CIS 6280 · FALL 2026

WORLD MODELS

LECTURE 4

State-Space Models

Jiatao Gu

University of Pennsylvania · September 3, 2026

Before We Begin: Three Course Updates

01

Ed Discussion is set up

Please let us know if you have any questions.

02

Final exam: last week of teaching

The final exam will be held during the final teaching week.

03

Assignment 1 release is delayed

The new release date is next Thursday, September 10 (09/10).

A Simulator Makes the Hidden State Explicit

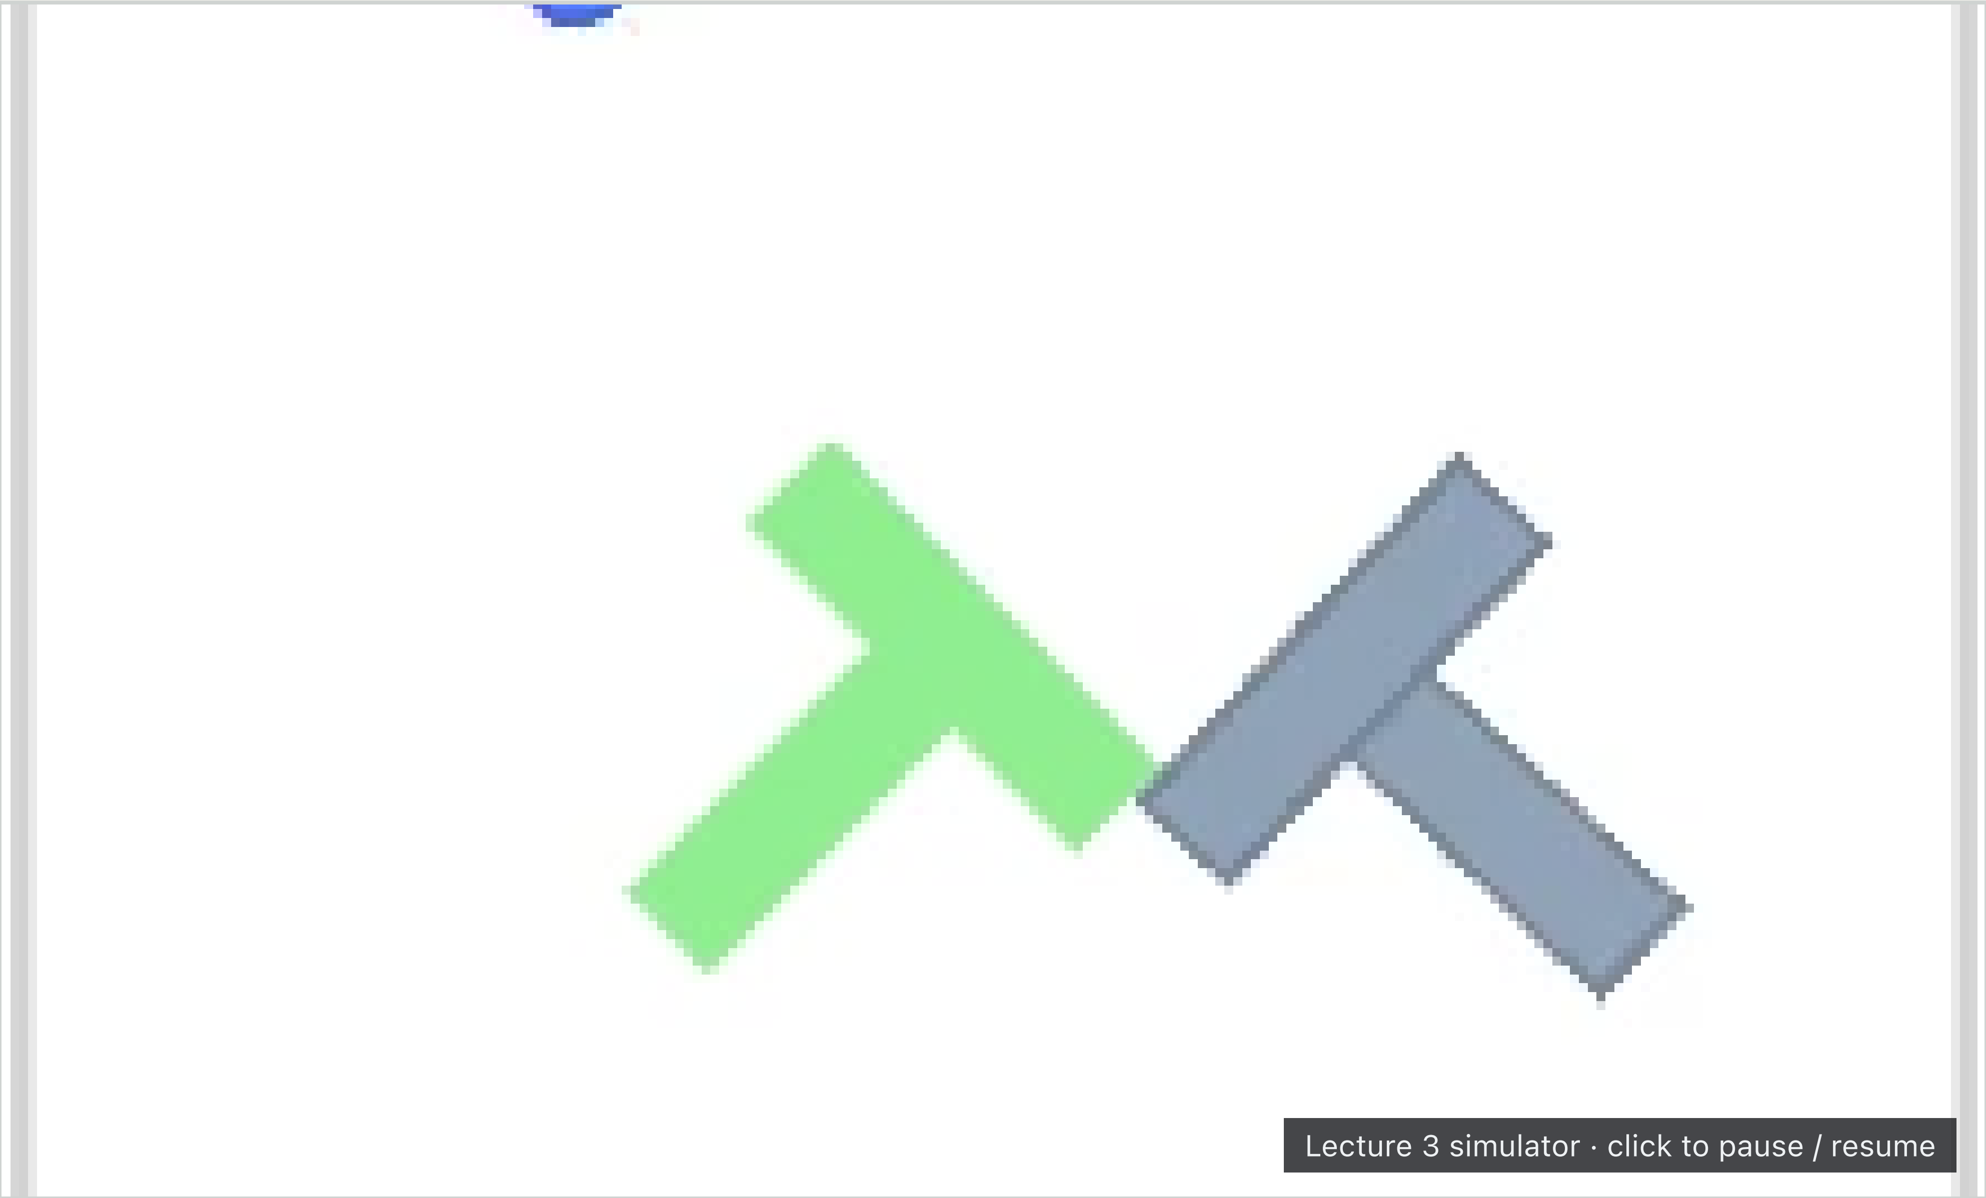

Pixels show outcomes.
Object pose and contact appear in o_t.

The simulator carries more.
Pose, velocity, and physical parameters determine what can happen next.

The step rule applies the action.

$$s_{t+1} = F(s_t, a_t)$$

Lecture 3 exposed s_t. A learner may receive only o_t.

History and Representation

$$p_\theta(Y_{\text{future}} \mid h_t, a_{\text{future}})$$

Predict future observations or other targets under a candidate action sequence.

$$h_t = (o_{0:t}, a_{0:t-1})$$

$$z_t = E_\phi(o_t;\, o_{<t}, a_{<t})$$

Contextualized representation: z_t represents the current observation using past context. Prediction may still need the sequence $z_{\le t}$.

Markov State and Contextualized Representation

MARKOV WORLD STATE

$$s_t$$

Given this state and the next action, the model can predict without earlier states.

