EXERCISE · SLIDE 4 · QUESTION

Flow Matching and VLA

Which clock is running?

k: physical robot step; $\tau \in [0, 1]$: generator time. One action chunk is refined while the camera observation stays fixed.

A. Every flow step moves the robot.

B. Flow steps refine a candidate chunk.

C. Flow time is the camera frame index.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples

Flow Matching and VLA

Which clock is running?

Solution 1 / 1

Separate computation from execution

k: robot time; τ: generation time; A: candidate action chunk.

1. The generator changes the candidate A while holding the observation at robot step k fixed.
2. The robot executes selected entries only after generation.
3. B is correct. A confuses numerical refinement with execution; C identifies two different clocks.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples

EXERCISE · SLIDE 4 · CHECK THE RESULT

Flow Matching and VLA

Which clock is running?

k: physical robot step; $\tau \in [0,1]$: generator time. One action chunk is refined while the camera observation stays fixed.

A. Every flow step moves the robot.

B. Flow steps refine a candidate chunk. ✓

C. Flow time is the camera frame index.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples

EXERCISE · SLIDE 9 · QUESTION

Probability Flows

Transfer: a different velocity

$X_0 \sim \mathcal{N}(1, 4)$ and $dx_\tau/d\tau = -3$. Find the law at $\tau = 1$.

A. Gaussian mean -2, variance 4

B. Gaussian mean -2, variance 16

C. Gaussian mean 4, variance 4

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 9 · SOLUTION 1/2

Probability Flows

Transfer: a different velocity

Solution 1 / 2

Integrate before transforming the law

X_0: initial random value; X_1: final value.

---

1. $X_1 = X_0 + \int_0^1 (-3)\, d\tau = X_0 - 3.$
2. $\mathbb{E}[X_1] = 1 - 3 = -2.$
3. $\mathrm{Var}(X_1) = \mathrm{Var}(X_0) = 4.$

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 9 · SOLUTION 2/2

Probability Flows

Transfer: a different velocity

Solution 2 / 2

Choose A

$\mathcal{N}(\mu, \sigma^2)$ uses variance as its second argument.

1. A follows because a constant shift leaves centered differences unchanged.
2. B incorrectly scales the spread. C reverses the velocity sign.

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 9 · CHECK THE RESULT

Probability Flows

Transfer: a different velocity

$X_0 \sim \mathcal{N}(1, 4)$ and $dx_\tau/d\tau = -3$. Find the law at $\tau = 1$.

- **A** Gaussian mean -2, variance 4 ✓
- **B** Gaussian mean -2, variance 16
- **C** Gaussian mean 4, variance 4

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 14 · QUESTION

Probability Flows

Does constant velocity freeze density?

$p_\tau(x) = p_0(x - 2\tau)$, $v = 2$. At a fixed x, does density generally stay constant?

A. Yes, because velocity is constant.

B. No, because the density profile moves.

C. No, because total probability grows.

Lipman et al. (2023), Appendix B · original course examples

EXERCISE · SLIDE 14 · SOLUTION 1/2

Probability Flows

Does constant velocity freeze density?

Solution 1 / 2

Differentiate the translated profile

p_0 is differentiable; prime means derivative of its scalar argument.

1. $\partial_\tau p_0(x - 2\tau) = -2p_0'(x - 2\tau).$
2. $-\partial_x(2p_\tau) = -2p_0'(x - 2\tau).$
3. These sides match, so translation satisfies continuity.

Lipman et al. (2023), Appendix B · original course examples

EXERCISE · SLIDE 14 · SOLUTION 2/2

Probability Flows

Does constant velocity freeze density?

Solution 2 / 2

Choose B

x stays fixed while the profile passes it.

1. B: the density at x can change even with constant velocity.
2. A confuses constant velocity with constant density. C violates conserved total mass.

Lipman et al. (2023), Appendix B · original course examples

EXERCISE · SLIDE 14 · CHECK THE RESULT

Probability Flows

Does constant velocity freeze density?

$p_\tau(x) = p_0(x - 2\tau)$, $v = 2$. At a fixed x, does density generally stay constant?

