Lecture 7:
A Vacuum-Cleaning Robot

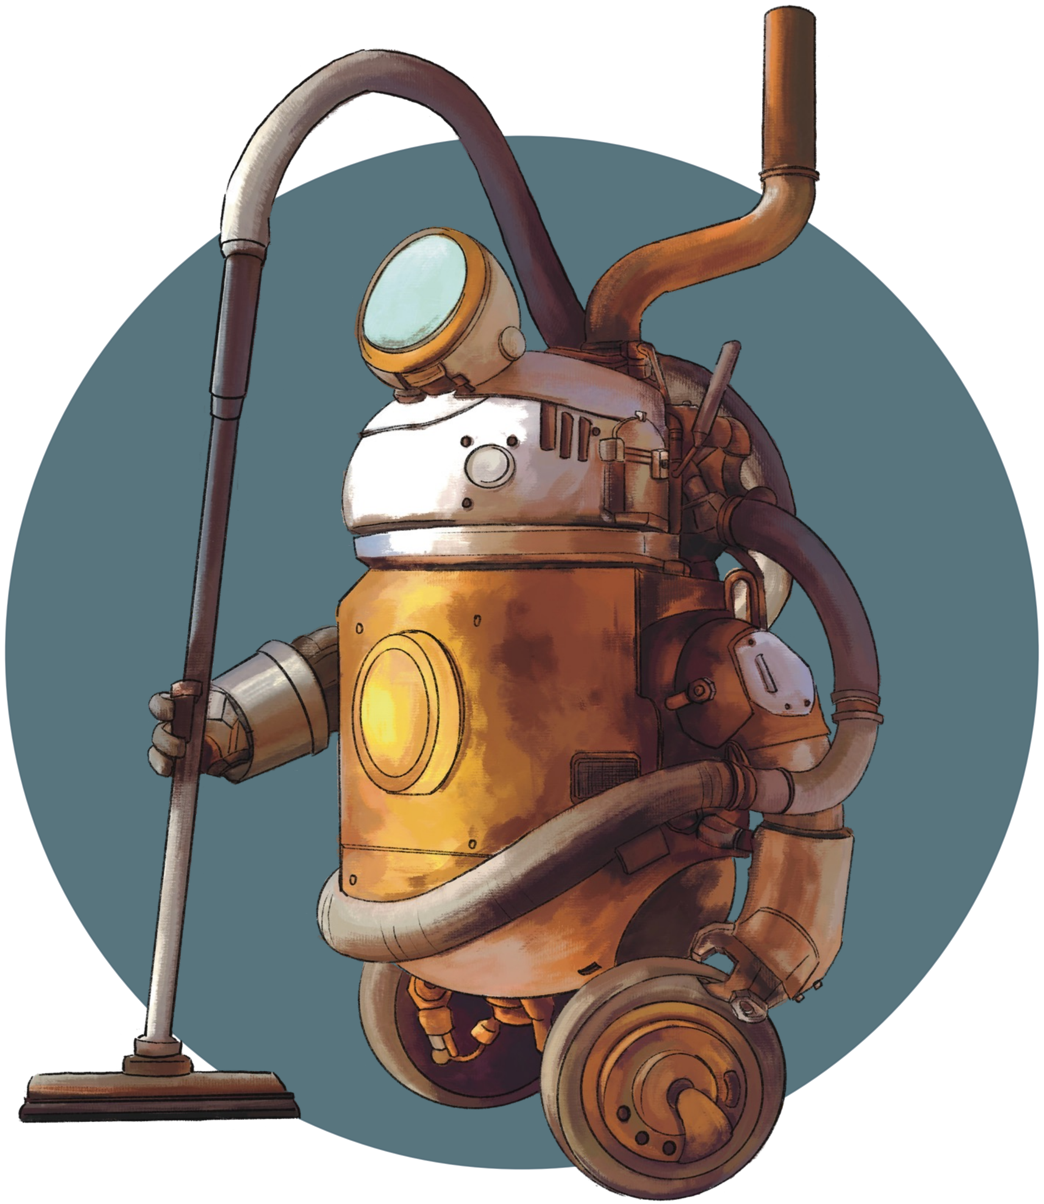

States, beliefs & actions

Frank Dellaert • Seth Hutchinson

Probability theory provides the mathematical foundation for dealing with uncertainties that confront robotic systems.

01 SENSE

Uncertain object, noisy sensors.

02 INFER

Bayes, MLE, MAP.

03 DECIDE

Choose a deterministic action.

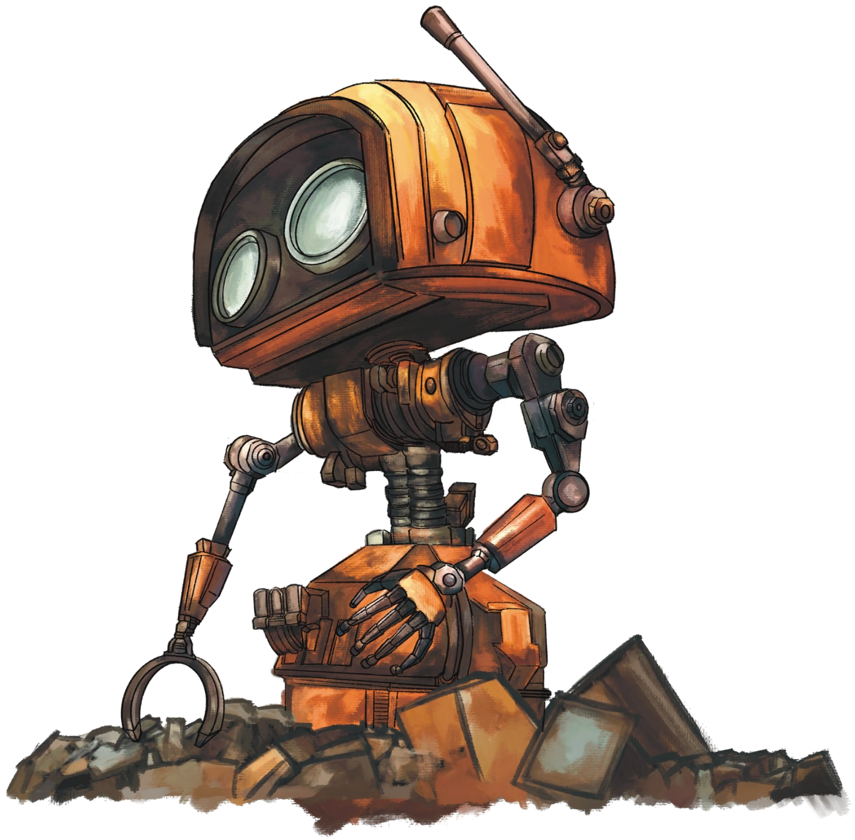

We can learn prior and sensor models from data we collect.

PRIORS + SENSOR MODELS

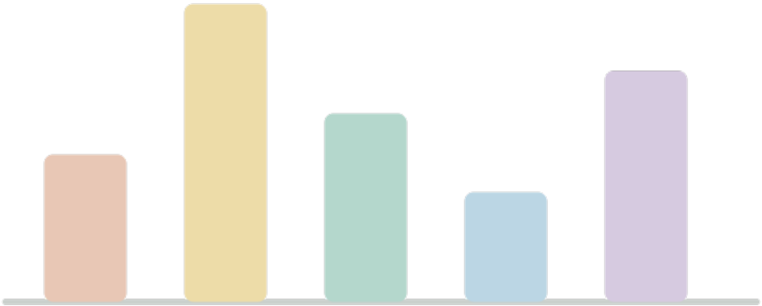

Count outcomes

then normalize

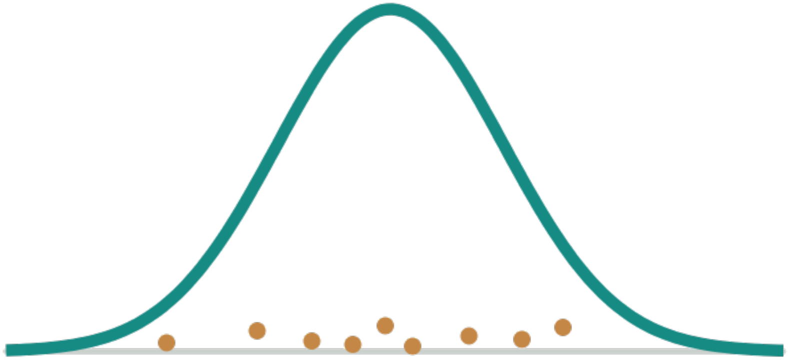

Fit a Gaussian

mean + variance

Our vacuum cleaner must reason about how the world changes over time.

Uncertain actions.
Consequences over time.

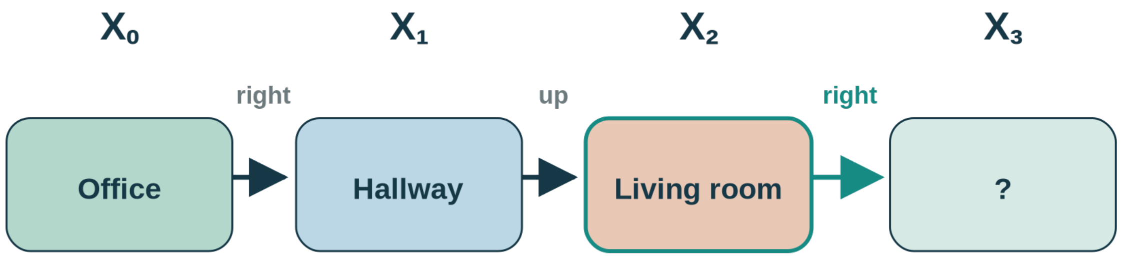

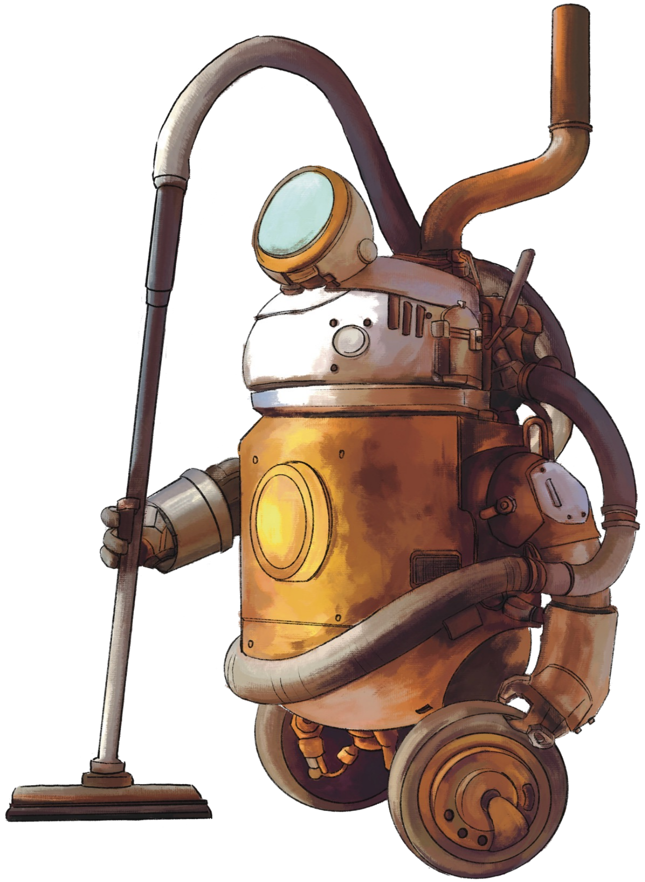

Low-level navigation and cleaning let us abstract away the geometry.