CONTEXTUALIZED OBSERVATION

$$z_t = E_\phi(o_t;\ o_{<t}, a_{<t})$$

No Markov assumption. Prediction may still need earlier representations.

$$p(s_{t+1} \mid s_{0:t}, a_{0:t}) = p(s_{t+1} \mid s_t, a_t)$$

This Markov property concerns prediction **when the state is known**.

$$(z_{\le t}, a_{<t}) \xrightarrow{\text{state inference}} p(s_t \mid z_{\le t}, a_{<t})$$

Inferring an unknown state can still require history. The inference result describes uncertainty about s_t.

An SSM Is a World Model with a Markov State

INITIALIZE AND OBSERVE

Start a trajectory

$$s_0 \sim p_\theta(s_0), \quad o_0 \sim p_\theta(o_0 \mid s_0)$$

→

INTERVENE

Choose an action

$$a_t$$

→

EVOLVE

Generate the next state

$$s_{t+1} \sim p_\theta(s_{t+1} \mid s_t, a_t)$$

→

OBSERVE

Generate evidence

$$o_{t+1} \sim p_\theta(o_{t+1} \mid s_{t+1})$$

THE MODEL

Three probability laws

Initial state, action-conditioned transition, and state-conditioned observation.

FROM DATA

Estimate θ from trajectories

Observed sequences constrain the dynamics, sensor model, and uncertainty.

A CLASSICAL STARTING POINT

STATE · TRANSITION · UNCERTAINTY

Linear-Gaussian State-Space Models

R. E. KALMAN

A New Approach to Linear Filtering and Prediction Problems[1]

R. E. Kalman · *Journal of Basic Engineering* · March 1960

LINEAR-GAUSSIAN STATE-SPACE MODEL · LGSSM

Continuous vector state

$s_t \in \mathbb{R}^n$

Linear transition and sensor

$$s_{t+1} = As_t + Ba_t + \epsilon_t^s$$
$$o_t = Cs_t + \epsilon_t^o$$

Gaussian uncertainty

Initial state, process noise, and observation noise.

Under these assumptions, prediction and correction have **closed-form Gaussian updates**. A mean and covariance represent the full posterior.

The Linear-Gaussian State-Space Model

INITIAL-STATE ASSUMPTION

$s_0 \sim \mathcal{N}(\mu_0, P_0)$

A Gaussian distribution over where the trajectory begins.

TRANSITION ASSUMPTION

$s_{t+1} = As_t + Ba_t + \epsilon_t^s$

$\epsilon_t^s \sim \mathcal{N}(0, Q)$: linear dynamics with Gaussian process noise.

OBSERVATION ASSUMPTION

$o_t = Cs_t + \epsilon_t^o$

$\epsilon_t^o \sim \mathcal{N}(0, R)$: a linear sensor with Gaussian noise.

LGSSM independence: s_0 and all process and observation noise terms are mutually independent; the noise is also independent across time.

A 1D Example: Position, Velocity, and One Sensor

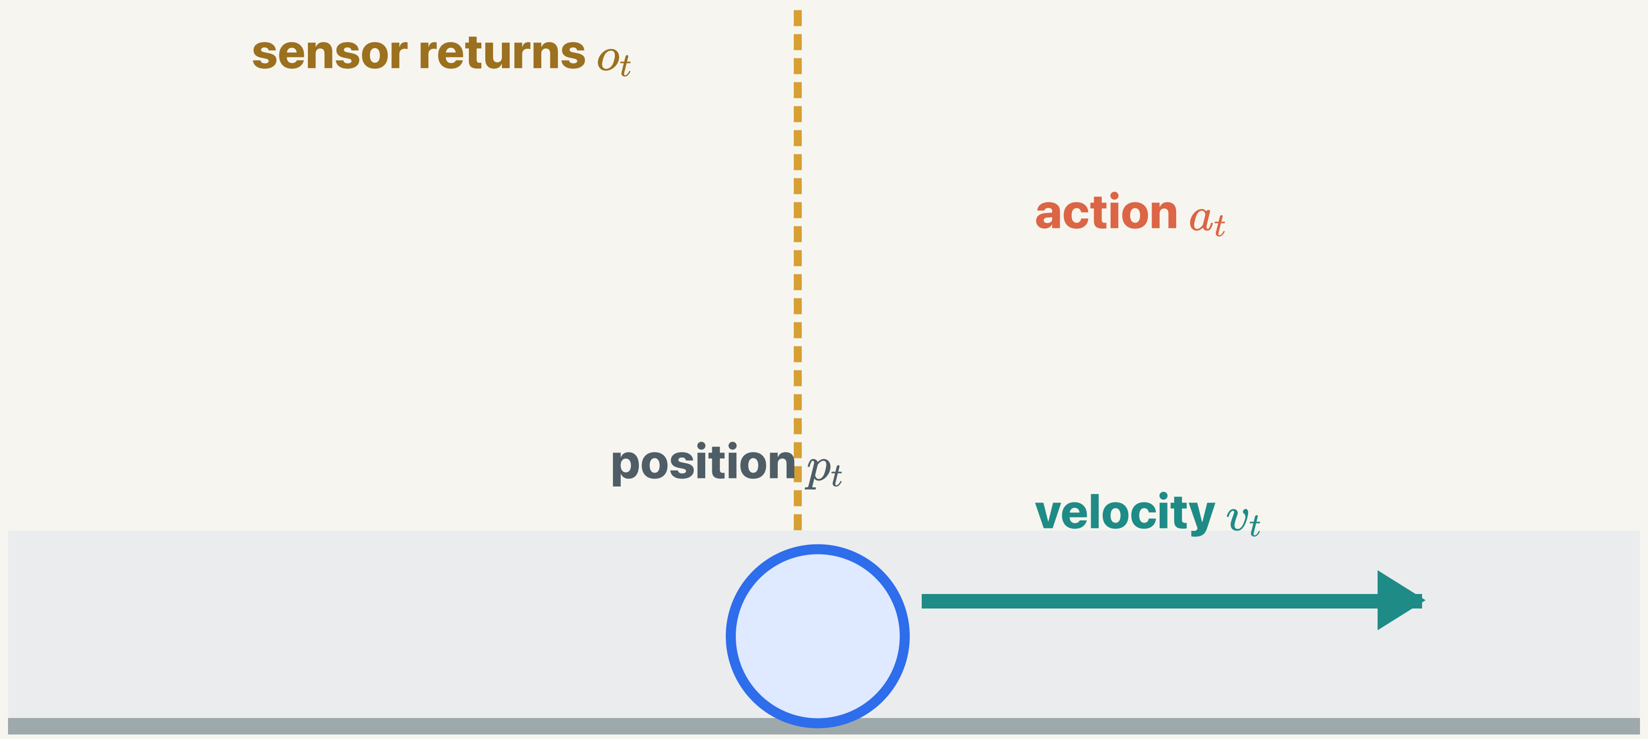

$$s_t = \begin{bmatrix} p_t \\ v_t \end{bmatrix}, \qquad o_t \in \mathbb{R}, \qquad a_t \in \mathbb{R}$$

STATE	OBSERVATION	ACTION
2D	**scalar**	**scalar**
position + velocity	noisy position	acceleration-like input

Here a_t is an acceleration-like input over one model step—not Lecture 3's relative target command.

The Object Moves; the Sensor Measures Position

HOW THE STATE EVOLVES

Position accumulates velocity and action

$$p_{t+1} = p_t + v_t + \frac{1}{2}a_t + \epsilon_t^p$$

$$v_{t+1} = v_t + a_t + \epsilon_t^v$$

WHAT THE SENSOR MEASURES

One noisy number: position

$$o_t = p_t + \epsilon_t^o$$

Velocity affects future positions, but it is not read directly by this sensor.

> Repeated position measurements can still reveal velocity because velocity leaves a trace in how position changes over time.

Read the Matrices as a Motion-and-Sensor Diagram

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \qquad B = \begin{bmatrix} \frac{1}{2} \\ 1 \end{bmatrix}, \qquad C = \begin{bmatrix} 1 & 0 \end{bmatrix}$$

$$s_{t+1} = As_t + Ba_t + \epsilon_t^s, \qquad o_t = Cs_t + \epsilon_t^o$$

One model step is one unit of time; Q and R describe process and sensor uncertainty.

CURRENT STATE

$s_t = [p_t, v_t]^\top$

A carry velocity into position
$p \leftarrow p + v,\ v \leftarrow v$

Ba_t add action to both coordinates
$p \leftarrow p + \frac{1}{2}a_t,\ v \leftarrow v + a_t$

NEXT STATE

$s_{t+1} = [p_{t+1}, v_{t+1}]^\top$

C read position; leave velocity hidden
$o_{t+1} = p_{t+1} + \epsilon_{t+1}^o$

The Same SSM Supports Modeling, Inference, and Learning

$$p_\theta(s_0), \qquad p_\theta(s_{t+1} \mid s_t, a_t), \qquad p_\theta(o_t \mid s_t)$$

$$p_\theta(s_{0:T}, o_{0:T} \mid a_{0:T-1}) = p_\theta(s_0) \prod_{t=0}^{T} p_\theta(o_t \mid s_t) \prod_{t=0}^{T-1} p_\theta(s_{t+1} \mid s_t, a_t)$$

WORLD MODEL / MODEL

Specify trajectory laws

The initial, transition, and observation laws define the trajectory distribution.

