

Introduction to Artificial Intelligence (INFO8006)

Exercise session 5

Inference tasks

For **unobserved** temporal variables X_t and related **observed** evidences E_t , we define the task of

- **prediction**: $P(X_{t+k} | e_{1:t})$ for $k > 0$,
- **filtering**: $P(X_t | e_{1:t})$,
- **smoothing**: $P(X_k | e_{1:t})$ for $0 \leq k < t$,
- **most likely explanation**: $\arg \max_{x_{1:t}} P(x_{1:t} | e_{1:t})$.

A **Bayes filter** is a dynamic filtering task where a **belief state** is updated through time as new **evidences** are collected. The latter can be decomposed in two recursive phases:

1. Prediction: using the transition model, update your belief about the unobserved state X_t .
2. Update: given new evidence, update your new state.

Kalman filter

The **Kalman filter** is a Bayes filter on **continuous variables** which assumes:

- Gaussian prior
- Linear Gaussian transition model
- Linear Gaussian sensor model

Under those assumptions, both the prediction and update steps leads to Gaussian distributions. Using Gaussian identities (cfr. previous session), we can derive the parameters μ_t and Σ_t of the **Kalman filter** $p(x_t | e_{1:t})$ in closed form.

Hidden Markov models (HMM)

A **HMM** is a **Markov process** in which state and evidence are **single discrete** random variables.

HMM allows to represent model's components as matrices and vectors, making the algorithms easy to frame as matrix/vector operations.

Formally, we define

- the prior belief $\mathbf{f}_0 = P(X_0)$,
- the **forward message** $\mathbf{f}_{1:t} = P(X_t | e_{1:t})$,
- the **backward message** $\mathbf{b}_{k+1:t} = P(e_{k+1:t} | X_k)$,
- the **transition matrix** \mathbf{T} s.t. $\mathbf{T}_{ij} = P(X_t = j | X_{t-1} = i)$,
- the **sensor matrix** \mathbf{B} s.t. $\mathbf{B}_{ij} = P(E_t = j | X_t = i)$.

Using those notations, we have

$$\begin{aligned}\mathbf{f}_{1:t+1} &= \alpha \mathbf{O}_{t+1} \mathbf{T}^T \mathbf{f}_{1:t} \\ \mathbf{b}_{k+1:t} &= \mathbf{O}_{k+1} \mathbf{b}_{k+2:t}\end{aligned}$$

with $\mathbf{b}_{t+1:t}$ an all-one vector, α a normalizing constant and \mathbf{O}_t a diagonal matrix whose elements are the column elements of \mathbf{B} corresponding to the realization of the evidence e_t .

Stationary distribution

A distribution is called a **stationary distribution** if, for $t \rightarrow \infty$,

$$P(X_t) = P(X_{t+1}).$$

Specifically, for **HMM**, the **stationary distribution** \mathbf{f} is a distribution such that

$$\mathbf{f} = \mathbf{T}^T \mathbf{f}.$$

In session exercises: Ex. 1, Ex. 2

Exercise 1 Oral exam (January 2023)

You are waiting for an oral exam for which the professor is known to be moody. Some students have difficult questions while others have easy ones. You learned from your seniors that if a student has difficult questions, there is a 40% chance that the next student will also have difficult ones. If a student has easy questions, there is 70% chance that the next student will also have easy ones. You observe the students that come out of the professor's office. They either look happy, neutral or sad. You assume that, if their questions were difficult, there are 20% and 60% chance that they look respectively happy and sad. If their questions were easy, there are 50% and 30% chance that they look respectively happy and sad.

1. Define the components of a hidden Markov model (HMM), where the state $X_t \in \{\text{difficult}, \text{easy}\}$ is the difficulty of the questions for the t -th student and the evidence $E_t \in \{\text{happy}, \text{neutral}, \text{sad}\}$ is the emotion you observe.
2. At some point, the professor decides that one in two students will come out of their office by another exit, which prevents you from observing them. Sketch a Bayesian network describing this new process.
3. Express, in terms of the HMM components, the distribution $P(X_t | e_{1:t-2})$ of the difficulty for t -th student given your observations $e_{1:t-2}$ of half of the previous students. Note that, $e_{1:6:2} = (e_1, e_3, e_5)$ and $e_{1:7:2} = (e_1, e_3, e_5, e_7)$. Separate the cases where t is even and odd.
4. Given the observations $e_{1:3:2} = (e_1, e_3) = (\text{sad}, \text{happy})$, compute the probability that the 3rd student has difficult questions.