ASSUMED CAPABILITIES

- **Any-direction motion**
- Navigation
- Room coverage
- Collision avoidance

Each of the five rooms is one possible state of our robot.

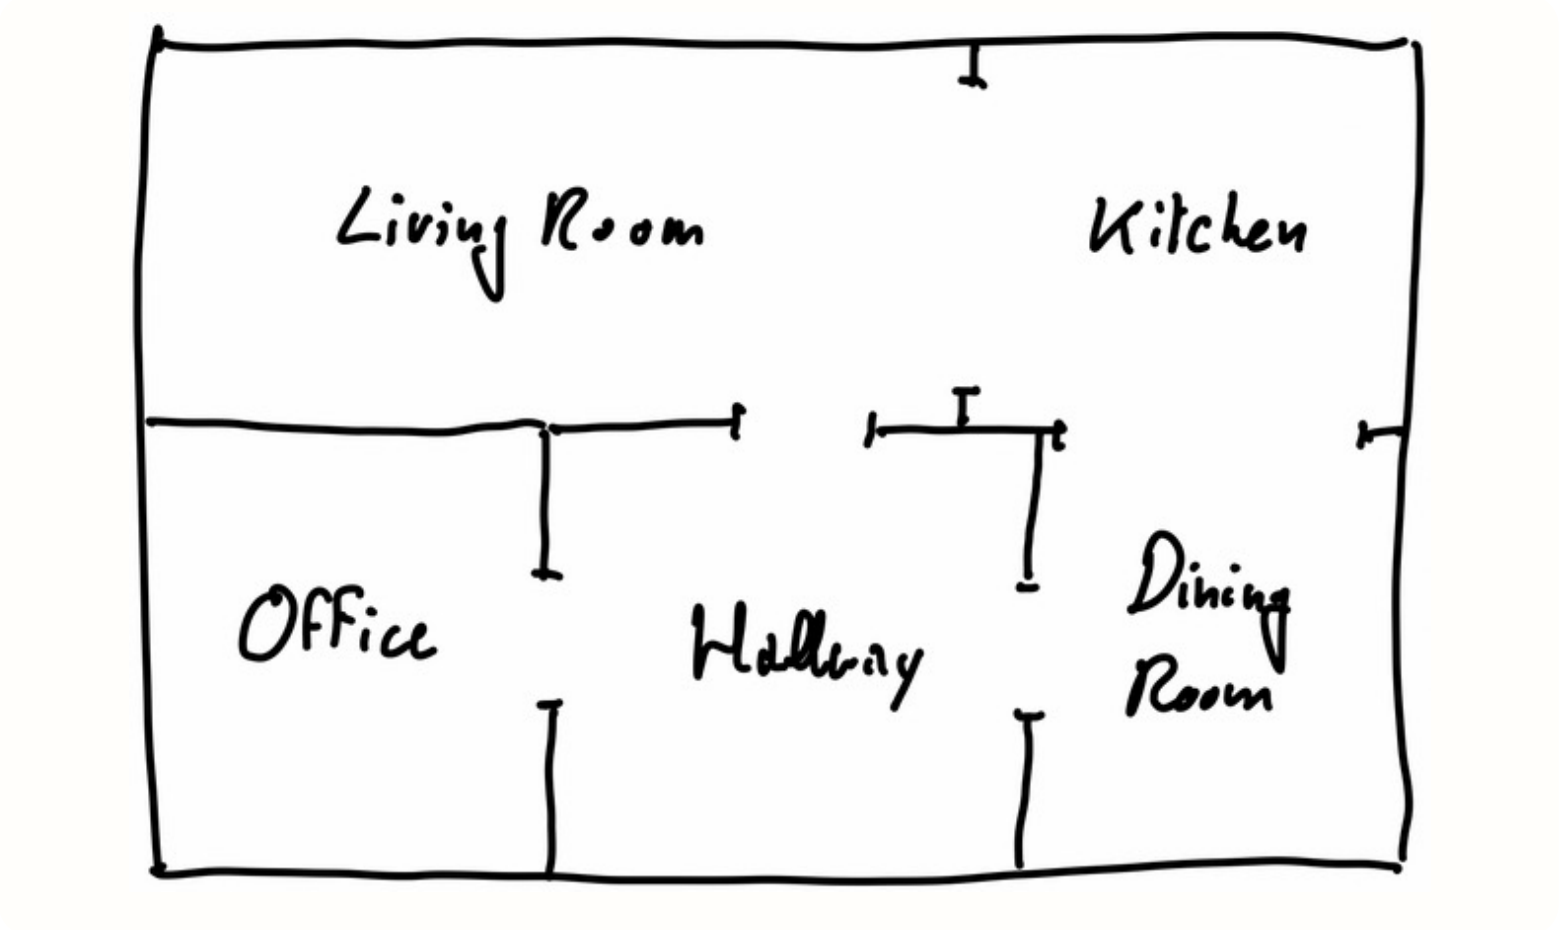

- **L** Living room
- **K** Kitchen
- **O** Office
- **H** Hallway
- **D** Dining room

The exact location within the living room is not relevant for this robot.

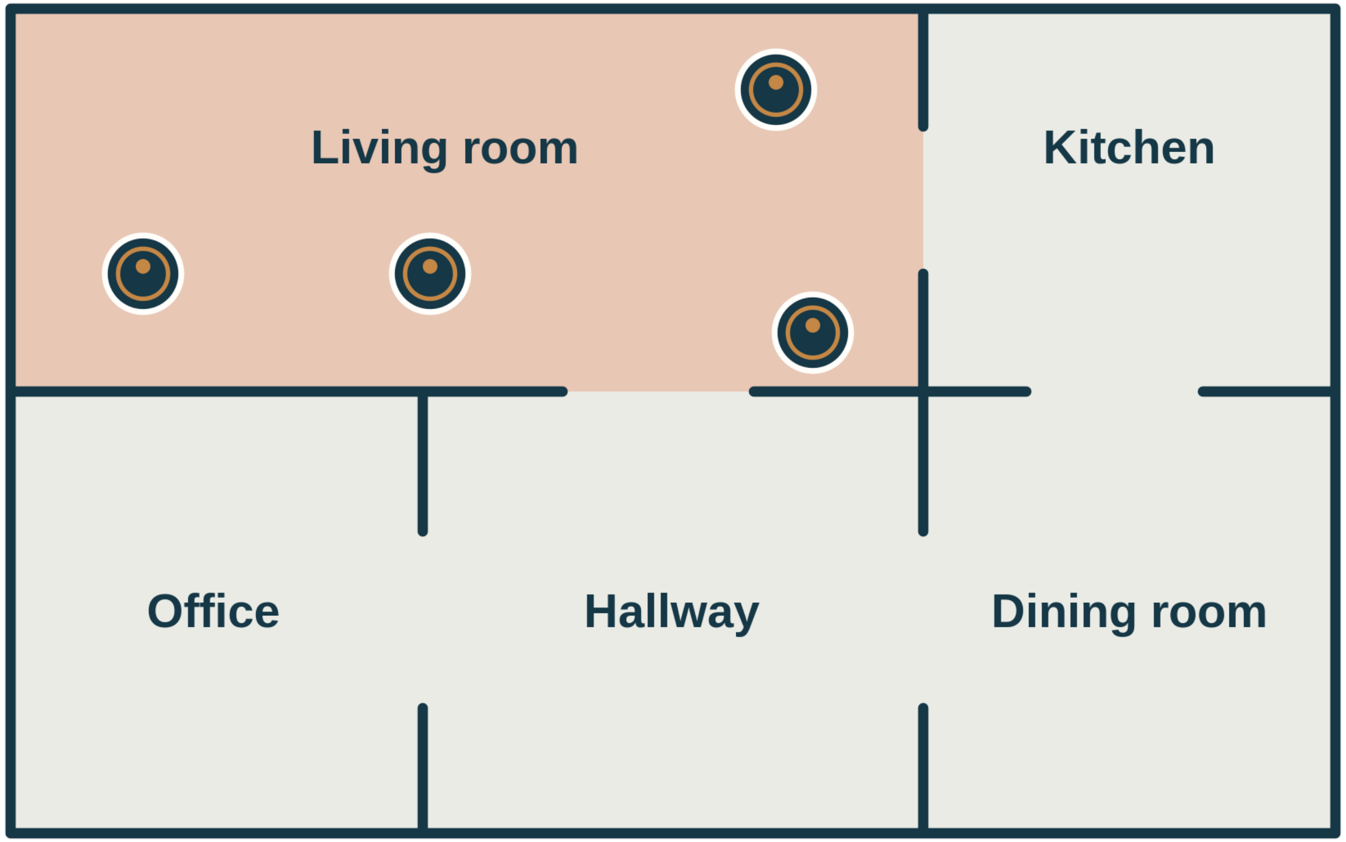

ALL FOUR LOCATIONS

$$X_t = \text{Living Room}$$

Exact position
is irrelevant here.

We represent the state space by an undirected graph.