INFERENCE

Compute a hidden-state posterior

With the model fixed, compute $p_\theta(s_t \mid o_{0:t}, a_{0:t-1})$. Here: Kalman filtering.

LEARNING

Estimate model parameters

Estimate $\mu_0, P_0, A, B, C, Q, R$ by maximum likelihood; EM is one algorithm for doing so.

Can the Agent Infer the State Used for Prediction?

STATE USED BY THE MODEL

$$s_t = \begin{bmatrix} p_t \\ v_t \end{bmatrix}$$

Given the action, position and velocity determine the next-state distribution.

INFORMATION AVAILABLE TO THE AGENT

$$h_t = (o_{0:t}, a_{0:t-1})$$

$$h_t \overset{?}{\Longrightarrow} s_t$$

Permanent ambiguity: some distinct states can produce the same entire observation history.

Estimation uncertainty: the history contains the information, but measurements are noisy.

Structural question: after collecting an ideal observation window, can we distinguish different states at its start? This is the observability test.

Noise-Free Observability Test

THE ACTUAL LGSSM INCLUDES NOISE

$$s_{t+1} = As_t + Ba_t + \epsilon_t^s, \qquad o_t = Cs_t + \epsilon_t^o$$

STRUCTURAL TEST

$$\epsilon_t^s = 0, \qquad \epsilon_t^o = 0$$

and

KNOWN QUANTITIES

$A,\ B,\ C$ and the action sequence

Can we uniquely determine the starting state from a sequence of observations?

One Position Measurement Leaves Velocity Unknown

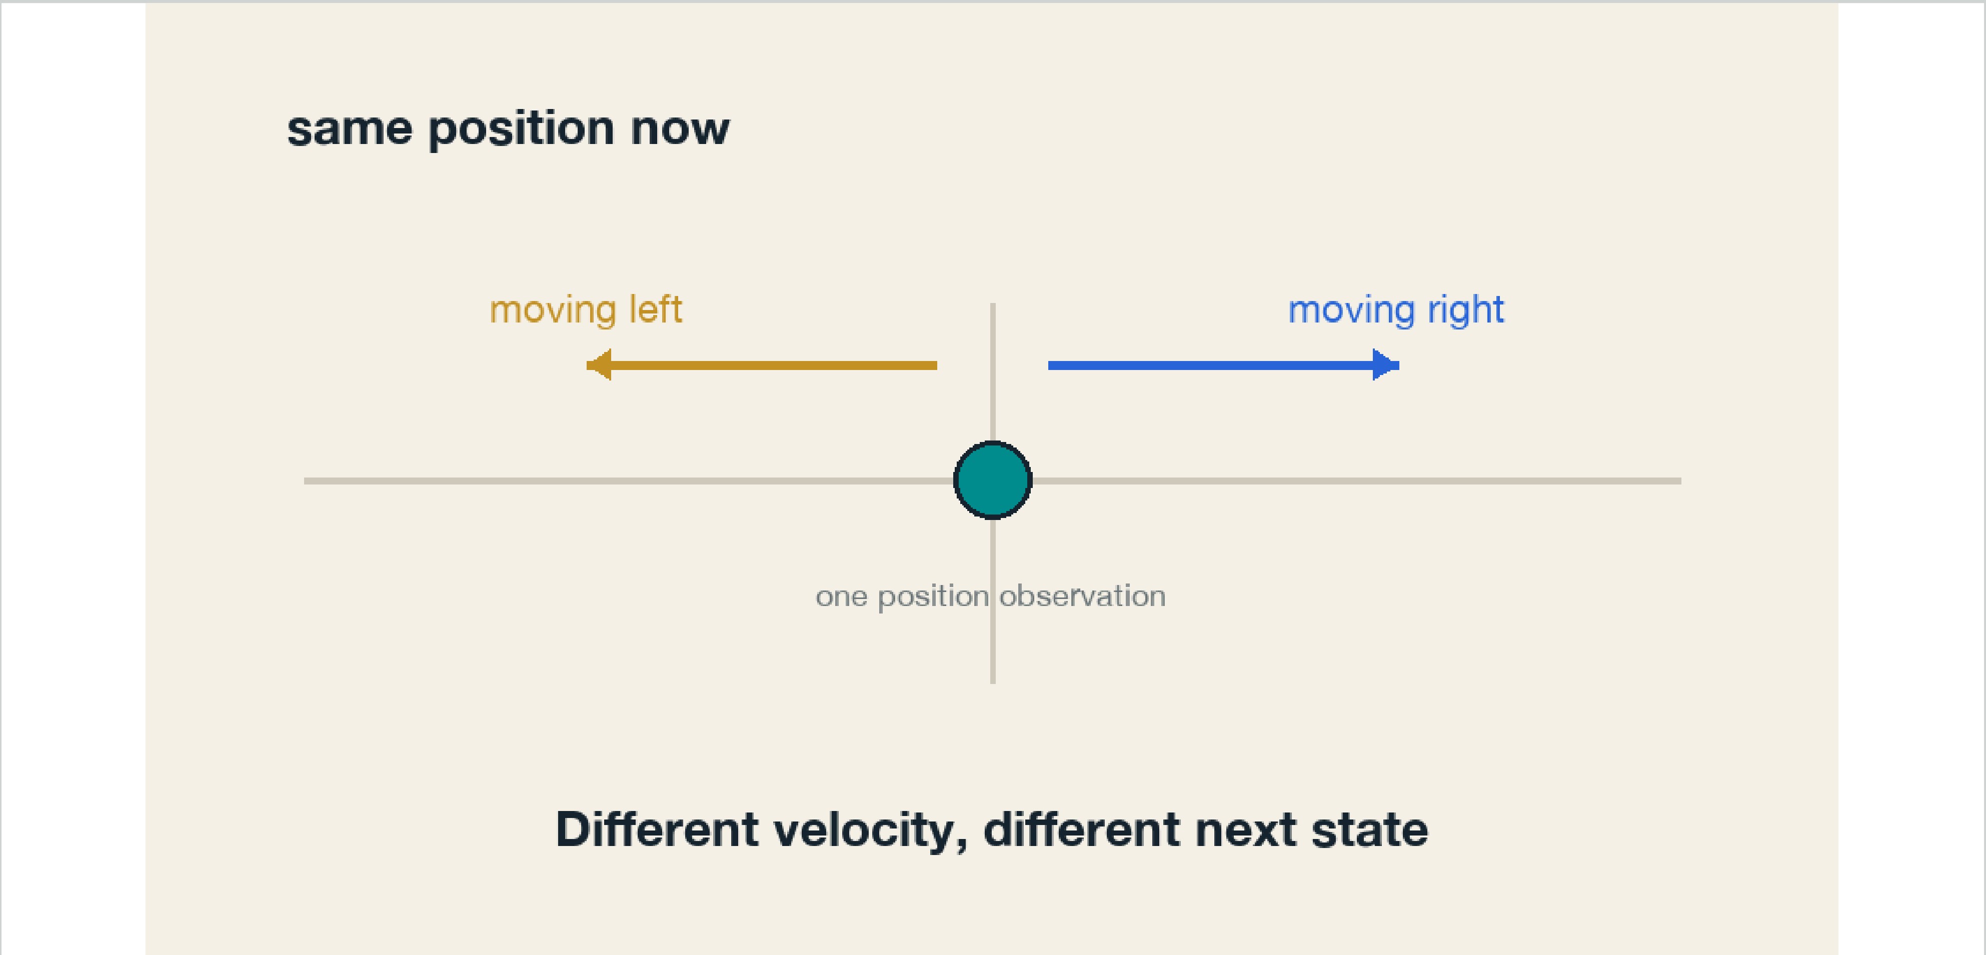

$$s_t = \begin{bmatrix} p_t \\ v_t \end{bmatrix}, \qquad o_t = p_t$$

$$\begin{bmatrix} p_t \\ v_t^{(1)} \end{bmatrix} \neq \begin{bmatrix} p_t \\ v_t^{(2)} \end{bmatrix}, \qquad o_t^{(1)} = o_t^{(2)} = p_t$$

The current measurement determines position.

Every velocity remains possible.

> A single frame does not contain enough information to recover the two-dimensional state.