A. Yes, because velocity is constant.

B. No, because the density profile moves. ✓

C. No, because total probability grows.

Lipman et al. (2023), Appendix B · original course examples

EXERCISE · SLIDE 16 · QUESTION

Probability Flows

Transfer: contracting toward a target mean

$X_\tau = 3\tau + (1 - \tau/2)X_0$, $X_0 \sim \mathcal{N}(0, 4)$. Find mean and variance at $\tau = 1$.

A Mean 3, variance 1

B Mean 3, variance 2

C Mean 0, variance 1

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 16 · SOLUTION 1/2

Probability Flows

Transfer: contracting toward a target mean

Solution 1 / 2

Apply affine moment rules

$X_1 = 3 + X_0/2$; initial variance is 4.

---

1. $\mathbb{E}[X_1] = 3 + (1/2)0 = 3.$
2. $\mathrm{Var}(X_1) = (1/2)^2(4) = 1.$
3. The endpoint scale 1/2 stays positive, so this affine map remains invertible.

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 16 · SOLUTION 2/2

Probability Flows

Transfer: contracting toward a target mean

Solution 2 / 2

Choose A

$\mathcal{N}(3, 1)$ is the final Gaussian.

---

1. A squares the scale when transforming variance.
2. B multiplies variance by the scale only. C ignores translation.

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 16 · CHECK THE RESULT

Probability Flows

Transfer: contracting toward a target mean

$X_\tau = 3\tau + (1 - \tau/2)X_0$, $X_0 \sim \mathcal{N}(0, 4)$. Find mean and variance at $\tau = 1$.

- **A** Mean 3, variance 1 ✓
- **B** Mean 3, variance 2
- **C** Mean 0, variance 1

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 21 · QUESTION

Conditional Paths

Transfer: another pair

$x_0 = (2, -1)^T$, $x_1 = (-2, 3)^T$, $\tau = 1/4$. Choose (x_τ, u).

A ((1, 0), (−4, 4))

B ((−1, 2), (−4, 4))

C ((1, 0), (−1, 1))

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 21 · SOLUTION 1/2

Conditional Paths

Transfer: another pair

Solution 1 / 2

Compute the rate, then the state

$u = x_1 - x_0$; $x_\tau = x_0 + \tau u$.

1. $u = (-2 - 2, 3 - (-1))^T = (-4, 4)^T$.
2. $x_{1/4} = (2, -1)^T + (1/4)(-4, 4)^T = (1, 0)^T$.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 21 · SOLUTION 2/2

Conditional Paths

Transfer: another pair

Solution 2 / 2

Choose A

$\tau u = (-1, 1)^T$ is an increment, not the rate.

---

1. A gives both quantities correctly.
2. B evaluates at τ=3/4. C replaces the rate by a quarter-step increment.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 21 · CHECK THE RESULT

Conditional Paths

Transfer: another pair

$x_0 = (2, -1)^T$, $x_1 = (-2, 3)^T$, $\tau = 1/4$. Choose (x_τ, u).

- **A** $((1, 0), (-4, 4))$ ✓
- B $((-1, 2), (-4, 4))$
- C $((1, 0), (-1, 1))$

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 23 · QUESTION

Conditional Paths

Is the noise coefficient a variance?

$X_0 \sim \mathcal{N}(0, 1)$, fixed $y = 4$, $X_\tau = (1 - \tau)X_0 + \tau y$. Use $\tau = 1/2$.

A Mean 2, variance 1/2

B Mean 2, variance 1/4

C Mean 4, variance 1/4

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 23 · SOLUTION 1/2

Conditional Paths

Is the noise coefficient a variance?

Solution 1 / 2

Transform the standard Gaussian

$X_{1/2} = X_0/2 + 2.$

---

1. $\mathbb{E}[X_{1/2}] = 2.$
2. $\mathrm{Var}(X_{1/2}) = (1/2)^2\,\mathrm{Var}(X_0) = 1/4.$
3. The standard deviation is 1/2.

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 23 · SOLUTION 2/2

Conditional Paths

Is the noise coefficient a variance?

Solution 2 / 2

Choose B

$y = 4$ is fixed; it supplies no conditional variance.

---

1. B keeps the standard deviation and variance distinct.
2. A uses the noise coefficient as variance. C uses the final rather than intermediate mean.