PHYSICAL LAYOUT

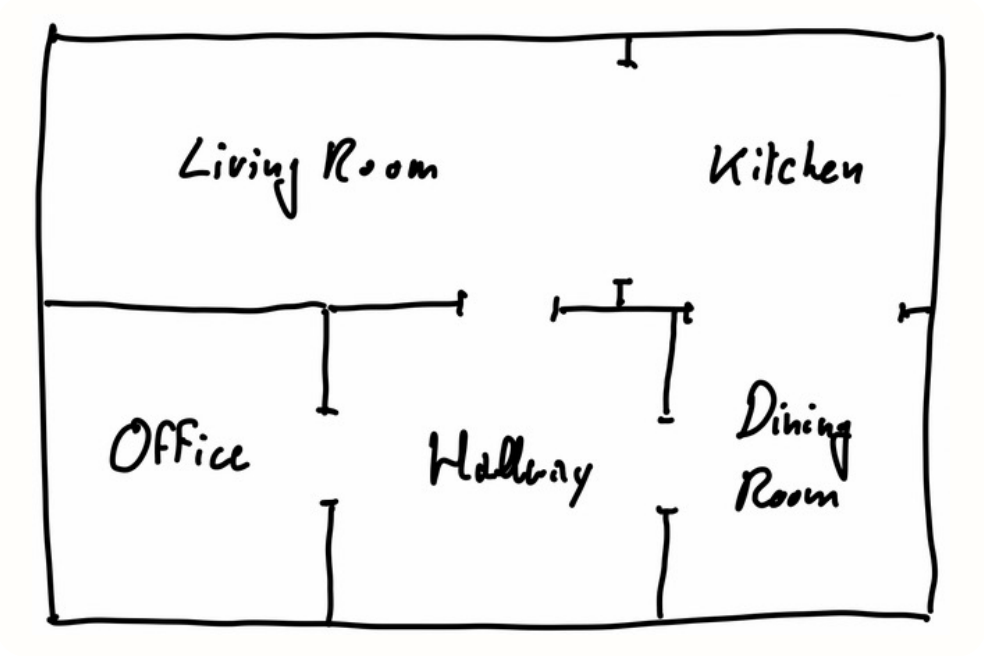

CONNECTIVITY GRAPH

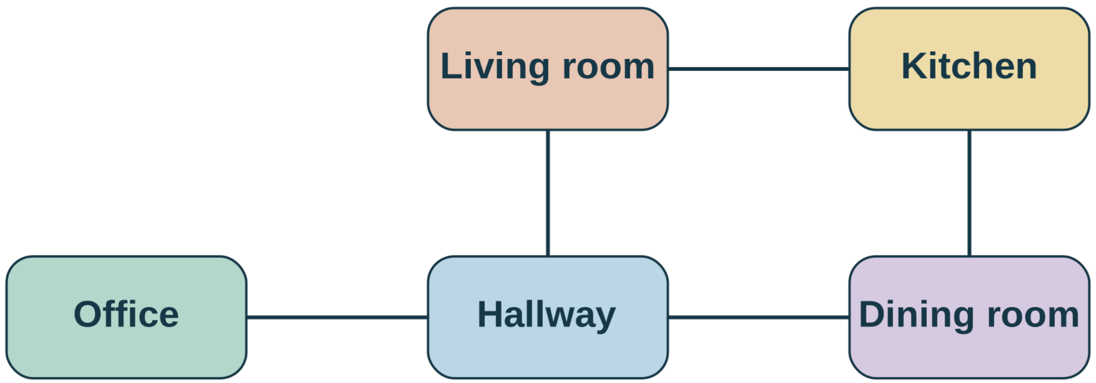

Rooms become nodes. Doorways become edges.

Teleportation for mobile robots has unfortunately not been invented yet!

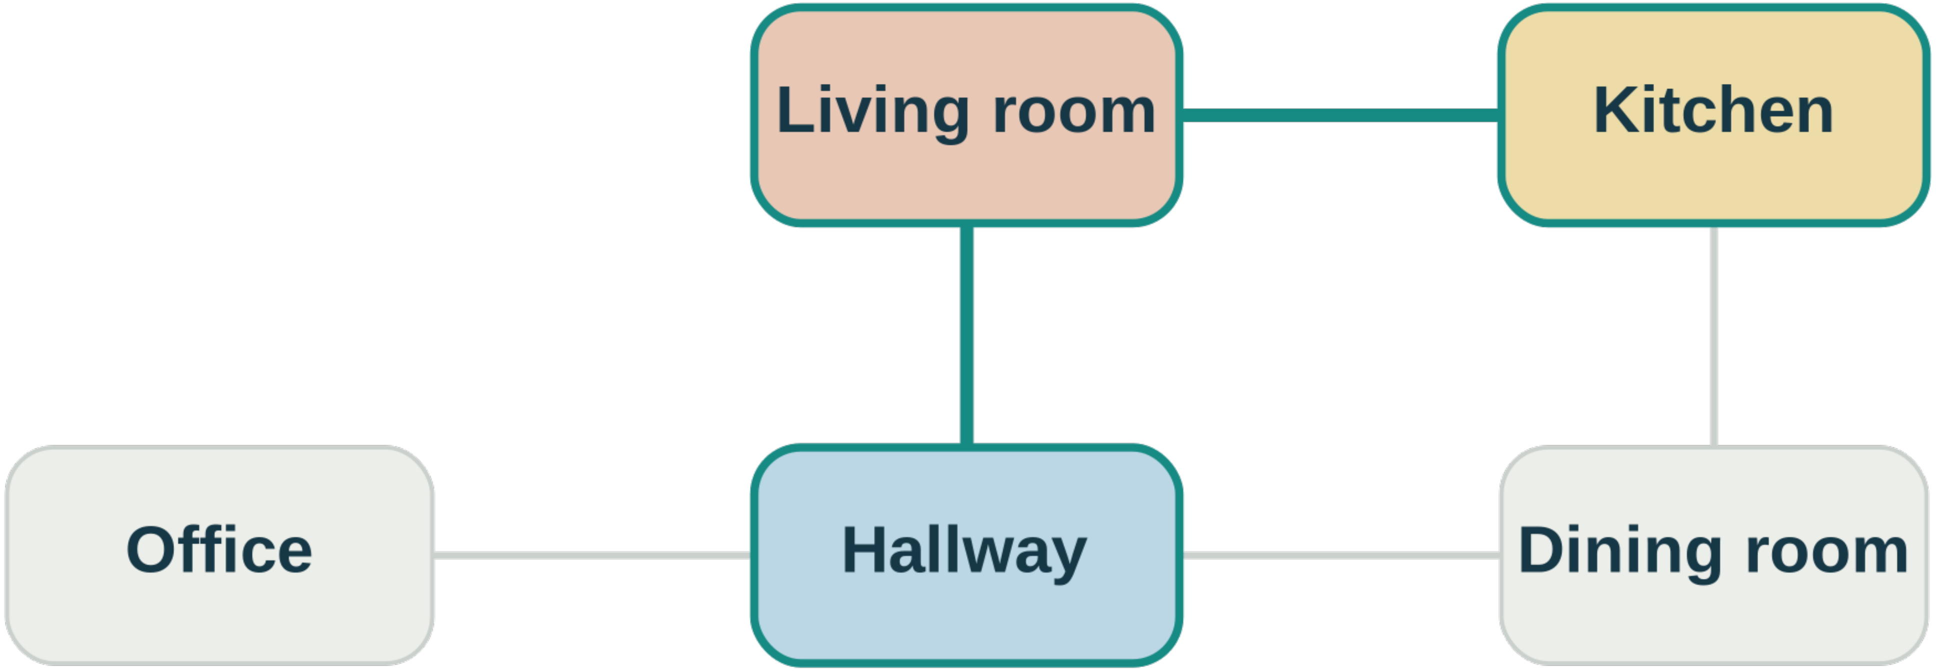

Living room → Kitchen: direct

Kitchen → Office: not direct

The robot starts in the office, so its initial state is known with certainty.

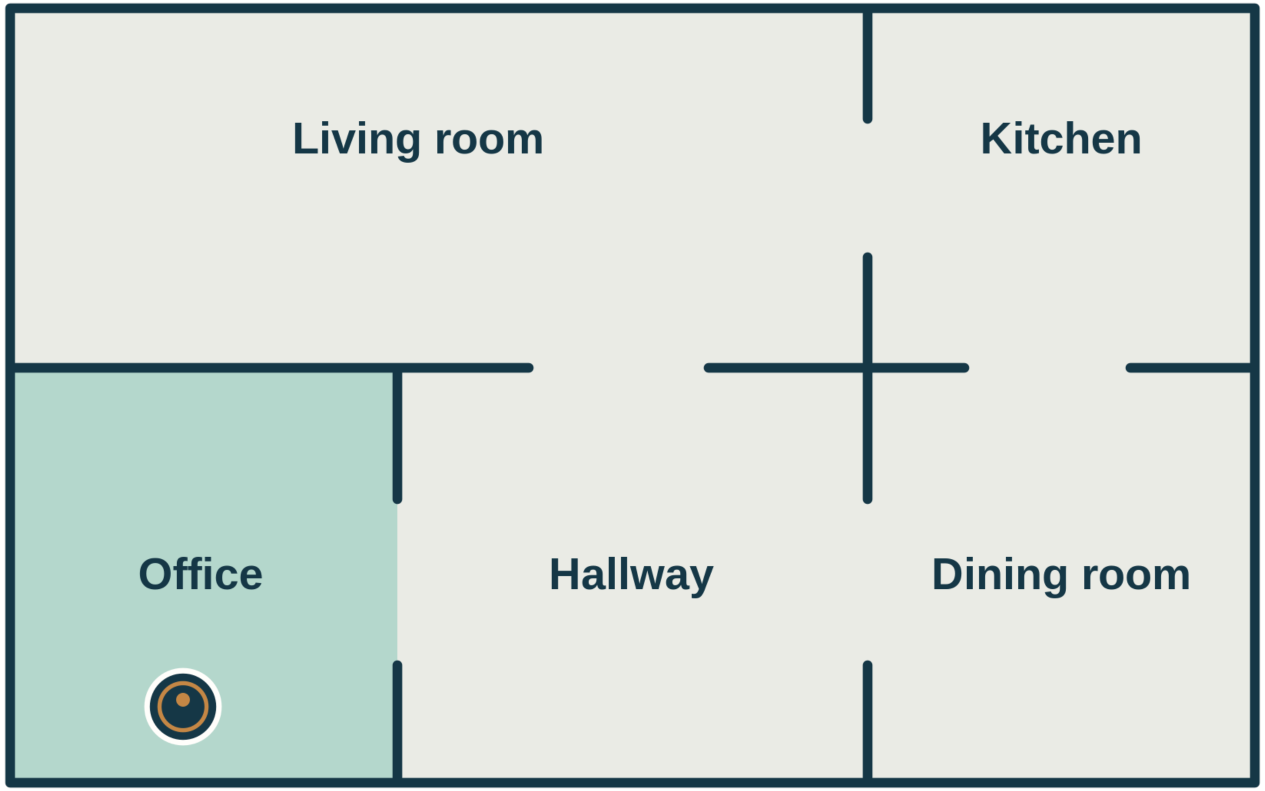

$$P(X_0 = O) = 1$$

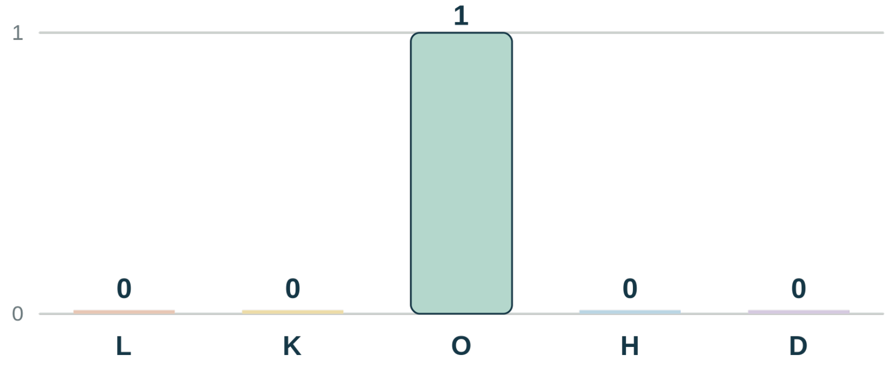

The charging station fixes the initial state.