1D Example: Two Observations

KNOWN MOTION AND SENSOR MODELS

$p_{t+1} = p_t + v_t + \frac{1}{2}a_t$

$$o_k = p_k$$

TWO IDEAL MEASUREMENTS

$o_t = p_t$

$$o_{t+1} - \tfrac{1}{2}a_t = p_t + v_t$$

$$p_t = o_t, \qquad v_t = o_{t+1} - o_t - \frac{1}{2}a_t$$

THE SAME TWO EQUATIONS IN MATRIX FORM

$$\begin{bmatrix} o_t \\ o_{t+1} - \frac{1}{2}a_t \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}}_{\mathcal{O}_2} \begin{bmatrix} p_t \\ v_t \end{bmatrix}, \qquad \det(\mathcal{O}_2) = 1 \neq 0, \quad \operatorname{rank}(\mathcal{O}_2) = 2$$

At time $t+1$, both observations are available and determine the earlier s_t. Velocity affects the later position even though the sensor measures only position.

The Observability Matrix

Collect $o_t, \ldots, o_{t+H-1}$ to distinguish states at the **start** of this noise-free window.

$$o_t = C s_t, \qquad o_{t+1} = CA s_t + CB a_t$$

$$\tilde{o}_{t:t+H-1} = \underbrace{\begin{bmatrix} C \\ CA \\ \vdots \\ CA^{H-1} \end{bmatrix}}_{\mathcal{O}_H \in \mathbb{R}^{dH \times n}} s_t, \qquad d = \dim(o_t)$$

Subtract the known action contributions, then stack the measurements.

$$\operatorname{rank}(\mathcal{O}_H) < n \;\Rightarrow\; \exists\, \delta \neq 0 :\; \mathcal{O}_H \delta = 0 \;\Rightarrow\; \mathcal{O}_H (s_t + \delta) = \mathcal{O}_H s_t$$

Dependent columns allow a nonzero state change that leaves all measurements unchanged. Full column rank removes this ambiguity.

> **One observation may not uniquely determine the state.** With a prior and an observation model, we can still estimate which states are more likely.

Hidden State and the Agent's Posterior

With noise, we estimate the current state and its uncertainty using the observations available so far.

The available history is $h_t = (o_{0:t}, a_{0:t-1})$.

WORLD STATE

$$s_t$$

One actual state of the world.

$\neq$

AGENT POSTERIOR

$$p(s_t \mid h_t)$$

Our estimate as a distribution over possible states.

> The transition law needs s_t. Because the agent does not observe it, the agent carries $p(s_t \mid h_t)$ into the next prediction.

One-Step State Prediction

$$p(s_{t+1} \mid h_t, a_t) = \int p(s_{t+1}, s_t \mid h_t, a_t)\, ds_t$$

Marginalize the unknown current state s_t.

$$= \int p(s_{t+1} \mid s_t, h_t, a_t)\, p(s_t \mid h_t, a_t)\, ds_t$$

Apply the product rule inside the integral.

SSM assumptions: the transition is Markov, and the action uses only h_t. Given that history, the action adds no information about s_t.

$$= \int \underbrace{p(s_{t+1} \mid s_t, a_t)}_{\text{transition law}} \; \underbrace{p(s_t \mid h_t)}_{\text{current posterior}} \; ds_t$$

Future Prediction from the Current Posterior

Definitions: Y_{future} is the future trajectory; a_{future} is a chosen open-loop action sequence.
Two reductions: $p(s_t \mid h_t, a_{\text{future}}) = p(s_t \mid h_t)$ and $Y_{\text{future}} \perp h_t \mid s_t, a_{\text{future}}$.

$$\begin{aligned} p(Y_{\text{future}} \mid h_t, a_{\text{future}}) &= \int p(Y_{\text{future}} \mid s_t, h_t, a_{\text{future}})\, p(s_t \mid h_t, a_{\text{future}})\, ds_t \\ &= \int p(Y_{\text{future}} \mid s_t, a_{\text{future}})\, p(s_t \mid h_t)\, ds_t. \end{aligned}$$

WORLD MODEL

Predict from each possible state

$$p(Y_{\text{future}} \mid s_t, a_{\text{future}})$$

AGENT POSTERIOR

Weight states using the history

$$p(s_t \mid h_t)$$

Bayesian Correction with the New Observation

$$h_{t+1} = (h_t, a_t, o_{t+1})$$

Observation assumption: $o_{t+1} \perp (h_t, a_t) \mid s_{t+1}$. The current state determines the sensor distribution.

$$\begin{aligned} p(s_{t+1} \mid h_{t+1}) &= \frac{p(o_{t+1} \mid s_{t+1}, h_t, a_t)\, p(s_{t+1} \mid h_t, a_t)}{p(o_{t+1} \mid h_t, a_t)} \\ &= \frac{\underbrace{p(o_{t+1} \mid s_{t+1})}_{\text{observation likelihood}} \; \underbrace{p(s_{t+1} \mid h_t, a_t)}_{\text{prediction}}}{\underbrace{p(o_{t+1} \mid h_t, a_t)}_{\text{normalizer}}}. \end{aligned}$$

The observation reweights the predicted next states. The result is the posterior for time $t+1$.

Predict–Correct Recursion

AFTER OBSERVING O_T

current posterior

$p(s_t \mid h_t)$

predict

AFTER ACTION A_T

next-state prediction

$p(s_{t+1} \mid h_t, a_t)$

correct

AFTER OBSERVING O_{T+1}

updated posterior

$p(s_{t+1} \mid h_{t+1})$

> $p(s_{t+1} \mid h_{t+1})$ has the same role at time $t+1$ that $p(s_t \mid h_t)$ had at time t. It becomes the input to the next prediction.

Exact Recursive Inference for an LGSSM

Rudolf E. Kálmán

A New Approach to Linear Filtering and Prediction Problems

PUBLISHED MARCH 1960

LGSSM:
the world-model probability laws.

Kalman filter:
exact recursive inference under those laws.

LGSSM Assumptions for Exact Filtering

$$s_{t+1} = As_t + Ba_t + \epsilon_t^s, \qquad o_t = Cs_t + \epsilon_t^o$$

$$s_0 \sim \mathcal{N}(\mu_0, P_0), \qquad \epsilon_t^s \sim \mathcal{N}(0, Q), \qquad \epsilon_t^o \sim \mathcal{N}(0, R)$$

s_0 and all noise variables are mutually independent across sources and time. Covariances satisfy $P_0 \succeq 0$, $Q \succeq 0$, and $R \succ 0$.

INITIALIZE AFTER O_0

$$p(s_0 \mid h_0) = p(s_0 \mid o_0) = \mathcal{N}(\mu_{0|0}, P_{0|0})$$

Gaussian conditioning gives the first filtered posterior.

PREDICT

Affine Gaussian closure

A linear transition plus independent Gaussian process noise gives a Gaussian prediction.

CORRECT

Gaussian conditioning

A linear-Gaussian observation gives another Gaussian posterior.

$$p(s_t \mid h_t) = \mathcal{N}(\mu_{t|t}, P_{t|t}) \quad \Longrightarrow \quad \text{repeat for every } t$$

Prediction and Filtering Notation

The first index is the state time. The second index is the latest observation time used.

BEFORE OBSERVING o_{T+1}

$$\mu_{t+1|t} := \mathbb{E}[s_{t+1} \mid h_t, a_t]$$

$$P_{t+1|t} := \mathrm{Cov}(s_{t+1} \mid h_t, a_t)$$

AFTER OBSERVING o_{T+1}

$$\mu_{t+1|t+1} := \mathbb{E}[s_{t+1} \mid h_{t+1}]$$

$$P_{t+1|t+1} := \mathrm{Cov}(s_{t+1} \mid h_{t+1})$$

$$(\mu^-, P^-) := (\mu_{t+1|t}, P_{t+1|t}), \qquad (\mu^+, P^+) := (\mu_{t+1|t+1}, P_{t+1|t+1})$$

(μ, P) represents the entire Gaussian posterior. The point estimate $\hat{s} = \mu$ alone does not represent its uncertainty.

Predicted State Mean

The policy selects a_t using only h_t. Therefore $p(s_t \mid h_t, a_t) = p(s_t \mid h_t)$ and $\mathbb{E}[s_t \mid h_t, a_t] = \mu_{t|t}$.

$$\begin{aligned} \mu_{t+1|t} &= \mathbb{E}[s_{t+1} \mid h_t, a_t] \\ &= \mathbb{E}[As_t + Ba_t + \epsilon_t^s \mid h_t, a_t]. \end{aligned}$$

STATE TERM

$$\mathbb{E}[s_t \mid h_t, a_t] = \mu_{t|t}$$

INDEPENDENT ZERO-MEAN PROCESS NOISE

$$\mathbb{E}[\epsilon_t^s \mid h_t, a_t] = 0$$

$$\mu_{t+1|t} = A\mu_{t|t} + Ba_t$$

Known Action and Conditional Covariance

$$\mu_{t+1|t} = A\mu_{t|t} + Ba_t$$

and

$$\operatorname{Cov}(Ba_t \mid h_t, a_t) = 0$$

A known action shifts every candidate state by the same vector. Uncertain actuation would require an additional random term or a larger process covariance.