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 23 · CHECK THE RESULT

Conditional Paths

Is the noise coefficient a variance?

$X_0 \sim \mathcal{N}(0,1)$, fixed $y = 4$, $X_\tau = (1-\tau)X_0 + \tau y$. Use $\tau = 1/2$.

A Mean 2, variance 1/2

B Mean 2, variance 1/4 ✓

C Mean 4, variance 1/4

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 24 · QUESTION

Conditional Paths

Can we use the endpoint during sampling?

Training stores (x_0, x_1, τ) and constructs x_τ. A new generation call has only noise and an observation.

A. Give the network the unknown x_1.

B. Use $x_1 - x_0$ directly at every step.

C. Predict velocity from x_τ, τ and observation.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 24 · SOLUTION 1/1

Conditional Paths

Can we use the endpoint during sampling?

Solution 1 / 1

Separate supervision from inputs

x_1: training target; v_θ: learned velocity.

1. During training x_1 constructs the target $u = x_1 - x_0$.
2. At generation time x_1 is unknown, so u cannot be supplied directly.
3. C is executable. A and B assume the answer is already known.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 24 · CHECK THE RESULT

Conditional Paths

Can we use the endpoint during sampling?

Training stores (x_0, x_1, τ) and constructs x_τ. A new generation call has only noise and an observation.

A. Give the network the unknown x_1.

B. Use $x_1 - x_0$ directly at every step.

C. Predict velocity from x_τ, τ and observation. ✓

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 27 · QUESTION

Conditional Paths

Which average does the field need?

At a fixed (x, τ), velocity targets are $+2$ with probability $3/4$ and -2 with probability $1/4$.

A. The unconditional endpoint average.

B. The target mean given this x and τ.

C. Always the most frequent target exactly.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 27 · SOLUTION 1/1

Conditional Paths

Which average does the field need?

Solution 1 / 1

Use the input-conditioned target law

$U \mid x, \tau$ is the stated two-valued random target.

---

1. $\mathbb{E}[U \mid x, \tau] = (3/4)2 + (1/4)(-2) = 1.$
2. B selects the relevant conditional distribution.
3. A discards the input condition. C gives +2, which need not minimize squared error.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 27 · CHECK THE RESULT

Conditional Paths

Which average does the field need?

At a fixed (x, τ), velocity targets are $+2$ with probability $3/4$ and -2 with probability $1/4$.

A. The unconditional endpoint average.

B. The target mean given this x and τ. ✓

C. Always the most frequent target exactly.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 32 · QUESTION

Learning the Velocity

Transfer: changed conditional weights

$P(U=-1 \mid x,\tau)=1/4$, $P(U=3 \mid x,\tau)=3/4$. Minimize conditional squared error.

A Predict 1

B Predict 2

C Predict 3

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 32 · SOLUTION 1/2

Learning the Velocity

Transfer: changed conditional weights

Solution 1 / 2

Compute the conditional mean

U: scalar rate target; a: prediction.

1. $a^* = (1/4)(-1) + (3/4)3 = 2.$
2. $R(a) = (1/4)(a+1)^2 + (3/4)(a-3)^2.$
3. $R(a) = a^2 - 4a + 7 = (a-2)^2 + 3.$

Tong et al. (2024), §§2.2–3.2 · original course examples

Learning the Velocity

Transfer: changed conditional weights

Solution 2 / 2

Choose B

$R(2) = 3$, while $R(1) = R(3) = 4$.

1. B minimizes the risk.
2. A uses equal rather than stated weights. C selects the more frequent label rather than the mean.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 32 · CHECK THE RESULT

Learning the Velocity

Transfer: changed conditional weights

$P(U=-1 \mid x,\tau)=1/4$, $P(U=3 \mid x,\tau)=3/4$. Minimize conditional squared error.

A Predict 1

B Predict 2 ✓

C Predict 3

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 35 · QUESTION

Learning the Velocity

Shapes and loss reduction

$B = 2$ samples, $d = 3$ coordinates. $V, U \in \mathbb{R}^{2\times 3}$. Use $L = \frac{1}{Bd}\sum_{i,j}(V_{ij} - U_{ij})^2$.