We only track the moments when the robot acts or takes a measurement.

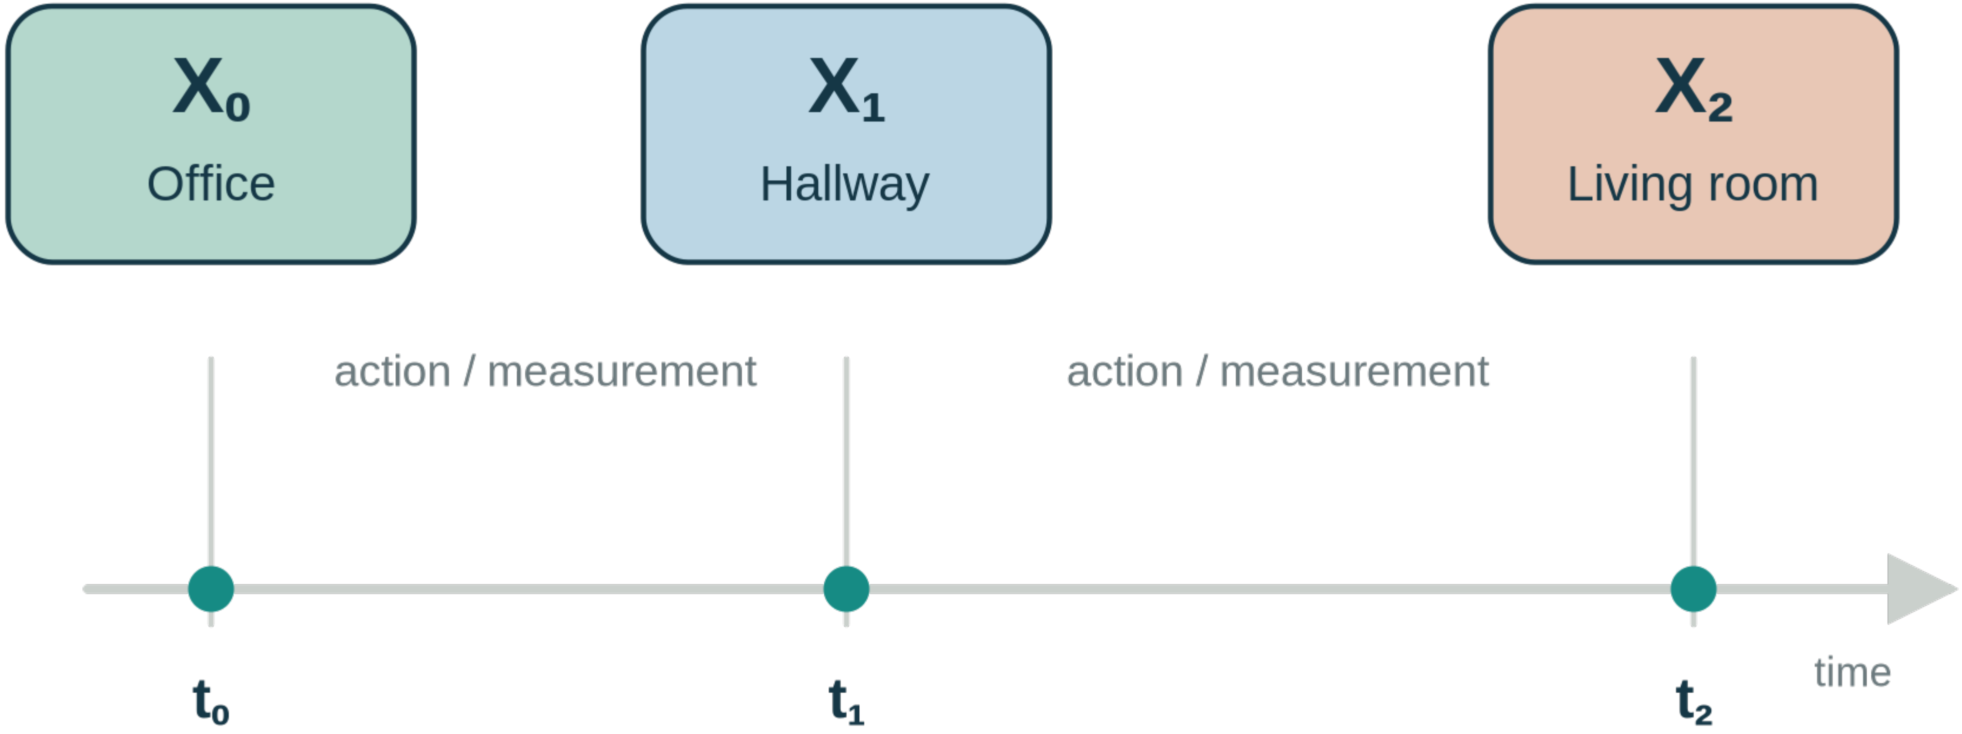

The belief state is a row vector whose elements correspond to the possible states.

$b_t(L)$	$b_t(K)$	$b_t(O)$	$b_t(H)$	$b_t(D)$
Living room	Kitchen	Office	Hallway	Dining room

$$b_t(x) = P(X_t = x \mid x_0, a_0, \ldots, a_{t-1})$$

$$\sum_x b_t(x) = 1$$

The belief state is a row vector.

We will characterize what we know and don’t know using probability distributions.

STATE X_t

one actual room

BELIEF B_t

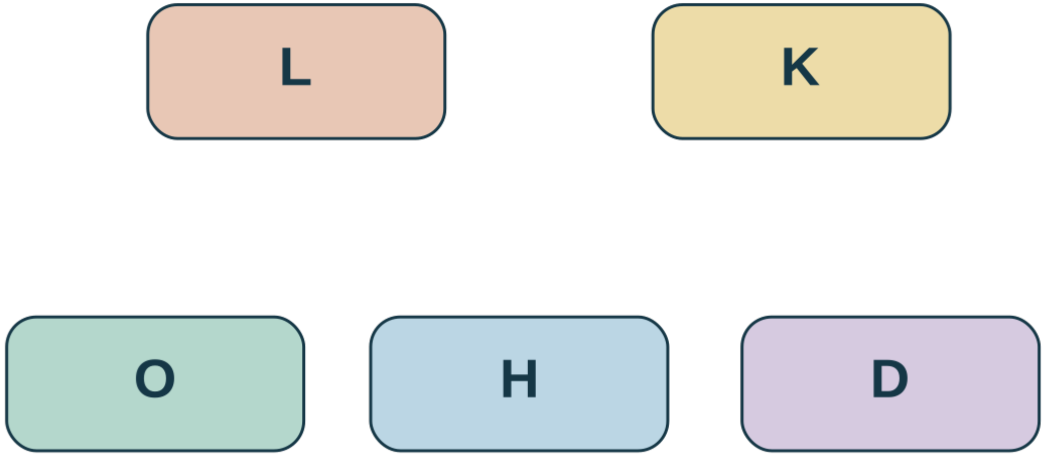

probability over rooms

Our robot is equipped with four primitive actions: Left, Right, Up, and Down.

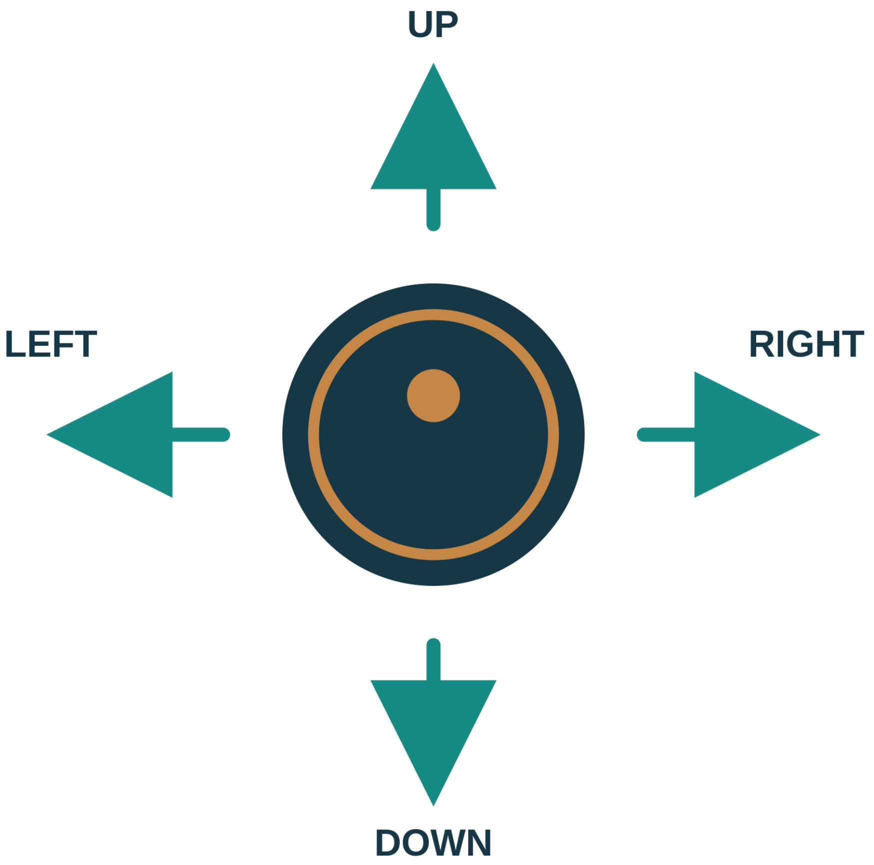

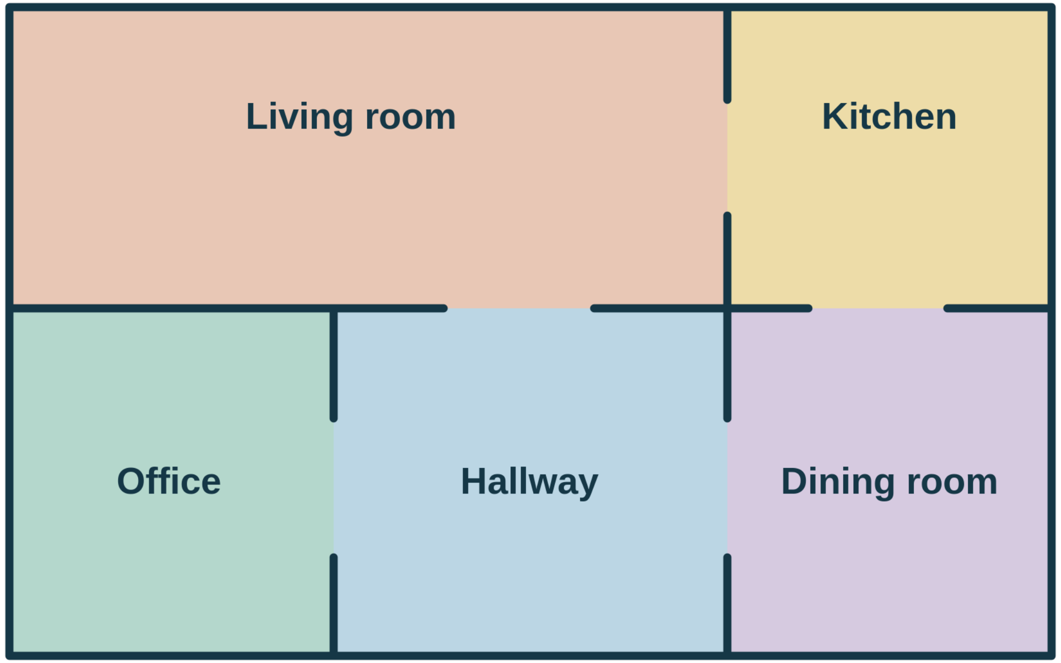

Directions are relative to the map.