Predicted State Covariance

$$e_t := s_t - \mu_{t|t}, \qquad e_{t+1|t} := s_{t+1} - \mu_{t+1|t} = Ae_t + \epsilon_t^s$$

$$\begin{aligned} P_{t+1|t} &= \mathbb{E}[e_{t+1|t}e_{t+1|t}^\top \mid h_t, a_t] \\ &= AP_{t|t}A^\top + Q \\ &\quad + A\mathbb{E}[e_t(\epsilon_t^s)^\top \mid h_t, a_t] + \mathbb{E}[\epsilon_t^s e_t^\top \mid h_t, a_t]A^\top. \end{aligned}$$

Process noise at time t is independent of the current error, so both cross expectations equal zero.

$$P_{t+1|t} = AP_{t|t}A^\top + Q$$

EXISTING UNCERTAINTY

$$AP_{t|t}A^\top$$

The dynamics rotate and scale the current error.

NEW UNCERTAINTY

$$Q$$

Independent process noise adds covariance.

Predictive State Distribution

$$p(s_t \mid h_t) = \mathcal{N}(\mu_{t|t}, P_{t|t}), \qquad p(s_{t+1} \mid s_t, a_t) = \mathcal{N}(As_t + Ba_t, Q)$$

$$p(s_{t+1} \mid h_t, a_t) = \int p(s_{t+1} \mid s_t, a_t)\, p(s_t \mid h_t)\, ds_t$$

$$= \mathcal{N}\left(s_{t+1}; \underbrace{A\mu_{t|t} + Ba_t}_{\mu_{t+1|t}}, \underbrace{AP_{t|t}A^\top + Q}_{P_{t+1|t}} \right)$$

> The previous pages computed the mean and covariance of this integral. Gaussian closure supplies the distribution family.

Predicted Observation Mean

$$\begin{aligned}\hat{o}_{t+1|t} &:= \mathbb{E}[o_{t+1} \mid h_t, a_t] \\ &= \mathbb{E}[Cs_{t+1} + \epsilon^o_{t+1} \mid h_t, a_t] \\ &= C\mu_{t+1|t}.\end{aligned}$$

$\mu_{t+1|t} \in \mathbb{R}^n$

C

$\hat{o}_{t+1|t} \in \mathbb{R}^d$

$C \in \mathbb{R}^{d\times n}$ maps latent state coordinates into observation coordinates. The hat denotes a prediction made before seeing o_{t+1}.

Predicted Observation Distribution

$$o_{t+1} - C\mu_{t+1|t} = C(s_{t+1} - \mu_{t+1|t}) + \epsilon^o_{t+1}$$

$$S_{t+1} := \mathrm{Cov}(o_{t+1} \mid h_t, a_t) = CP_{t+1|t}C^\top + R$$

$$p(o_{t+1} \mid h_t, a_t) = \mathcal{N}\left(C\mu_{t+1|t}, S_{t+1}\right)$$

BEFORE THE OBSERVATION ARRIVES

Predict its distribution

This is the denominator in the generic Bayes correction.

AFTER THE OBSERVATION ARRIVES

$$\nu_{t+1} := o_{t+1} - C\mu_{t+1|t}$$

The realized observation-space residual.

Joint State–Observation Distribution

$$\mu^- := \mu_{t+1|t}, \quad P^- := P_{t+1|t}, \quad S := S_{t+1}, \quad \nu := \nu_{t+1}$$

$$\begin{aligned}\mathrm{Cov}(s_{t+1}, o_{t+1} \mid h_t, a_t) &= \mathrm{Cov}(s_{t+1}, Cs_{t+1} + \epsilon^o_{t+1} \mid h_t, a_t) \\ &= P^- C^\top.\end{aligned}$$

$$\begin{bmatrix} s_{t+1} \\ o_{t+1} \end{bmatrix} \Bigg| \, h_t, a_t \sim \mathcal{N}\left(\begin{bmatrix} \mu^- \\ C\mu^- \end{bmatrix}, \begin{bmatrix} P^- & P^- C^\top \\ CP^- & S \end{bmatrix} \right)$$

STATE BLOCK

$$P^-$$

STATE–SENSOR BLOCK

$$P^- C^\top$$

SENSOR BLOCK

$$S = CP^- C^\top + R$$

State Error and Observation Error

Goal: split the state error into a part determined by the observation and a residual independent of it.

All expectations and covariances below condition on the known h_t, a_t.

$$e_s := s_{t+1} - \mu^-, \qquad e_o := o_{t+1} - C\mu^-$$

Gaussian property: jointly Gaussian variables with zero cross-covariance are independent.

For any matrix $L \in \mathbb{R}^{n \times d}$, write:

$$e_s = Le_o + r, \qquad r := e_s - Le_o$$

Choose L to make r independent of e_o. For this Gaussian pair, zero cross-covariance is enough:

$$0 = \mathrm{Cov}(r, e_o) = \underbrace{P^- C^\top}_{\mathrm{Cov}(e_s, e_o)} - L \underbrace{S}_{\mathrm{Cov}(e_o)} \quad \Rightarrow \quad L = P^- C^\top S^{-1} =: K^{\mathrm{KF}}$$

Conditional Mean and Covariance

Keep h_t, a_t fixed. From the previous page, $r = e_s - K^{\mathrm{KF}} e_o$ is independent of o_{t+1} and has mean zero.

$$s_{t+1} = \mu^- + K^{\mathrm{KF}}(o_{t+1} - C\mu^-) + r$$

$$\mu^+ := \mathbb{E}[s_{t+1} \mid h_t, a_t, o_{t+1}] = \mu^- + K^{\mathrm{KF}} \nu$$

After observing o_{t+1}, only r remains uncertain. Write $K = K^{\mathrm{KF}}$ in the expansion below.

$$\begin{aligned} P^+ &= \mathrm{Cov}(r) = P^- - KCP^- - P^- C^\top K^\top + KSK^\top \\ &= P^- - KSK^\top, \qquad KS = P^- C^\top. \end{aligned}$$

> r is Gaussian, so the conditional distribution is Gaussian with mean μ^+ and covariance P^+.