A. Gradient is 2(V−U)

B. Gradient is (V−U)/3

C. Gradient has shape 3×2

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 35 · SOLUTION 1/1

Learning the Velocity

Shapes and loss reduction

Solution 1 / 1

Differentiate the stated mean

V: predictions; U: fixed labels; $Bd = 6$.

1. $\partial L / \partial V_{ij} = 2(V_{ij} - U_{ij})/6$.
2. $\nabla_V L = (V - U)/3 \in \mathbb{R}^{2\times 3}$.
3. B is correct. A omits the reduction; C transposes the gradient shape.

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 35 · CHECK THE RESULT

Learning the Velocity

Shapes and loss reduction

$B = 2$ samples, $d = 3$ coordinates. $V, U \in \mathbb{R}^{2\times 3}$. Use $L = \frac{1}{Bd}\sum_{i,j}(V_{ij} - U_{ij})^2$.

A. Gradient is 2(V−U)

B. Gradient is (V−U)/3 ✓

C. Gradient has shape 3×2

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 37 · QUESTION

Learning the Velocity

Transfer: a negative input

$v_w(x) = wx$, one pair $(x, u) = (-2, 3)$, $L = \frac{1}{2}(wx - u)^2$, $w = 1$, $\eta = .1$.

A New w is 0

B New w is 2

C New w is 1.5

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 37 · SOLUTION 1/2

Learning the Velocity

Transfer: a negative input

Solution 1 / 2

Differentiate before updating

$e = wx - u$: residual; w: weight; x, u: fixed input and target.

1. $e = 1(-2) - 3 = -5.$
2. $dL/dw = ex = (-5)(-2) = 10.$
3. $w_{new} = 1 - .1(10) = 0.$

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 37 · SOLUTION 2/2

Learning the Velocity

Transfer: a negative input

Solution 2 / 2

Choose A

Old loss $25/2$; new loss $9/2$.

1. A follows gradient descent and decreases this example loss.
2. B uses the wrong update sign. C omits the input factor in the chain rule.

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 37 · CHECK THE RESULT

Learning the Velocity

Transfer: a negative input

$v_w(x) = wx$, one pair $(x, u) = (-2, 3)$, $L = \frac{1}{2}(wx - u)^2$, $w = 1$, $\eta = .1$.

- **A** New w is 0 ✓
- **B** New w is 2
- **C** New w is 1.5

Goodfellow et al. (2016), §§3.8–3.9.3 · original course examples

EXERCISE · SLIDE 39 · QUESTION

Sampling the Learned Flow

What remains fixed while sampling?

$dx_\tau/d\tau = v_\theta(x_\tau, \tau, c)$; draw x_0 once. The condition c is one observation.

A. Optimize θ at every flow step.

B. Hold θ and c fixed; update x and τ.

C. Replace x with fresh noise every step.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 39 · SOLUTION 1/1

Sampling the Learned Flow

What remains fixed while sampling?

Solution 1 / 1

Read the differential equation

θ: learned parameters; c: condition; x_τ: generated state.

1. The ODE evolves x with generation time τ under one fixed field.
2. B follows the stated ODE. A adds test-time training; C defines a different stochastic process.
3. The initial draw supplies randomness even with deterministic integration.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 39 · CHECK THE RESULT

Sampling the Learned Flow

What remains fixed while sampling?

$dx_\tau/d\tau = v_\theta(x_\tau, \tau, c)$; draw x_0 once. The condition c is one observation.

A. Optimize θ at every flow step.

B. Hold θ and c fixed; update x and τ. ✓

C. Replace x with fresh noise every step.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 42 · QUESTION

Sampling the Learned Flow

Transfer: a time-dependent rate

$v(x, \tau) = 2\tau$, $x_0 = 1$, $h = 1/2$. Use two forward Euler steps from $\tau = 0$ to 1.

- (A) 1.5
- (B) 2
- (C) 2.5

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 42 · SOLUTION 1/2

Sampling the Learned Flow

Transfer: a time-dependent rate

Solution 1 / 2

Evaluate at the beginning of each step

x_j: Euler state; $\tau_j = j/2$.

1. $v(x_0, 0) = 0$, so $x_1 = 1 + (1/2)0 = 1$.
2. $v(x_1, 1/2) = 1$, so $x_2 = 1 + (1/2)1 = 1.5$.