Exercise 2 Umbrella World (AIMA, Section 15.1.1)

You are a security guard stationed at a secret underground installation. You want to know whether it is raining today, but your only access to the outside world occurs each morning when you see the director coming in with, or without, an umbrella. For each day t , the evidence is a single variable $Umbrella_t \in \{1, 0\}$, *i.e.* whether the umbrella appears or not, and the (hidden) state is a single variable $Rain_t \in \{1, 0\}$, *i.e.* whether it is raining or not.

You believe that from one day $t - 1$ to the next t , the chances that the weather stays the same are 70%. You also believe that the director brings his umbrella 90% of the time when it is raining, and 20% of the time otherwise.

1. You would like to represent your umbrella world as a Markov model. What formal assumptions correspond to your beliefs?
2. Sketch a Bayesian network structure describing the umbrella world and provide the transition and sensor models.
3. Express the distributions $P(R_{t+1}|R_{t-1})$, $P(U_t|R_{t-1})$ and $P(R_t|R_{t-1}, U_t)$ in terms of the transition and sensor models.
4. Suppose you observe an unending sequence of days on which the umbrella appears. Show that, as the days go by, the probability of rain on the current day tends monotonically towards a fixed point. Calculate this fixed point.
5. Now consider forecasting further and further into the future, given t umbrella observations. Is there a fixed point? If yes, compute its exact value.

Exercise 3 The coins

You are in a room containing 3 precious biased gold coins a , b and c . You inspect the coins and notice that the coins a , b and c have a head probability of 80 %, 50 % and 20 %, respectively.

Another person enters the room, takes the coins and put them into a bag. They draw a coin from the bag and tell you that they will repeat 4 times the same routine: hide their hand in the bag and either keep the current coin with probability $\frac{2}{3}$ or replace it by another, then toss it and show you the result. They proceed and the sequence of results are heads, heads, tail, heads. If you answer right to the following questions they will give you the coins.

1. Provide a hidden Markov model (HMM) that describes the process.
2. What are the probabilities of the last coin given the sequence of evidences?
3. What are the probabilities of the first coin given the sequence of evidences? And of the first coin tossed?
4. What is the most likely sequence of tossed coins?

Exercise 4 Super Spring Ultra Pro Max XXL

The “Foire de Liège” has a new attraction called “Super Spring Ultra Pro Max XXL” which consists in a ball attached to a spring on a platform. The participants take seat in the ball locked at some position. The ball is then released and pulled by the spring which makes it oscillate back and forth. After a few seconds, the ball is stopped by magnetic brakes. You notice that the final position of the ball is different each time. As the ball is opaque, you wonder if it is possible for the participants to guess where the ball stopped, given their perception of acceleration.

You more or less remember your Newtonian mechanics class and model the movement of the ball as a series of transitions

$$p_t = p_{t-1} + \Delta t \dot{p}_{t-1} + \frac{1}{2} \Delta t^2 \ddot{p}_{t-1}$$

$$\dot{p}_t = \dot{p}_{t-1} + \Delta t \ddot{p}_{t-1}$$

$$\ddot{p}_t = \rho w_{t-1} - \kappa p_{t-1} - \eta \dot{p}_{t-1}$$


$$w_t \sim \mathcal{N}(\alpha w_{t-1}, \sigma_w^2)$$

where p_t , \dot{p}_t and \ddot{p}_t are respectively the position, velocity and acceleration of the ball at timestep t , w_t is the wind at timestep t , Δt is the time elapsed between $t - 1$ and t , ρ is the thrust coefficient of the wind, κ is the stiffness of the spring, η is the friction coefficient of the platform and α is the persistence of the wind. You estimate that the ball starts 10.0 ± 0.5 m at the left of the spring with a negligible speed 0.0 ± 0.1 m s⁻¹ and acceleration 0.0 ± 0.1 m s⁻². You assume that the wind has a stationary distribution. Finally, you assume that the human perception of acceleration follows an unbiased Gaussian distribution.

1. You wish to predict the state of the ball given the perceptions of a participant. Define the components of a Kalman filter in this context.
2. Express the distribution $p(x_t | e_{1:t})$ with respect to the components defined previously.
3. Represent the transition and sensor models as a dynamic Bayesian network.

Exercise 5 Robots (UC Berkeley CS188, Fall 2017)

In the near future, autonomous robots will live among us. Therefore, the robots need to know how to act appropriately in the presence of humans. In this question, we explore a simplified model of this interaction. We assume that we can observe the robot's actions at time t , R_t , and an evidence observation, E_t , directly caused by the (hidden) human state, H_t . Robot actions from the current time-step affect the human state in the next-time step, as illustrated in the Bayesian network given below.