Kalman Correction

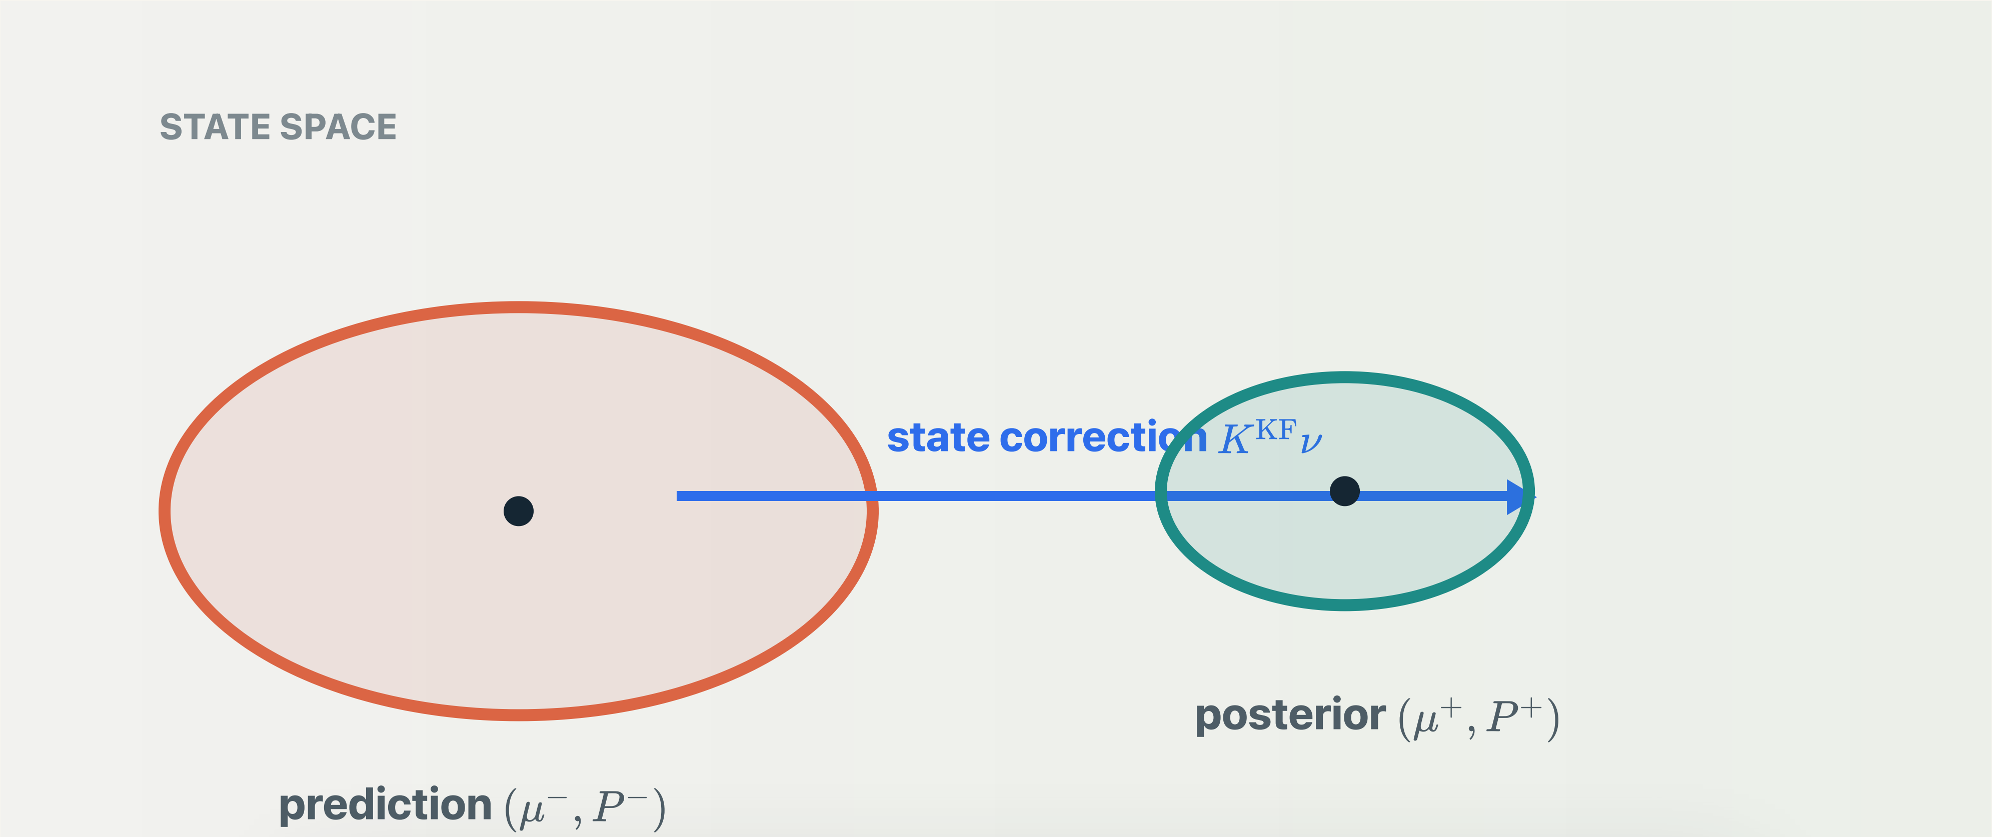

$$\nu \in \mathbb{R}^d \xrightarrow{K^{\mathrm{KF}}} K^{\mathrm{KF}}\nu \in \mathbb{R}^n$$

$$K^{\mathrm{KF}} = P^- C^\top S^{-1}$$

$$\mu^+ = \mu^- + K^{\mathrm{KF}}\nu$$

$$\begin{aligned} P^+ &= P^- - P^- C^\top S^{-1} C P^- \\ &= P^- - K^{\mathrm{KF}} S (K^{\mathrm{KF}})^\top. \end{aligned}$$

The observation residual ν lives in observation space. The mean moves in state space by $K^{\mathrm{KF}}\nu$.

One Complete Kalman Step

Start with $p(s_t \mid h_t) = \mathcal{N}(\mu_{t|t}, P_{t|t})$. The model parameters remain fixed.

1. PREDICT AFTER ACTION A_T

$$\mu^- = \mu_{t+1|t} = A\mu_{t|t} + Ba_t$$

$$P^- = P_{t+1|t} = AP_{t|t}A^\top + Q$$

2. CORRECT AFTER OBSERVATION O_{T+1}

$$\nu = o_{t+1} - C\mu^-, \quad S = CP^-C^\top + R$$

$$K = P^-C^\top S^{-1}$$

$$\mu_{t+1|t+1} = \mu^- + K\nu$$

$$P_{t+1|t+1} = P^- - KSK^\top$$

$$(\mu_{t|t}, P_{t|t}) \xrightarrow{a_t,\, o_{t+1}} (\mu_{t+1|t+1}, P_{t+1|t+1})$$

Use this new pair as the starting pair for the next time step. It summarizes the history needed to predict the future under this LGSSM.

Position–Velocity Cross-Covariance

$$s_{t+1} = \begin{bmatrix} p_{t+1} \\ v_{t+1} \end{bmatrix}, \qquad C = \begin{bmatrix} 1 & 0 \end{bmatrix}, \qquad o_{t+1} = p_{t+1} + \epsilon^o_{t+1}$$

$$P_{t+1|t} = \begin{bmatrix} \mathrm{Var}(p_{t+1} \mid h_t, a_t) & \mathrm{Cov}(p_{t+1}, v_{t+1} \mid h_t, a_t) \\ \mathrm{Cov}(v_{t+1}, p_{t+1} \mid h_t, a_t) & \mathrm{Var}(v_{t+1} \mid h_t, a_t) \end{bmatrix}$$

DIAGONAL ENTRIES

Predicted position and velocity uncertainty

How uncertain each next-state component is before o_{t+1}.

OFF-DIAGONAL ENTRIES

How the two prediction errors vary together

$$P^-_{vp} := \mathrm{Cov}(v_{t+1}, p_{t+1} \mid h_t, a_t)$$

When $P^-_{vp} \neq 0$, the velocity entry of $P^- C^\top$ is nonzero. A nonzero position innovation can therefore change the velocity mean.

Numerical Prediction

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \qquad B = \begin{bmatrix} 1/2 \\ 1 \end{bmatrix}, \qquad C = \begin{bmatrix} 1 & 0 \end{bmatrix}$$

$$\mu_{0|0} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad P_{0|0} = \begin{bmatrix} 1/2 & 0 \\ 0 & 1 \end{bmatrix}, \quad Q = \begin{bmatrix} 3/2 & 0 \\ 0 & 1 \end{bmatrix}, \quad R = 1, \quad a_0 = 1$$

Initialization: the prior is $\mathcal{N}([0,1]^\top, I)$. Observing $o_0 = 0$ gives the filtered mean and covariance above.

$$\mu_{1|0} = A\mu_{0|0} + Ba_0 = \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \begin{bmatrix} 1/2 \\ 1 \end{bmatrix} = \begin{bmatrix} 3/2 \\ 2 \end{bmatrix}$$

$$P_{1|0} = AP_{0|0}A^\top + Q = \begin{bmatrix} 3/2 & 1 \\ 1 & 1 \end{bmatrix} + \begin{bmatrix} 3/2 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 3 & 1 \\ 1 & 2 \end{bmatrix}$$

The known action changes the mean. The covariance combines propagated uncertainty with process noise.

Numerical Correction

$$\mu_{1|0} = \begin{bmatrix} 3/2 \\ 2 \end{bmatrix}, \quad P_{1|0} = \begin{bmatrix} 3 & 1 \\ 1 & 2 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 \end{bmatrix}, \quad R = 1$$

INNOVATION

$$o_1 = \frac{5}{2}$$

$$\nu_1 = \frac{5}{2} - \frac{3}{2} = 1$$

$$S_1 = C P_{1|0} C^\top + R = 4$$

GAIN

$$K_1^{\mathrm{KF}} = P_{1|0} C^\top S_1^{-1} = \begin{bmatrix} 3/4 \\ 1/4 \end{bmatrix}$$

The velocity entry is nonzero because predicted position and velocity are correlated.

POSTERIOR MEAN

$$\mu_{1|1} = \begin{bmatrix} 3/2 \\ 2 \end{bmatrix} + \begin{bmatrix} 3/4 \\ 1/4 \end{bmatrix} (1) = \begin{bmatrix} 9/4 \\ 9/4 \end{bmatrix}$$

The position-only observation changes the velocity estimate.

$$P_{1|1} = P_{1|0} - K_1^{\mathrm{KF}} S_1 (K_1^{\mathrm{KF}})^\top = \begin{bmatrix} 3/4 & 1/4 \\ 1/4 & 7/4 \end{bmatrix}$$

Cross-covariance carries information: a position-only sensor can update velocity.

Marginal Likelihood for an LGSSM

Inference: we treated θ as fixed and estimated hidden states. **Learning:** θ is unknown, the states remain hidden, and observed trajectories determine the fit.