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 42 · SOLUTION 2/2

Sampling the Learned Flow

Transfer: a time-dependent rate

Solution 2 / 2

Choose A and compare the exact integral

$x(1) = 1 + \int_0^1 2\tau \, d\tau = 2.$

1. A is the forward Euler approximation.
2. B is exact here, not two-step forward Euler. C uses right-end rates.

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 42 · CHECK THE RESULT

Sampling the Learned Flow

Transfer: a time-dependent rate

$v(x, \tau) = 2\tau$, $x_0 = 1$, $h = 1/2$. Use two forward Euler steps from $\tau = 0$ to 1.

- **A** 1.5 ✓
- **B** 2
- **C** 2.5

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 45 · QUESTION

Sampling the Learned Flow

Transfer: a fixed evaluation budget

Budget: 8 calls to v_θ. Euler uses 1 call/step; explicit midpoint uses 2. Both integrate $[0,1]$.

A. Euler h=1/8; midpoint h=1/4

B. Both h=1/8

C. Midpoint always gives the smaller error

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 45 · SOLUTION 1/1

Sampling the Learned Flow

Transfer: a fixed evaluation budget

Solution 1 / 1

Count steps from calls per step

N: accepted steps; $h = 1/N$.

1. Euler: $N = 8/1 = 8$, so $h = 1/8$.
2. Midpoint: $N = 8/2 = 4$, so $h = 1/4$.
3. A uses the full budget. B needs 16 midpoint calls; C needs assumptions and an error comparison.

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 45 · CHECK THE RESULT

Sampling the Learned Flow

Transfer: a fixed evaluation budget

Budget: 8 calls to v_θ. Euler uses 1 call/step; explicit midpoint uses 2. Both integrate $[0,1]$.

- **A** Euler h=1/8; midpoint h=1/4 ✓
- **B** Both h=1/8
- **C** Midpoint always gives the smaller error

Lipman et al. (2023), §§2–4 · original course examples

EXERCISE · SLIDE 48 · QUESTION

Sampling the Learned Flow

Reversing the generation clock

$s = 1 - \tau$; $y_s = x_{1-s}$; $dx_\tau/d\tau = v(x_\tau, \tau)$. Find dy_s/ds.

A $v(y_s, 1-s)$

B $-v(y_s, 1-s)$

C $-v(y_s, s)$ always

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 48 · SOLUTION 1/1

Sampling the Learned Flow

Reversing the generation clock

Solution 1 / 1

Apply the chain rule

s: reversed clock; $y_s = x_{1-s}$.

1. $dy_s/ds = (dx_\tau/d\tau)(d\tau/ds)$.
2. $d\tau/ds = -1$, hence $dy_s/ds = -v(y_s, 1-s)$.
3. B changes both the sign and time argument. A misses the sign; C generally uses the wrong time.

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 48 · CHECK THE RESULT

Sampling the Learned Flow

Reversing the generation clock

$s = 1 - \tau$; $y_s = x_{1-s}$; $dx_\tau/d\tau = v(x_\tau, \tau)$. Find dy_s/ds.

A $v(y_s, 1-s)$

B $-v(y_s, 1-s)$ ✓

C $-v(y_s, s)$ always

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 52 · QUESTION

Pairing and Optimal Transport

Transfer: compute the assignment cost

Sources $0, 2$; targets $1, 4$; each has mass $1/2$. Compare aligned $(0 \to 1, 2 \to 4)$ and crossed pairs.

A. Aligned cost 2.5; crossed 8.5

B. Aligned cost 5; crossed 17

C. Both cost 2.5

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 52 · SOLUTION 1/2

Pairing and Optimal Transport

Transfer: compute the assignment cost

Solution 1 / 2

Square each displacement, then average

C is expected squared cost, not a sum.

---

1. $C_{aligned} = [(1-0)^2 + (4-2)^2]/2 = (1+4)/2 = 2.5.$
2. $C_{crossed} = [(4-0)^2 + (1-2)^2]/2 = (16+1)/2 = 8.5.$

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 52 · SOLUTION 2/2

Pairing and Optimal Transport

Transfer: compute the assignment cost

Solution 2 / 2

Choose A

Both assignments preserve one-half mass at each endpoint.

