, \quad \sigma_k = 0.6, \quad \hat{\epsilon} = 0.5$$

$\alpha_0 = 1, \sigma_0 = 0$ clean endpoint

x_0^{out} deterministic output to find

> **Check** $x_0^{\text{out}} = (0.9 - 0.6(0.5))/0.8 = 0.75.$

DDIM Eq.12; course exercise · Song, Meng & Ermon (2021)

EXERCISE · SLIDE 60 · QUESTION

Reverse sampling

A step budget trades computation for approximation

Toy measured cost: 8 ms per network evaluation.
Compare 20 evaluations with 5.
Which claims require an experiment?

Diffusion Policy §3.4; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 60 · SOLUTION 1/1

Reverse sampling

A step budget trades computation for approximation

Solution 1 / 1

Separate timing arithmetic from quality evidence

The hypothetical cost is 8 ms per sequential network evaluation.

1. 20 evaluations: 20 × 8 = 160 ms. Five evaluations: 5 × 8 = 40 ms.
2. Fixed overhead and hardware details were omitted, so these are model-only estimates.
3. Equal task success does not follow from the timing calculation.

Diffusion Policy §3.4; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 60 · CHECK THE RESULT

Reverse sampling

A step budget trades computation for approximation

Toy measured cost: 8 ms per network evaluation.
Compare 20 evaluations with 5.
Which claims require an experiment?

> **Check** Model time: 160 ms vs 40 ms; quality must be measured.

Diffusion Policy §3.4; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 65 · QUESTION

Conditioning

Transfer: calculate and interpret guidance

Compute guided noise and score for each w.

$$\hat{\epsilon}_u = 0.2, \quad \hat{\epsilon}_c = -0.1, \quad \sigma_k = 0.5$$

$w = 0, 1, 2$ scales to compare

$s_g = -\hat{\epsilon}_g / \sigma_k$ guided score

Classifier-Free Guidance §3.2; course exercise · Ho & Salimans (2022)

EXERCISE · SLIDE 65 · SOLUTION 1/1

Conditioning

Transfer: calculate and interpret guidance

Solution 1 / 1

Compute the difference and apply the scale

Conditional minus unconditional noise equals -0.3.

1. $w = 0$: $\hat{\epsilon}_g = 0.2$, so $s_g = -0.4$.
2. $w = 1$: $\hat{\epsilon}_g = -0.1$, so $s_g = 0.2$.
3. $w = 2$: $\hat{\epsilon}_g = -0.4$, so $s_g = 0.8$; this extrapolates.

Classifier-Free Guidance §3.2; course exercise · Ho & Salimans (2022)

EXERCISE · SLIDE 65 · CHECK THE RESULT

Conditioning

Transfer: calculate and interpret guidance

Compute guided noise and score for each w.

$$\hat{\epsilon}_u = 0.2, \quad \hat{\epsilon}_c = -0.1, \quad \sigma_k = 0.5$$

$w = 0, 1, 2$ scales to compare

$s_g = -\hat{\epsilon}_g / \sigma_k$ guided score

> **Check** w=1 gives the ordinary conditional prediction.

Classifier-Free Guidance §3.2; course exercise · Ho & Salimans (2022)

EXERCISE · SLIDE 67 · QUESTION

Diffusion Policy

Three horizons serve different purposes

Count observations, predicted scalars, and executed steps.

$$H_o = 2, \quad H_p = 4, \quad H_a = 2, \quad d_a = 2$$

- H_o observations supplied
- H_p actions predicted
- H_a actions executed before replanning

Diffusion Policy §2.3; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 67 · SOLUTION 1/1

Diffusion Policy

Three horizons serve different purposes

Solution 1 / 1

Read each horizon in its own clock

Each of four predicted actions has two components.

1. Observation input: two recent observations.
2. Prediction: $4 \times 2 = 8$ scalar components.
3. Execution: apply the first two actions, then obtain new observations and replan.