T: final time index · observed: $o_{0:T}$ · known inputs: $a_{0:T-1}$ · hidden: $s_{0:T}$

$$\theta = (\mu_0, P_0, A, B, C, Q, R)$$

$$p_\theta(s_{0:T}, o_{0:T} \mid a_{0:T-1}) = p_\theta(s_0) \prod_{t=0}^{T-1} p_\theta(s_{t+1} \mid s_t, a_t) \prod_{t=0}^{T} p_\theta(o_t \mid s_t)$$

$$p_\theta(o_{0:T} \mid a_{0:T-1}) = \int p_\theta(s_{0:T}, o_{0:T} \mid a_{0:T-1})\, ds_{0:T}$$

$$\theta^\star = \arg\max_\theta \ \ell(\theta), \qquad \ell(\theta) := \log p_\theta(o_{0:T} \mid a_{0:T-1})$$

Kalman Predictions Compute the Marginal Likelihood

Likelihood convention: treat $a_{0:T-1}$ as fixed, exogenous inputs.
For $t \geq 1$, $h_{t-1} := (o_{0:t-1}, a_{0:t-2})$; let $d := \dim(o_t)$.

$$p_\theta(o_{0:T} \mid a_{0:T-1}) = p_\theta(o_0) \prod_{t=1}^{T} p_\theta(o_t \mid h_{t-1}, a_{t-1})$$

$$p_\theta(s_t \mid h_{t-1}, a_{t-1}) = \mathcal{N}(\mu_{t|t-1}, P_{t|t-1}), \qquad p_\theta(o_t \mid s_t) = \mathcal{N}(o_t;\, C s_t, R)$$

REUSE THE EARLIER PREDICTED-OBSERVATION GAUSSIAN

$$\begin{aligned} p_\theta(o_t \mid h_{t-1}, a_{t-1}) &= \int p_\theta(o_t \mid s_t)\, p_\theta(s_t \mid h_{t-1}, a_{t-1})\, ds_t \\ &= \mathcal{N}\big(o_t;\, C\mu_{t|t-1},\, S_t\big) \\ S_t &:= C P_{t|t-1} C^\top + R \end{aligned}$$

After the sensor reports o_t, its residual from the predicted mean is $\nu_t := o_t - C\mu_{t|t-1}$.

One-Observation Negative Log-Likelihood

$$\ell_t(\theta) := \log p_\theta(o_t \mid h_{t-1}, a_{t-1}), \qquad \mathcal{L}_t(\theta) := -\ell_t(\theta), \quad t \geq 1$$

$$\mathcal{L}_t(\theta) = \frac{1}{2}\left[d \log(2\pi) + \log|S_t| + \nu_t^\top S_t^{-1} \nu_t\right]$$

NORMALIZED PREDICTION ERROR

$$\nu_t^\top S_t^{-1} \nu_t$$

Large when the observation is far from its predicted mean relative to uncertainty.

+

PREDICTIVE WIDTH

$$\log|S_t|$$

Large when the model assigns a broad distribution to outcomes.

> $S_t = C P_{t|t-1} C^\top + R$: increasing predictive covariance can reduce the normalized error, but it increases $\log|S_t|$.

Sequence Negative Log-Likelihood

INITIAL OBSERVATION

$$\ell_0(\theta) = \log \mathcal{N}(o_0;\, C\mu_0, S_0)$$

$$\nu_0 = o_0 - C\mu_0, \qquad S_0 = CP_0C^\top + R$$

LATER OBSERVATIONS · $T \geq 1$

$$\ell_t(\theta) = \log \mathcal{N}(o_t;\, C\mu_{t|t-1}, S_t)$$

The forward Kalman recursion supplies ν_t, S_t.

$$\ell(\theta) = \ell_0(\theta) + \sum_{t=1}^{T} \ell_t(\theta), \qquad \mathcal{L}(\theta) := -\ell(\theta)$$

$$\mathcal{L}(\theta) = \frac{1}{2} \sum_{t=0}^{T} \left[d \log(2\pi) + \log|S_t| + \nu_t^\top S_t^{-1} \nu_t \right]$$

The prior $\mathcal{N}(\mu_0, P_0)$ defines the initial $t = 0$ factor; history-indexed predictive factors begin at $t = 1$.

Differentiate Through the Kalman Forward Pass

$$\mathcal{L}(\theta) = -\ell(\theta), \qquad \theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}(\theta), \qquad \eta > 0$$

PARAMETERS

θ

$\mu_0, P_0, A, B, C, Q, R$

→

FORWARD PASS

Kalman recursion

Predict, evaluate the likelihood of o_t under that prediction, then correct with o_t. Repeat.

→

LOSS

$\mathcal{L}(\theta)$

Sum terms from $\nu_t(\theta), S_t(\theta)$ saved before each correction.

→

GRADIENT STEP

Automatic differentiation

Propagate derivatives through the full recursion.

> Here η is the learning rate. Changing θ changes every later prediction, so the gradient passes through all predict–correct steps. The covariance constraints are $P_0, Q \succeq 0$ and usually $R \succ 0$, consistent with an invertible innovation covariance S_t.

Next: EM optimizes the same marginal-likelihood objective with an explicit posterior over hidden trajectories.

Expectation–maximization (EM)

$$\max_{\theta} \log p_{\theta}(o_{0:T} \mid a_{0:T-1})$$

The same MLE objective. Direct gradients differentiate the Kalman likelihood. EM alternates the following steps.

E-STEP

Estimate hidden-state uncertainty

Fix the parameters. Use the complete recorded sequence to estimate a posterior over hidden trajectories. This offline inference is **smoothing**.

M-STEP

Update the model parameters

Fit the model using posterior expectations of the hidden states and their products. Linear-Gaussian models often allow closed-form updates.

Why it works: EM improves a likelihood lower bound that touches the current likelihood. Exact E- and M-steps cannot decrease that likelihood.

Tradeoff: convenient updates for classical LGSSMs, but no guarantee of faster convergence or the same local solution as gradients.

LGSSM: Model, Inference, and Learning

World model	Initial, transition, and observation distributions parameterized by θ
Prediction	Future observations $p(ot{+}1 \mid ht, at; \theta)$
Optimizer	Direct gradient optimization or expectation–maximization
Filtering	Online posterior $p(st \mid o0{:}t, a0{:}t{-}1; \theta)$
Objective	Marginal log likelihood $\ell(\theta) = \log p(o0{:}T \mid a0{:}T{-}1; \theta)$

Hidden Markov Model (HMM)

A discrete-state SSM with L possible hidden states.

STATE SPACE

$$s_t \in \{1, \dots, L\}$$

A finite set of latent modes.

INITIAL LAW

$$p(s_0)$$

How a sequence begins.

TRANSITION

$$p(s_{t+1} \mid s_t)$$

Move between modes.

OBSERVATION

$$p(o_t \mid s_t)$$

Each mode explains observations.

$$p(s_t \mid o_{0:t}) \longleftrightarrow \big(p(s_t = 1 \mid o_{0:t}), \dots, p(s_t = L \mid o_{0:t})\big)$$

HMM Filtering with Finite Sums

$$p(s_{t+1} = j \mid o_{0:t}) = \sum_{i=1}^{L} p(s_{t+1} = j \mid s_t = i)\, p(s_t = i \mid o_{0:t})$$

$$p(s_{t+1} = j \mid o_{0:t+1}) \propto p(o_{t+1} \mid s_{t+1} = j)\, p(s_{t+1} = j \mid o_{0:t})$$

LGSSM POSTERIOR

$$(\mu_{t|t}, P_{t|t})$$

Gaussian parameters

HMM POSTERIOR

$$\big(p(s_t = 1 \mid o_{0:t}), \ldots, p(s_t = L \mid o_{0:t})\big)$$

Categorical probabilities

Learning: maximize the marginal likelihood using direct gradients or EM. For an HMM, EM uses forward–backward inference over discrete states.

A Valid SSM State Makes the Dynamics Markov

1

Position alone can omit the velocity needed for prediction.

$$p(s_{t+1} \mid s_{0:t}, a_{0:t}) = p(s_{t+1} \mid s_t, a_t)$$

If this equality fails, the chosen coordinates are not a sufficient Markov state.

State repair: augment the state or learn a representation that retains the missing predictive information.

Nonlinear Dynamics

2

$$s_{t+1} = As_t + Ba_t + \epsilon_t^s$$

Contact and saturation change how the next state responds at different parts of the state space.

Nonlinear SSM: replace the linear map with $s_{t+1} = f_\theta(s_t, a_t) + \epsilon_t^s$, where f_θ can model contact or saturation. The prediction may no longer be Gaussian.

Image Observations Need a Nonlinear Observation Model

3

$$o_t = Cs_t + \epsilon_t^o$$

Pixels depend nonlinearly on geometry, perspective, lighting, visibility, and occlusion.

Nonlinear observation SSM: use $o_t \sim p_\theta(o_t \mid s_t)$. The standard Kalman correction is generally no longer exact.

A Single Gaussian Merges Distinct Future Modes

4

Kalman prediction represents the next-state distribution as one Gaussian.

$$p(s_{t+1} \mid h_t, a_t) = \mathcal{N}(\mu_{t+1|t}, P_{t+1|t})$$

Separated plausible futures can place the Gaussian mean on an implausible trajectory.

Richer uncertainty: represent multiple modes rather than forcing single-Gaussian closure.