The effects of actions depend on the current world state.

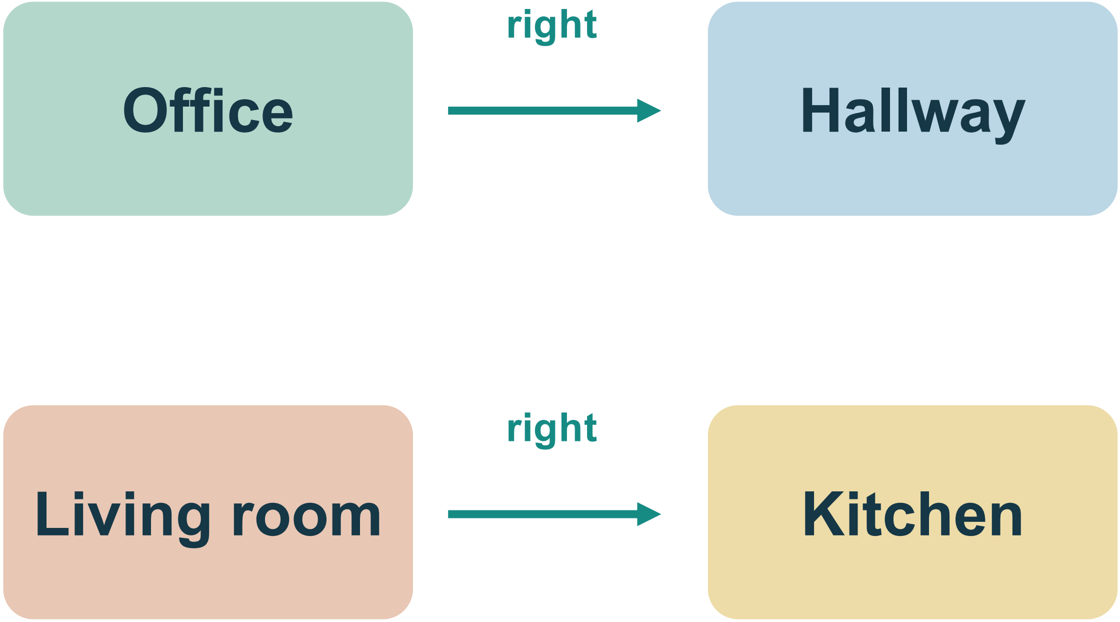

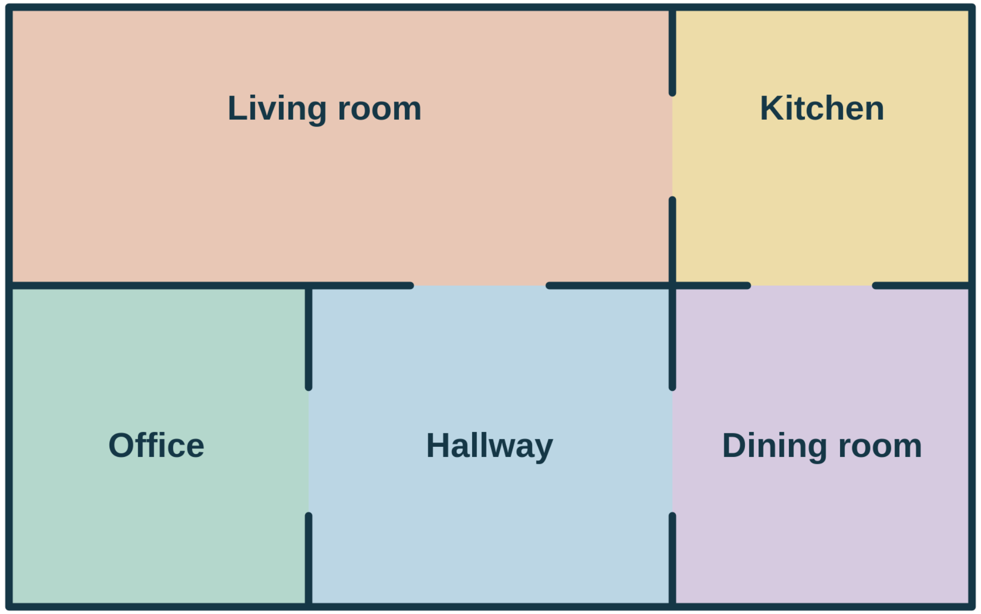

Nominal destinations

We turn room connections into directed edges labeled by actions.

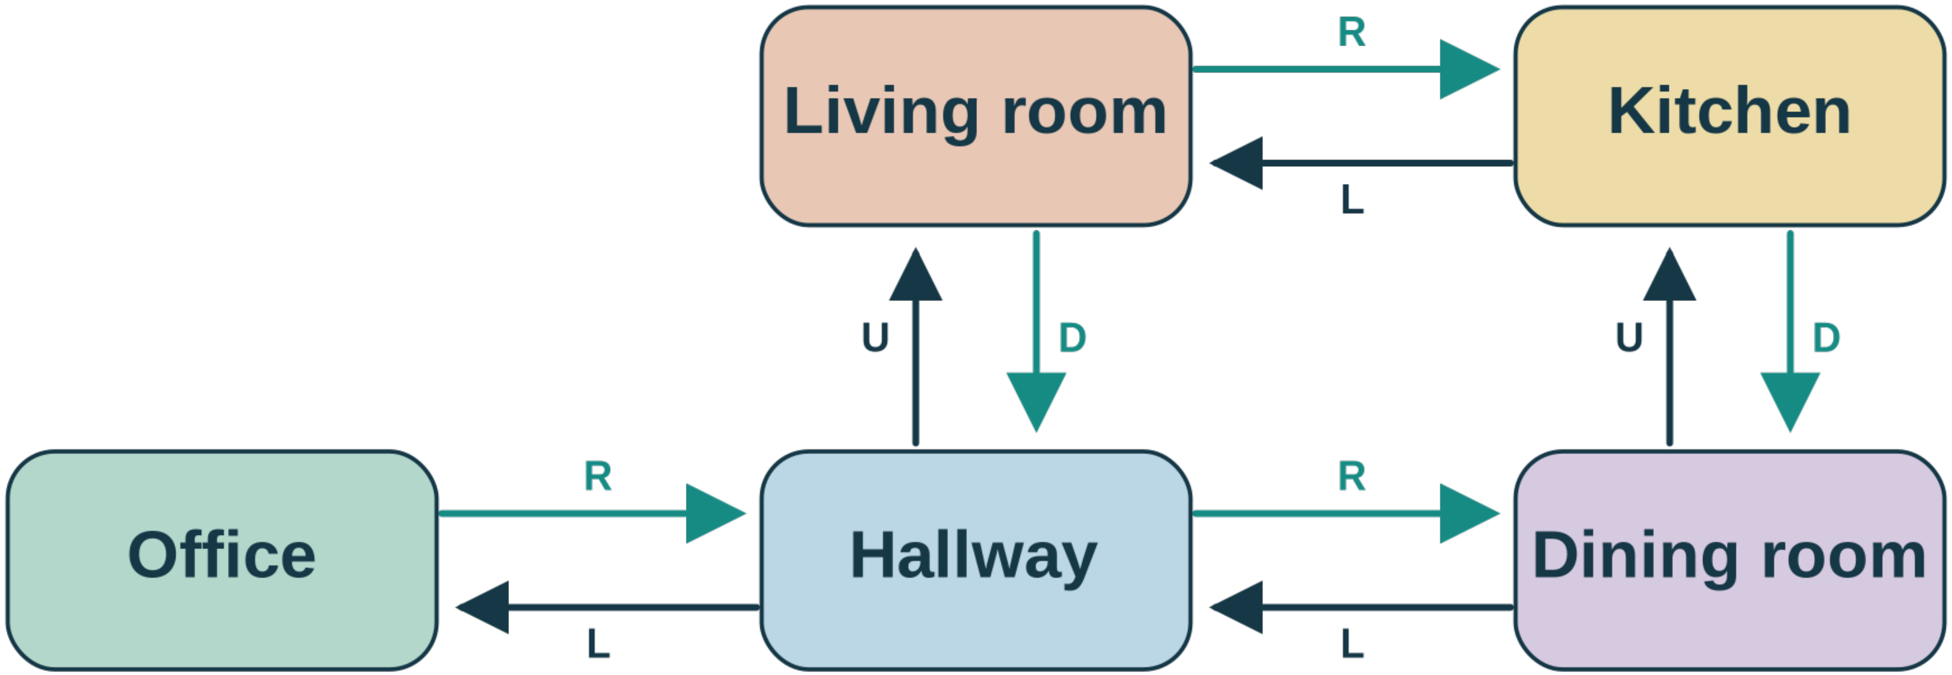

L / R / U / D = left / right / up / down

Robots do not always execute actions perfectly.