---

1. A includes the probability weights.
2. B reports sums, not expected costs. C assumes fixed marginals determine the joint pairing.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 52 · CHECK THE RESULT

Pairing and Optimal Transport

Transfer: compute the assignment cost

Sources $0, 2$; targets $1, 4$; each has mass $1/2$. Compare aligned $(0 \to 1, 2 \to 4)$ and crossed pairs.

- **A** Aligned cost 2.5; crossed 8.5 ✓
- **B** Aligned cost 5; crossed 17
- **C** Both cost 2.5

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 53 · QUESTION

Pairing and Optimal Transport

Are learned trajectories necessarily straight?

Training paths are straight. The learned field is $v^*(x, \tau) = \mathbb{E}[U \mid x, \tau]$.

A. Yes, every generated path is one training segment.

B. No; the conditional average can vary along a path.

C. Every learned path must cross another path.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 53 · SOLUTION 1/1

Pairing and Optimal Transport

Are learned trajectories necessarily straight?

Solution 1 / 1

Track the field along the generated state

U: paired velocity label; v^*: conditional mean field.

1. As x and τ change, the compatible hidden endpoint pairs can change.
2. Their conditional average can therefore change direction along a generated trajectory.
3. B is correct. A confuses paired paths with the learned flow; C contradicts uniqueness for a regular ODE.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 53 · CHECK THE RESULT

Pairing and Optimal Transport

Are learned trajectories necessarily straight?

Training paths are straight. The learned field is $v^*(x, \tau) = \mathbb{E}[U \mid x, \tau]$.

A. Yes, every generated path is one training segment.

B. No; the conditional average can vary along a path. ✓

C. Every learned path must cross another path.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 54 · QUESTION

Pairing and Optimal Transport

What does a conditional path establish?

For each data point y, Gaussian noise contracts toward y along a simple conditional path. The data point may be sampled independently.

A. This alone proves globally optimal endpoint pairing.

B. A simple conditional path and global OT are different claims.

C. Every coupling has equal expected displacement cost.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 54 · SOLUTION 1/1

Pairing and Optimal Transport

What does a conditional path establish?

Solution 1 / 1

Identify what was optimized

y: fixed endpoint in a conditional path; π: joint endpoint pairing.

1. A simple path conditional on y does not choose the best joint pairing π.
2. Global quadratic OT additionally minimizes expected cost over couplings with fixed marginals.
3. B is correct. A skips that optimization; C fails in the 1-versus-9 cost example.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 54 · CHECK THE RESULT

Pairing and Optimal Transport

What does a conditional path establish?

For each data point y, Gaussian noise contracts toward y along a simple conditional path. The data point may be sampled independently.

A. This alone proves globally optimal endpoint pairing.

B. A simple conditional path and global OT are different claims. ✓

C. Every coupling has equal expected displacement cost.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 55 · QUESTION

Pairing and Optimal Transport

Re-pair examples with different conditions

Condition c says “move left”. Its data actions are negative. A batch also contains “move right” examples.

A. Re-pair all targets freely and retain old c.

B. Keep targets compatible with each condition.

C. Remove c from training and still claim conditioning.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 55 · SOLUTION 1/1

Pairing and Optimal Transport

Re-pair examples with different conditions

Solution 1 / 1

Preserve the conditional target law

$p_{data}(x_1 \mid c)$: action distribution for a fixed condition.

1. The conditional generator must learn endpoint actions paired with their observation c.
2. B retains the intended conditional law.
3. A can relabel a right action as left. C trains an unconditional model instead.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 55 · CHECK THE RESULT

Pairing and Optimal Transport

Re-pair examples with different conditions

Condition c says “move left”. Its data actions are negative. A batch also contains “move right” examples.

A. Re-pair all targets freely and retain old c.

B. Keep targets compatible with each condition. ✓

C. Remove c from training and still claim conditioning.

Tong et al. (2024), §§2.2–3.2 · original course examples

EXERCISE · SLIDE 60 · QUESTION

The Bridge to Diffusion

Worked bridge: a known Gaussian path

Fixed data $Y = 2$, $\tau = 1/2$, $x = 0$. Then $p_\tau = \mathcal{N}(1, 1/4)$. Find score and velocity.