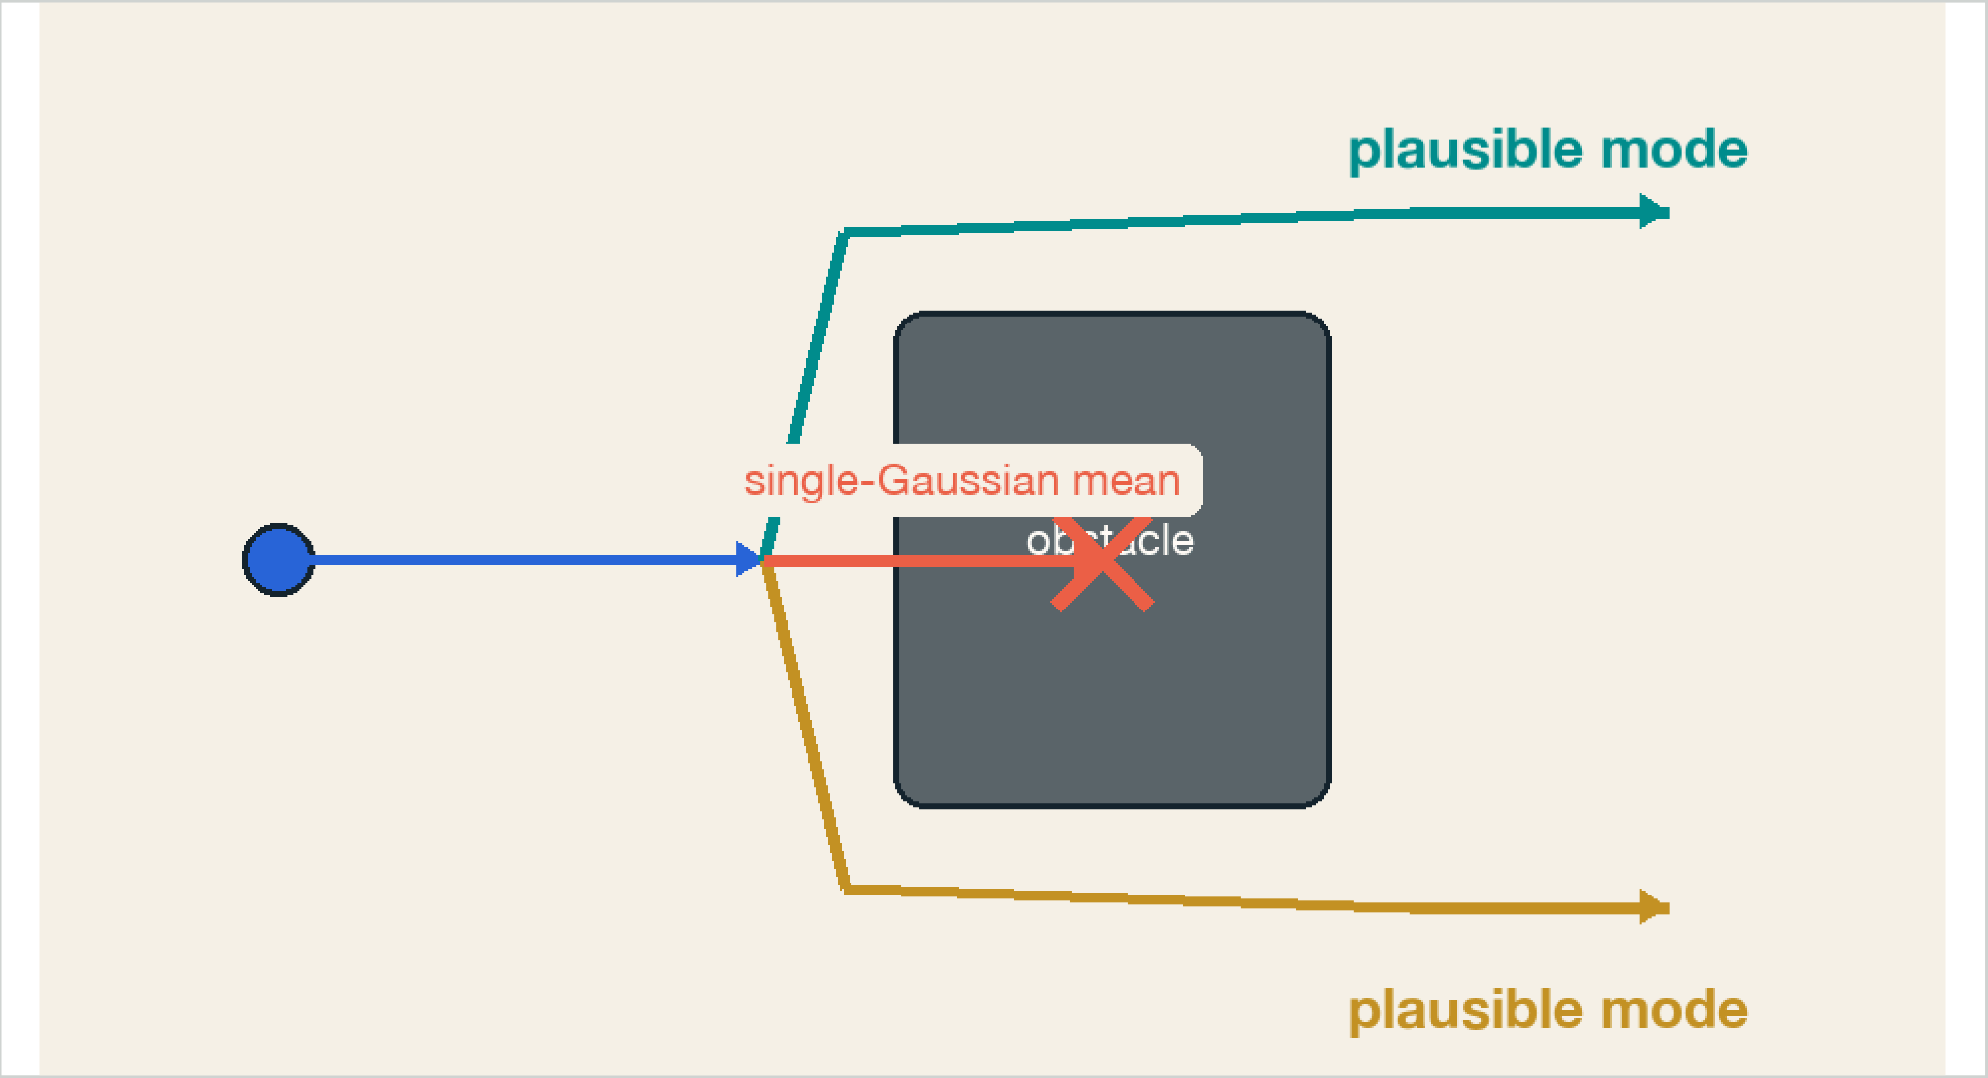

Time-Varying LGSSMs Use Time-Indexed Parameters

5

$$s_{t+1} = A_t s_t + B_t a_t + \epsilon_t^s, \qquad o_t = C_t s_t + \epsilon_t^o$$

Known extension: A_t, B_t, C_t, Q_t, R_t still define a linear-Gaussian SSM with exact Kalman filtering.

Unknown drift or context must be represented in the state, modeled adaptively, or learned from additional data.

Explicit States and Recurrent Predictive Memories

Sequential models can carry predictive information in an explicit generative state or in memory updated from observation history.

EXPLICIT GENERATIVE STATE

Transition and observation laws are named

Simulator: physical state advanced by a step rule

LGSSM: continuous vector state

HMM: finite categorical state

RECURRENT PREDICTIVE MEMORY

History is compressed inside a predictor

RNN: hidden memory updates as new inputs arrive

Video model: learned latent memory supports frame prediction

This memory may be agent-side rather than a world-side latent variable.

SSM interpretation: initial law + Markov transition under actions + observation law conditioned on state.

SSM Requirements, LGSSM Choices, and Algorithms

SSM REQUIREMENT

Sufficient Markov state

Initial, transition, and observation laws define a trajectory model.

LGSSM CHOICES

Linear + Gaussian

Continuous state, linear maps, and Gaussian uncertainty; known parameters may vary with time.

EXACT INFERENCE

Kalman filtering

Prediction and conditioning remain Gaussian under the LGSSM assumptions.

LEARNING

Maximum likelihood

Direct gradients or EM estimate model parameters from trajectories.

Hierarchy: choose a valid state and probability laws first; then choose inference and learning algorithms that match those assumptions.

Self-Supervised Representation Learning

Learn useful representations from observations without state labels.

$$z_t = E_\phi(o_t;\, o_{<t}, a_{<t})$$

Today: infer the hidden state s_t under an LGSSM. Next: learn which observation features to preserve for prediction.

INVARIANCE

What should stay unchanged?

EQUIVARIANCE

What should transform predictably?

COLLAPSE

What keeps the representation informative?

> A contextualized representation z_t need not be Markov. A predictor may still need the sequence $z_{\le t}$.