Assuming discrete random variables and given the probability tables $P(H_0)$, $P(E_t|H_t)$, $P(R_t|E_t)$ and $P(H_{t+1}|H_t, R_t)$, your goal is to derive a procedure to maintain a belief $P(H_t|e_{1:t}, r_{1:t})$ about the state of the human at time t .

1. Derive an update equation for incorporating a pair of observations (e_t, r_t) to a given belief state $P(H_t)$.
2. Derive an update equation for predicting the future state H_{t+1} of the human at time $t + 1$ given a belief state $P(H_t|e_{1:t}, r_{1:t})$.
3. Combine both equations to derive a recursive update equation of the belief state $P(H_t|e_{1:t}, r_{1:t})$, as observations are collected and time passes.
4. Let us now assume that all variables are continuous. Discuss how you would compute or approximate the belief state on a computer.

Exercise 6 September 2019 (AIMA, Ex 15.13 and 15.14)

A professor wants to know if students are getting enough sleep. Each day, the professor observes whether the students sleep in class, and whether they have red eyes. The professor has the following hypotheses:

- The prior probability of getting enough sleep, with no observations, is 0.7.
- The probability of getting enough sleep on night t is 0.8 given that the student got enough sleep the previous night, and 0.3 otherwise.
- The probability of having red eyes is 0.2 if the student got enough sleep, and 0.7 otherwise.
- The probability of sleeping in class is 0.1 if the student got enough sleep, and 0.3 otherwise.

The professor asks you to answer the following questions:

1. Formulate the environment and hypotheses as a dynamic Bayesian network that the professor could use to detect sleep deprived students, from a sequence of observations. Provide the associated probability tables.
2. Reformulate the dynamic Bayesian network as a hidden Markov model that has only a single observation variable. Give the complete probability tables for the model.
3. For the sequence $e_{1:3}$ of observations “no red eyes, not sleeping in class”, “red eyes, not sleeping in class” and “red eyes, sleeping in class”, calculate the distributions $P(EnoughSleep_t|e_{1:t})$ and $P(EnoughSleep_t|e_{1:3})$ for $t \in \{1, 2, 3\}$.

Quiz

The Kalman filter requires the specification of ...

- a prior $\mathbf{P}(\mathbf{X}_0)$ and a transition model $\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{x}_t)$.
- a transition model $\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{x}_t)$ and an observation model $\mathbf{P}(\mathbf{E}_t|\mathbf{x}_t)$.
- a prior $\mathbf{P}(\mathbf{X}_0)$, a transition model $P(\mathbf{X}_{t+1}|\mathbf{x}_t)$, and an observation model $P(\mathbf{E}_t|\mathbf{x}_t)$.
- a prior $\mathbf{P}(\mathbf{X}_0)$, a prior $\mathbf{P}(\mathbf{E}_0)$, a transition model $P(\mathbf{X}_{t+1}|\mathbf{x}_t)$, and an observation model $P(\mathbf{E}_t|\mathbf{x}_t)$.

In filtering, the belief state is updated at each timestep using ...

- the current belief and the transition model.
- the current belief, the observation model and the new evidence.
- the transition model, the observation model and the new evidence.
- the current belief, the transition model, the observation model and the new evidence.

Let us assume a Markov process with hidden state variables \mathbf{X}_t , evidence variables \mathbf{E}_t , transition model $\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{X}_t)$ and sensor model $\mathbf{P}(\mathbf{E}_t|\mathbf{X}_t)$. Which of the following is true?

- The Bayes filter estimates $\mathbf{P}(\mathbf{X}_{1:t}|\mathbf{e}_{1:t})$.
- Smoothing consists in computing $\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1})$ by pushing the belief state $\mathbf{P}(\mathbf{X}_t|\mathbf{e}_{1:t})$ forward through the transition model and then updating this new belief state with the evidence \mathbf{e}_{t+1} .
- The Kalman filter computes $\mathbf{P}(\mathbf{X}_{t+k}|\mathbf{e}_{1:t})$ by pushing the belief state $\mathbf{P}(\mathbf{X}_t|\mathbf{e}_{1:t})$ forward k times through the transition model.
- The Viterbi algorithm computes the most likely sequence of state values $\mathbf{x}_{1:t}$ given the evidence $\mathbf{e}_{1:t}$.