A s=4, v=4

B s=2, v=2

C s=−4, v=−4

Lipman et al. (2023), §4 and Appendix D · original course examples

EXERCISE · SLIDE 60 · SOLUTION 1/2

The Bridge to Diffusion

Worked bridge: a known Gaussian path

Solution 1 / 2

Differentiate the Gaussian log density

$\mu = 1$, $\sigma^2 = 1/4$; $s(x) = -(x - \mu)/\sigma^2$.

1. $s(0) = -(0 - 1)/(1/4) = 4.$
2. $v(0, 1/2) = 0/(1/2) + [(1/2)/(1/2)]4 = 4.$

Lipman et al. (2023), §4 and Appendix D · original course examples

EXERCISE · SLIDE 60 · SOLUTION 2/2

The Bridge to Diffusion

Worked bridge: a known Gaussian path

Solution 2 / 2

Check by the endpoint-conditioned field

$y = 2$, $x = 0$, $\tau = 1/2$.

1. $u = (y - x)/(1 - \tau) = 2/(1/2) = 4$.
2. A agrees with both derivations. B uses standard deviation as variance; C reverses the Gaussian score sign.

Lipman et al. (2023), §4 and Appendix D · original course examples

EXERCISE · SLIDE 60 · CHECK THE RESULT

The Bridge to Diffusion

Worked bridge: a known Gaussian path

Fixed data $Y = 2$, $\tau = 1/2$, $x = 0$. Then $p_\tau = \mathcal{N}(1, 1/4)$. Find score and velocity.

- **A** s=4, v=4 ✓
- **B** s=2, v=2
- **C** s=−4, v=−4

Lipman et al. (2023), §4 and Appendix D · original course examples

EXERCISE · SLIDE 61 · QUESTION

The Bridge to Diffusion

Transfer: reverse the clock consistently

At course time $\tau = 1/2$, velocity is $+4$. Code time is $s = 1 - \tau$ and a sampling step has $\Delta s = -.1$.

A. Code velocity -4, state increment $+.4$

B. Code velocity $+4$, state increment $-.4$

C. Code velocity -4, state increment $-.4$

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 61 · SOLUTION 1/1

The Bridge to Diffusion

Transfer: reverse the clock consistently

Solution 1 / 1

Transform rate and increment together

$w(x, s) = -v(x, 1 - s)$; $\Delta s = -.1$.

1. $w = -4$.
2. $\Delta x \approx \Delta s\, w = (-.1)(-4) = +.4$.
3. A preserves the same forward generation direction. B misses the field sign; C multiplies signs incorrectly.

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 61 · CHECK THE RESULT

The Bridge to Diffusion

Transfer: reverse the clock consistently

At course time $\tau = 1/2$, velocity is $+4$. Code time is $s = 1 - \tau$ and a sampling step has $\Delta s = -.1$.

- **A** Code velocity -4, state increment $+.4$ ✓
- **B** Code velocity $+4$, state increment $-.4$
- **C** Code velocity -4, state increment $-.4$

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 66 · QUESTION

Vision-Language-Action Models

A code step toward the data

Action target $A = 2$, noise $\varepsilon = -1$. Code clock starts $s = 1$, $z = -1$; use exact target rate and $\Delta s = -1/4$.

A u=−3, new z=−.25

B u=+3, new z=−1.75

C u=−3, new z=−1.75

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 66 · SOLUTION 1/2

Vision-Language-Action Models

A code step toward the data

Solution 1 / 2

Follow the code convention

$u = \varepsilon - A$; $z_{new} = z + \Delta s\, u$.

1. $u = -1 - 2 = -3.$
2. $z_{new} = -1 + (-1/4)(-3) = -1 + 3/4 = -1/4.$

openpi: Pi0.compute_loss / sample_actions · original course examples

Vision-Language-Action Models

A code step toward the data

Solution 2 / 2

Check the interpolation at the new time

$s_{new} = 3/4$; $Z_s = (1-s)A + s\varepsilon$.

1. $Z_{3/4} = (1/4)2 + (3/4)(-1) = -1/4$.
2. A agrees. B reverses the target; C reverses the increment sign.