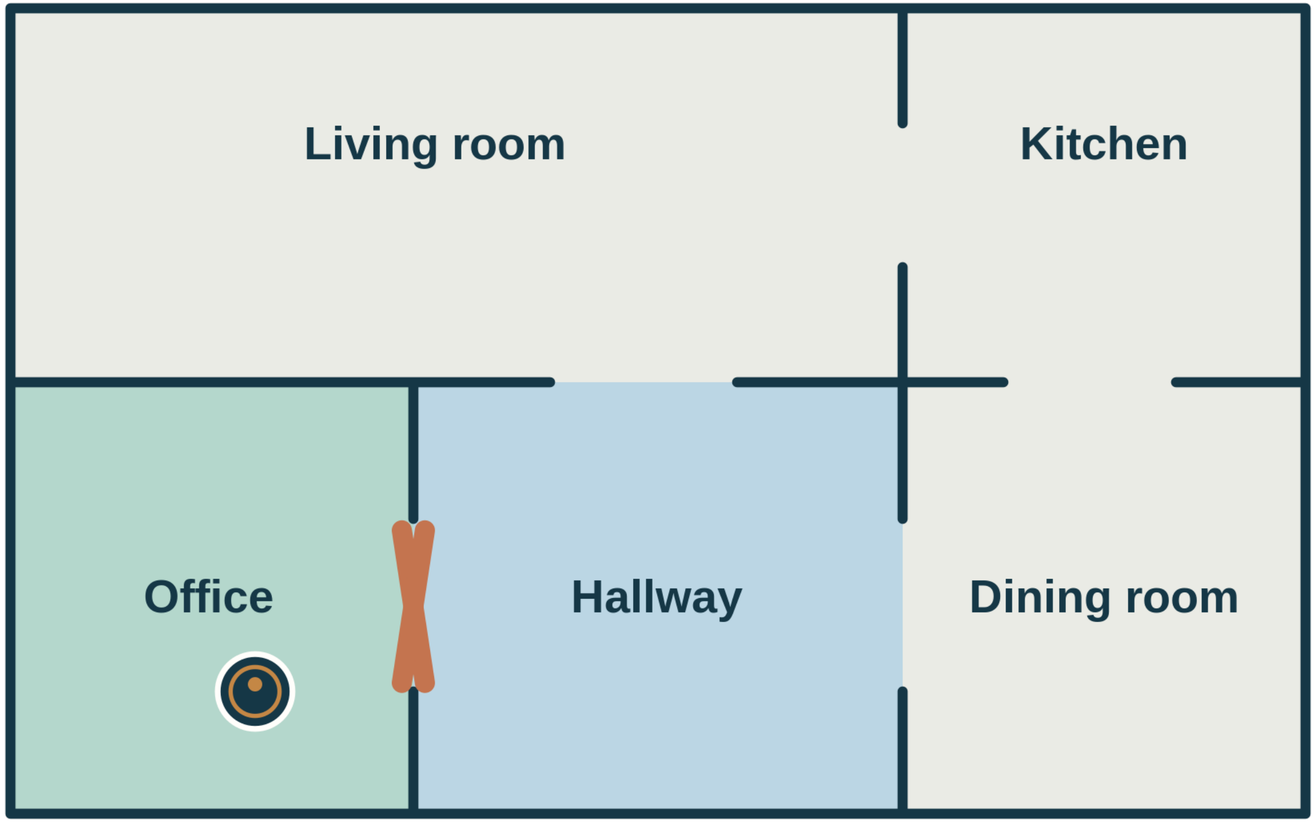

RIGHT ≠ ARRIVE

Blocked doorway.
Navigation failure.
Unavailable direction.

We will model uncertainty in the effects of actions by using conditional probability distributions.

$$P(X_{t+1} = x' \mid X_t = x,\ a_t = a)$$

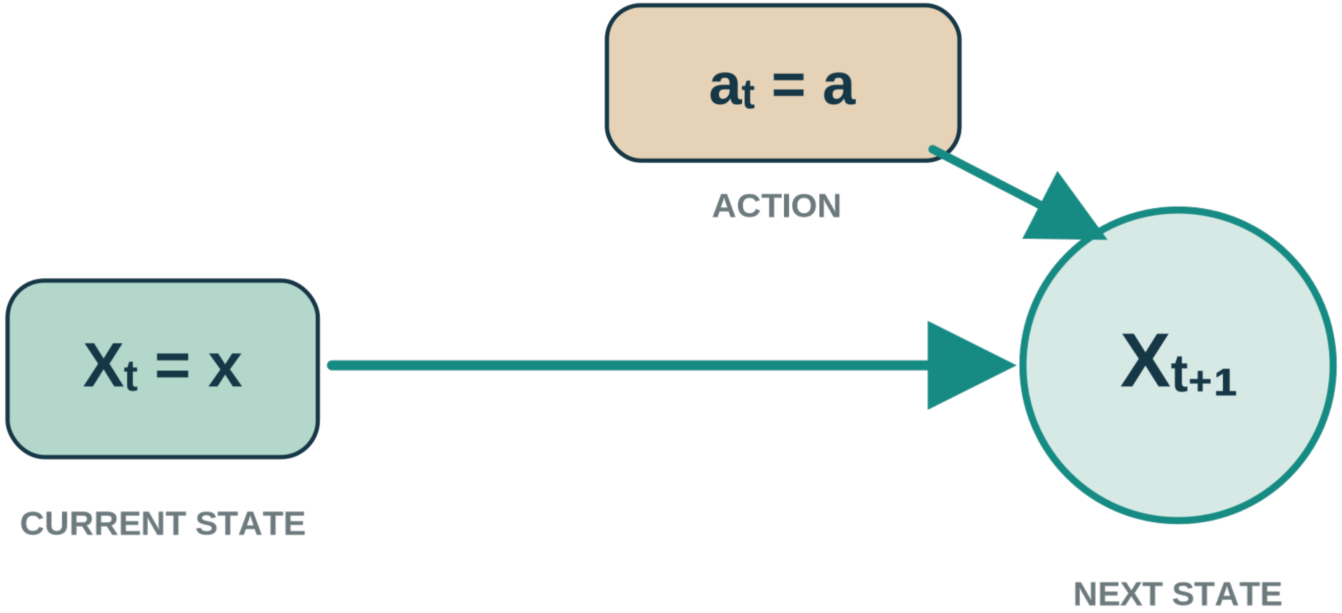

The current room and action determine a distribution over the next room.

$$P(X_{t+1} = x' \mid X_t = x,\ a_t = a)$$

X_{t+1}	next-state random variable
x'	a candidate room
$X_t = x, \quad a_t = a$	what we condition on

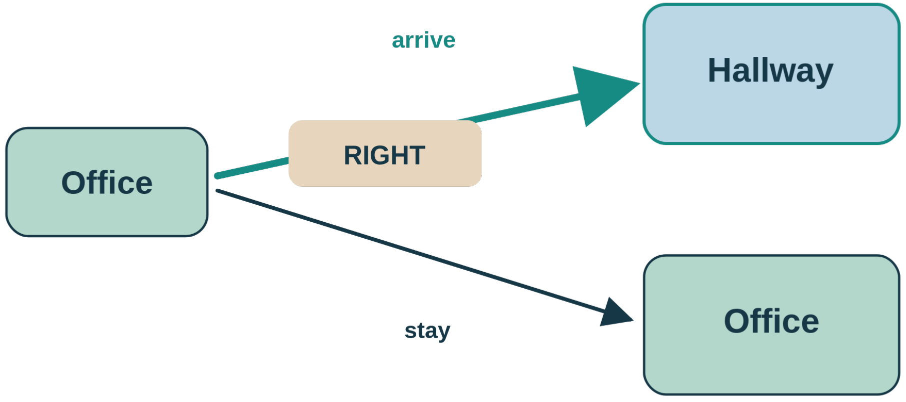

Moving right from the office reaches the hallway with 80% probability.

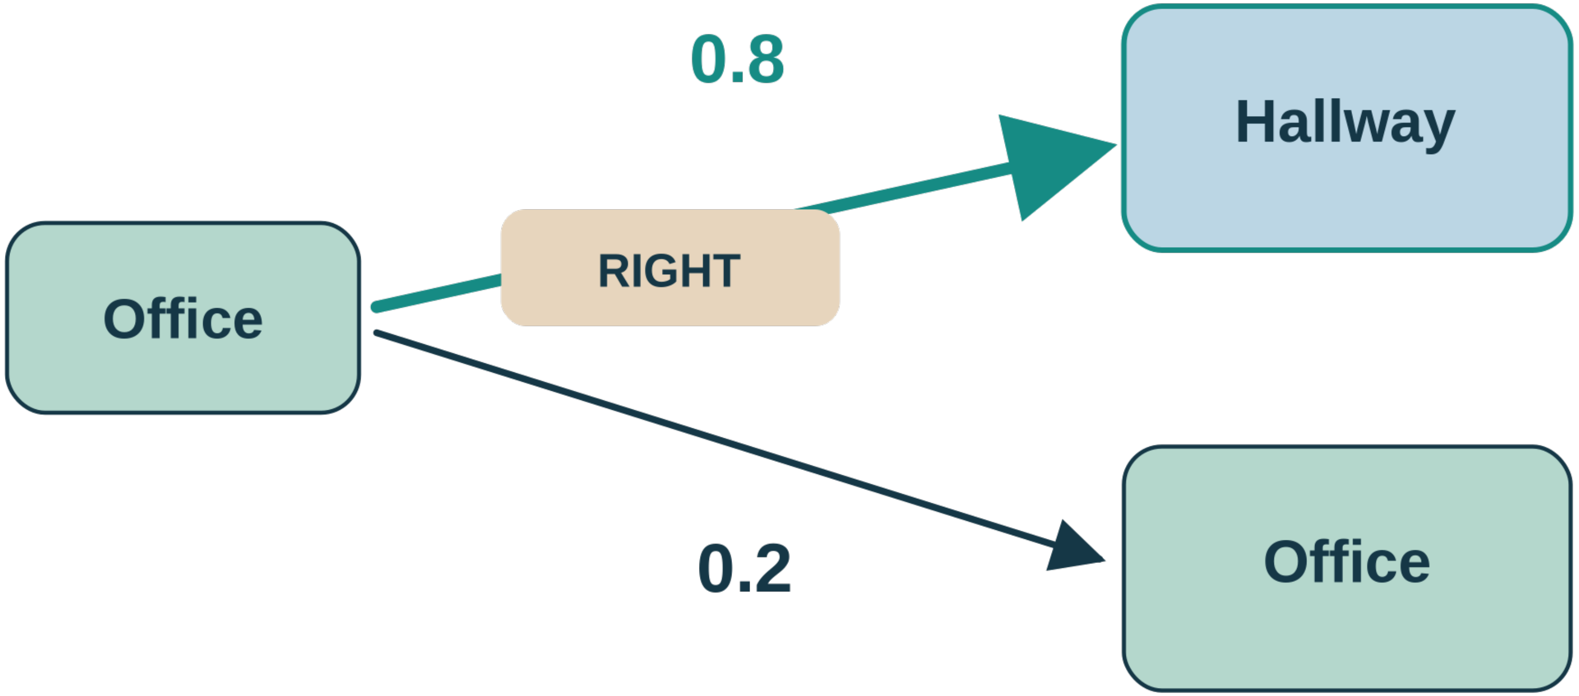

Book motion model: 80% arrival, 20% stay.

Our belief changes when we act, even before we take a measurement.