Diffusion Policy §2.3; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 67 · CHECK THE RESULT

Diffusion Policy

Three horizons serve different purposes

Count observations, predicted scalars, and executed steps.

$$H_o = 2, \quad H_p = 4, \quad H_a = 2, \quad d_a = 2$$

- H_o observations supplied
- H_p actions predicted
- H_a actions executed before replanning

> **Check** 2 observations; 8 predicted scalars; execute 2 steps.

Diffusion Policy §2.3; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 69 · QUESTION

Diffusion Policy

A receding horizon closes the feedback loop

Observe → sample chunk → execute prefix → observe again
What happens to the unused suffix?

Diffusion Policy §2.3 · Chi et al. (2023/2024)

EXERCISE · SLIDE 69 · SOLUTION 1/1

Diffusion Policy

A receding horizon closes the feedback loop

Solution 1 / 1

Close the feedback loop

The learned parameters remain fixed during deployment.

1. Collect the current observation context, then denoise an initialized action chunk.
2. Execute its first H_a actions using the robot controller.
3. Refresh the context and generate a new chunk; the unexecuted suffix is not a commitment.

Diffusion Policy §2.3 · Chi et al. (2023/2024)

EXERCISE · SLIDE 69 · CHECK THE RESULT

Diffusion Policy

A receding horizon closes the feedback loop

Observe → sample chunk → execute prefix → observe again
What happens to the unused suffix?

> **Check** Replan using fresh observations; do not blindly replay the old suffix.

Diffusion Policy §2.3 · Chi et al. (2023/2024)

EXERCISE · SLIDE 73 · QUESTION

Diffusion Policy

Transfer: change execution without changing the sampler

Targets: 0.6 m, then 0.3 m. Compare H_a=1 and H_a=2. An obstacle moves after the first action: which policy can refresh its second target?

Diffusion Policy §2.3; course exercise · Chi et al. (2023/2024)

Diffusion Policy

Transfer: change execution without changing the sampler

Solution 1 / 1

Separate generation from execution

The clean predicted chunk and sampler are unchanged.

1. $H_a = 1$: execute 0.6 m target, observe, then regenerate the next target.
2. $H_a = 2$: execute both planned targets before the policy replans.
3. The first policy can incorporate the new observation earlier; neither horizon alone guarantees safety.

Diffusion Policy §2.3; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 73 · CHECK THE RESULT

Diffusion Policy

Transfer: change execution without changing the sampler

Targets: 0.6 m, then 0.3 m. Compare H_a=1 and H_a=2. An obstacle moves after the first action: which policy can refresh its second target?

> **Check** Shorter execution horizon refreshes policy feedback sooner.

Diffusion Policy §2.3; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 75 · QUESTION

Diffusion Policy

A low denoising loss is not a deployment certificate

Training metric: noise-prediction error.
Deployment question: does the closed-loop task succeed?
Give two missing checks.

Diffusion Policy §§5,9; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 75 · SOLUTION 1/1

Diffusion Policy

A low denoising loss is not a deployment certificate

Solution 1 / 1

Trace the gap to the physical task

Noise loss is a surrogate measured on a specified data distribution.

1. New observations can lie outside that distribution; evaluate held-out environments.
2. Wrong units, action alignment, controller tracking or latency can break an otherwise correct denoiser.
3. Measure closed-loop success and uncertainty; sampling plausible actions alone does not enforce all constraints.

Diffusion Policy §§5,9; course exercise · Chi et al. (2023/2024)

EXERCISE · SLIDE 75 · CHECK THE RESULT

Diffusion Policy

A low denoising loss is not a deployment certificate

Training metric: noise-prediction error.
Deployment question: does the closed-loop task succeed?
Give two missing checks.

> **Check** Validate the entire observation-to-action loop.

Diffusion Policy §§5,9; course exercise · Chi et al. (2023/2024)