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 66 · CHECK THE RESULT

Vision-Language-Action Models

A code step toward the data

Action target $A = 2$, noise $\varepsilon = -1$. Code clock starts $s = 1$, $z = -1$; use exact target rate and $\Delta s = -1/4$.

- **A** u=−3, new z=−.25 ✓
- **B** u=+3, new z=−1.75
- **C** u=−3, new z=−1.75

openpi: Pi0.compute_loss / sample_actions · original course examples

EXERCISE · SLIDE 67 · QUESTION

Vision-Language-Action Models

Generation, execution, and replanning

Generate $H = 6$ actions. Execute the first $r = 2$, observe again, then generate a new chunk.

A. Six internal steps imply six robot actions.

B. Physical time advances by two executed actions.

C. The remaining four must always be executed.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples

Vision-Language-Action Models

Generation, execution, and replanning

Solution 1 / 1

Follow the physical loop

H: generated horizon; r: executed prefix length; $1 \leq r \leq H$.

1. The internal solver first produces a six-action candidate.
2. Executing two entries advances the robot by two action periods.
3. B is correct. A confuses clocks. C rules out the stated replanning procedure.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples

EXERCISE · SLIDE 67 · CHECK THE RESULT

Vision-Language-Action Models

Generation, execution, and replanning

Generate $H = 6$ actions. Execute the first $r = 2$, observe again, then generate a new chunk.

A Six internal steps imply six robot actions.

B Physical time advances by two executed actions. ✓

C The remaining four must always be executed.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples

EXERCISE · SLIDE 70 · QUESTION

Vision-Language-Action Models

What follows from a richer training mixture?

A VLA receives additional image-language and robot examples. Its low-level action model still uses flow matching.

- **A** The ODE becomes physical robot dynamics.
- **B** Representation learning changes; the flow objective can remain.
- **C** Generalization is guaranteed for every new home.

Physical Intelligence, $\pi_{0.5}$ (2025), §§IV-A–D · original course examples

EXERCISE · SLIDE 70 · SOLUTION 1/1

Vision-Language-Action Models

What follows from a richer training mixture?

Solution 1 / 1

Separate the mathematical objective from evidence

VLA: vision-language-action model; ODE: numerical action-generation flow.

1. Additional data changes the learned representation and parameters.
2. It does not redefine generation time as physical time or prove every new task succeeds.
3. B follows. A confuses the two systems; C overstates empirical evidence.

Physical Intelligence, $\pi_{0.5}$ (2025), §§IV-A–D · original course examples

EXERCISE · SLIDE 70 · CHECK THE RESULT

Vision-Language-Action Models

What follows from a richer training mixture?

A VLA receives additional image-language and robot examples. Its low-level action model still uses flow matching.

A. The ODE becomes physical robot dynamics.

B. Representation learning changes; the flow objective can remain. ✓

C. Generalization is guaranteed for every new home.

Physical Intelligence, $\pi_{0.5}$ (2025), §§IV-A–D · original course examples

EXERCISE · SLIDE 73 · QUESTION

From Data to an Executed Action

Transfer: the observation changes mid-execution

After one executed action, a new image shows an obstacle. Old condition c described a clear route.

A. More ODE steps under old c incorporate the new image.

B. Update the condition and replan under the policy.

C. Low training MSE guarantees the old chunk is safe.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples

EXERCISE · SLIDE 73 · SOLUTION 1/1

From Data to an Executed Action

Transfer: the observation changes mid-execution

Solution 1 / 1

Identify the changed input

c_{old} and c_{new} represent different observations.

1. Continuing integration under c_old keeps using the old information.
2. B passes the new condition into a new policy computation.
3. A confuses numerical accuracy with updated observations; C confuses regression fit with a safety guarantee.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples

EXERCISE · SLIDE 73 · CHECK THE RESULT

From Data to an Executed Action

Transfer: the observation changes mid-execution

After one executed action, a new image shows an obstacle. Old condition c described a clear route.

A. More ODE steps under old c incorporate the new image.

B. Update the condition and replan under the policy. ✓

C. Low training MSE guarantees the old chunk is safe.

Black et al., π_0 (2024/2026), §IV; App. A-D · original course examples