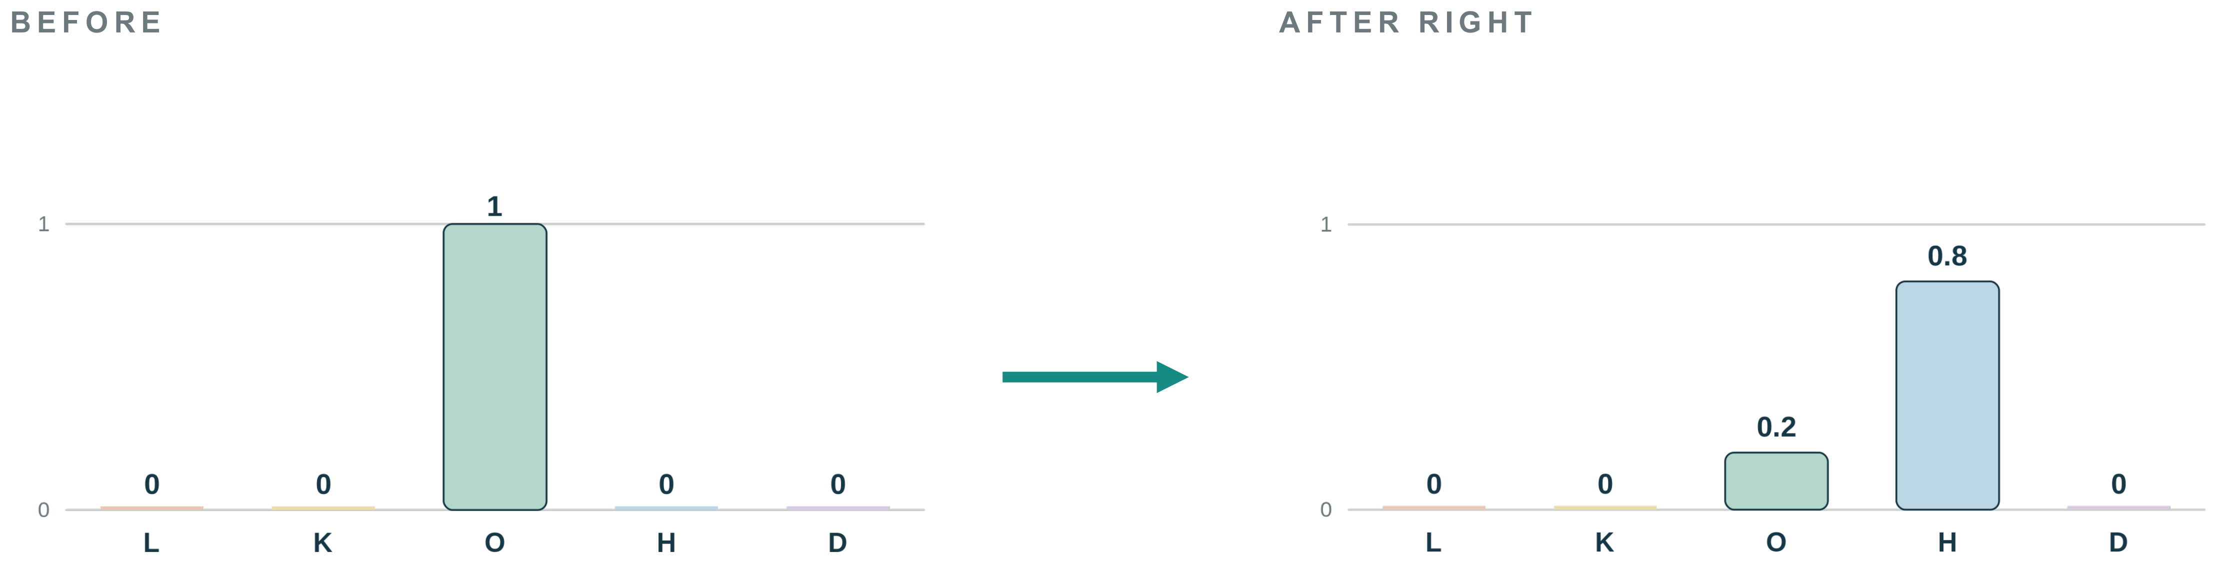

$$b_0 = [0 \quad 0 \quad 1 \quad 0 \quad 0]$$

$$b_1 = [0 \quad 0 \quad 0.2 \quad 0.8 \quad 0]$$

No sensor measurement yet.

Must we remember every previous state and action to predict what happens next?

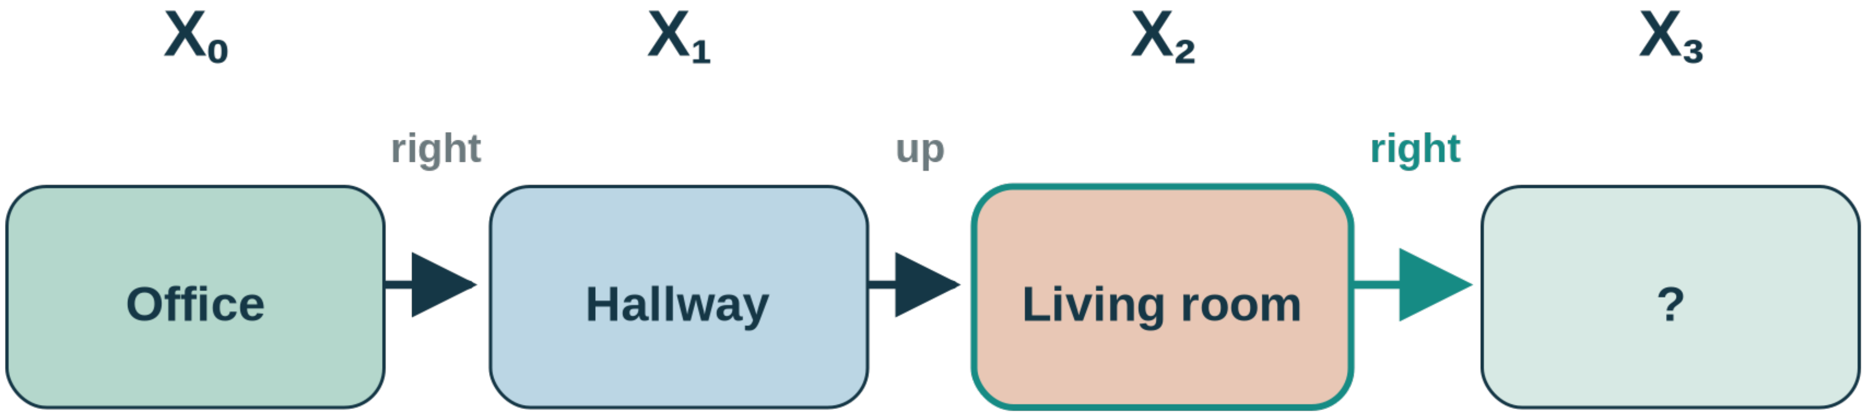

Predict X_3.

Once we know the current room and action, the route taken to get here is irrelevant.

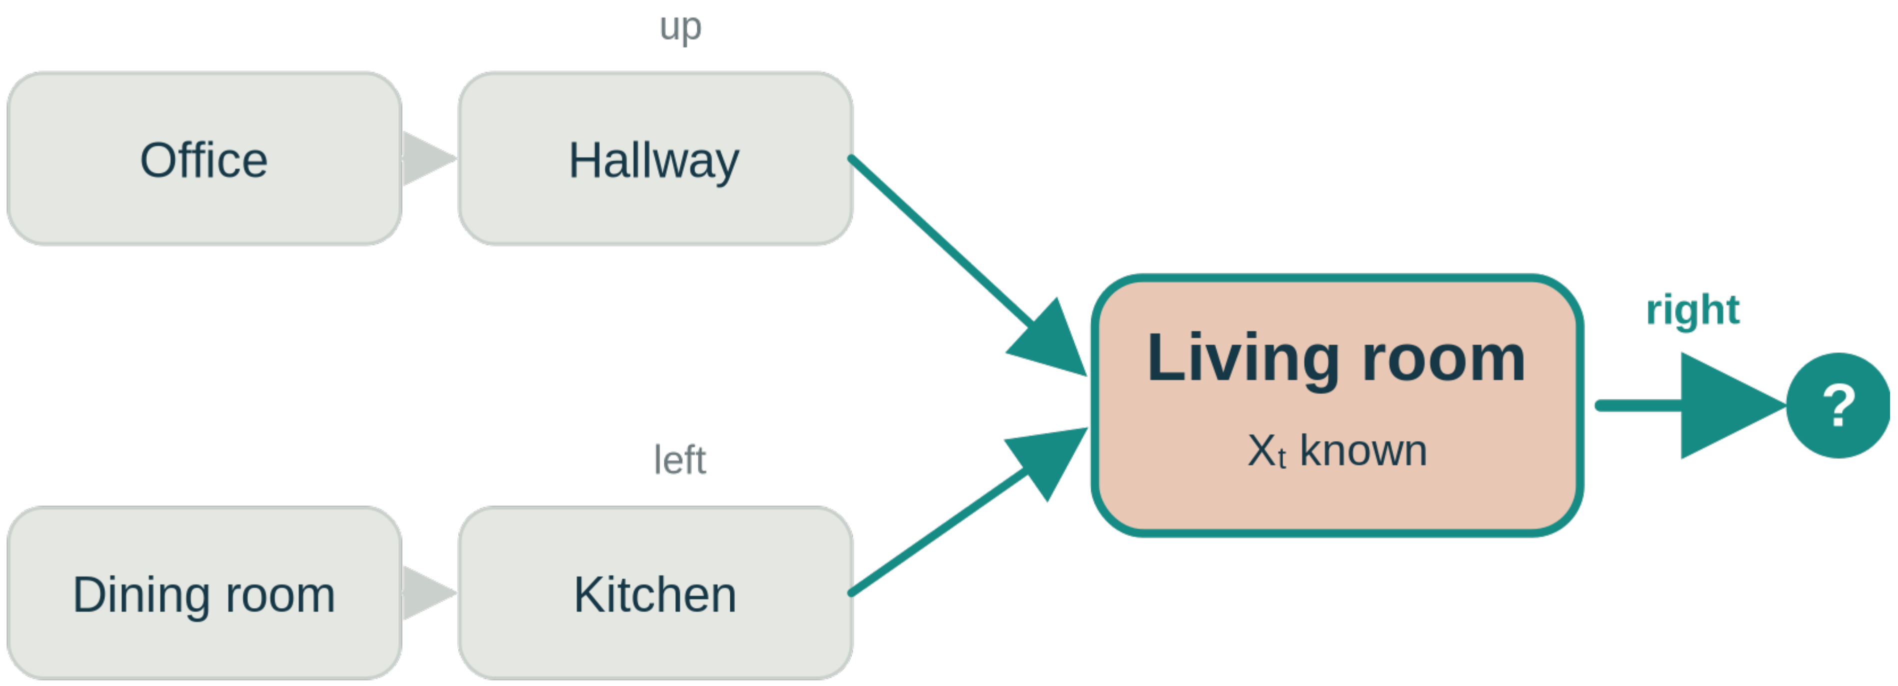

Given Living room + right, earlier routes are irrelevant.

Markov: The next state is independent of the past, given the current state and action.

$$P(X_{t+1} \mid X_{0:t}, a_{0:t}) = P(X_{t+1} \mid X_t, a_t)$$

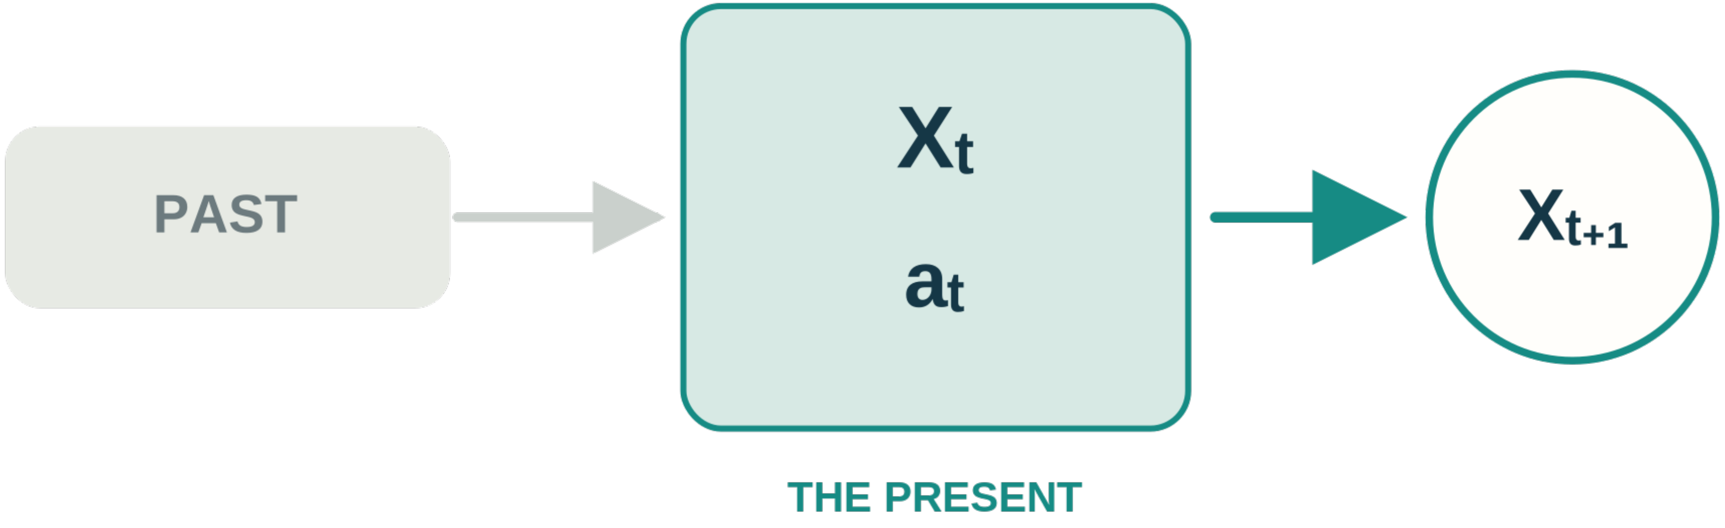

Condition on the current state and current action.

When the current state is unknown, history still informs our belief about it.

History informs the belief about the current room.

Rooms, uncertain actions, and the Markov property give us a model for reasoning over time.

01

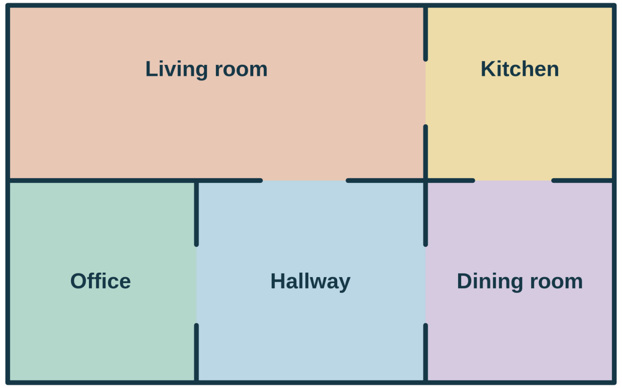

Abstract the state.

One room label.

02

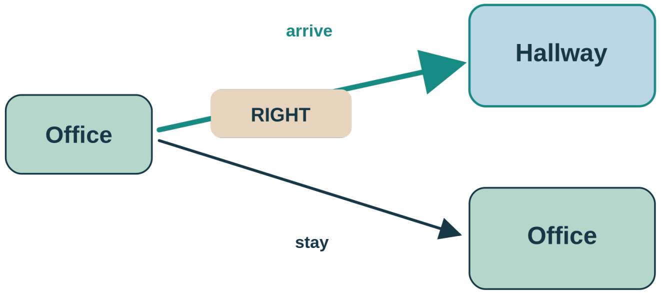

Model uncertain actions.

A distribution.

03

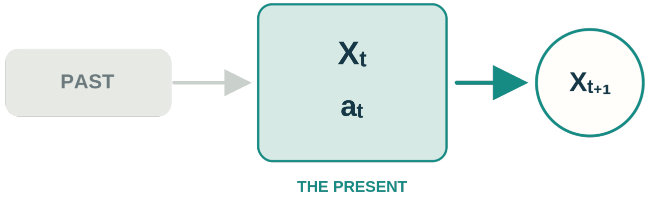

Use the present.

State + action predict next.

COMING NEXT

Next, we will use this model to perceive, plan, and learn from experience.

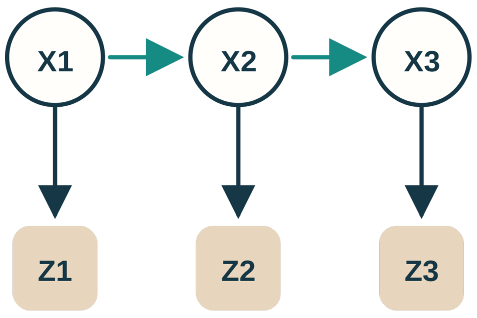

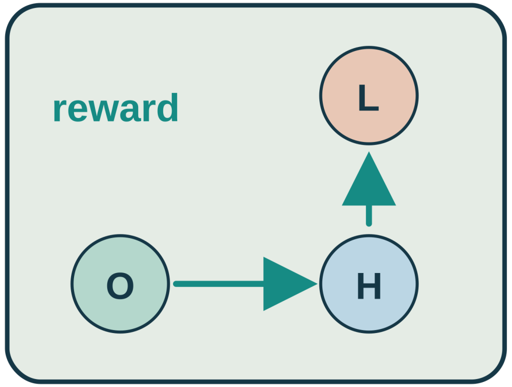

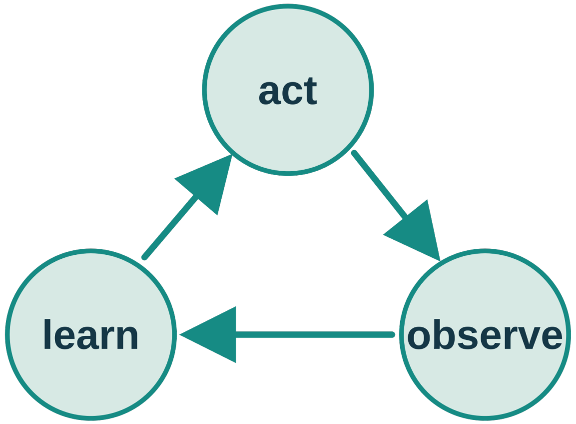

HMMs

Measurements over time

MDPs

Value iteration

Reinforcement learning

Learn by operating

The book's Sections 3.1 and 3.2 explain the state and action models behind this lecture.

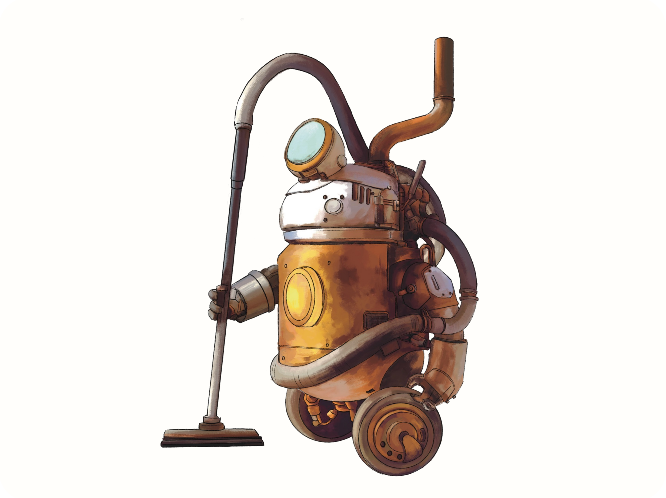

SECTION 3.0

A robot vacuum cleaner

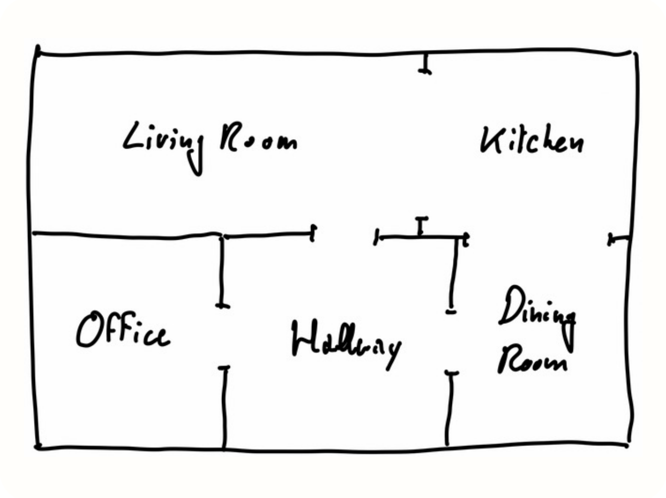

SECTION 3.1

Modeling the state

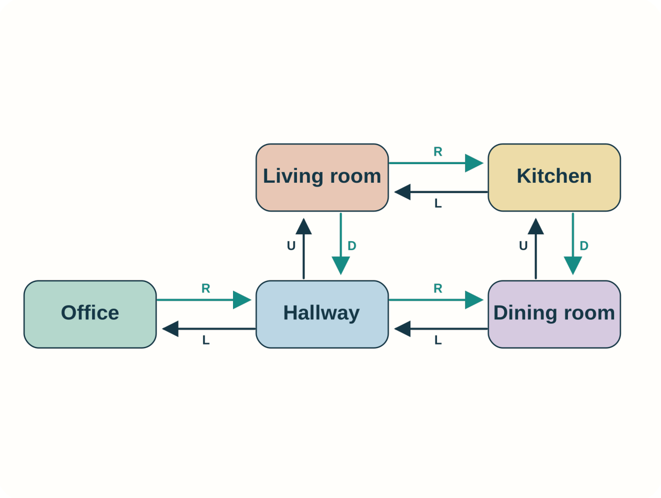

SECTION 3.2

Actions over time

Book: Frank Dellaert & Seth Hutchinson.